SJÄLVSTÄNDIGA ARBETEN I MATEMATIK

MATEMATISKA INSTITUTIONEN, STOCKHOLMS UNIVERSITET

Hurwitz Theory, the ELSV Formula and Beyond

av

Daole Zhang

2026 - No M11

MATEMATISKA INSTITUTIONEN, STOCKHOLMS UNIVERSITET, 106 91 STOCKHOLM

Hurwitz Theory, the ELSV Formula and Beyond

Daole Zhang

Självständigt arbete i matematik 30 högskolepoäng, avancerad nivå

Handledare: Prof. Dan Petersen

2026

Abstract

The ELSV formula, introduced by Ekedahl, Lando, Shapiro, and Vainshtein in 2001, is a foundational result in the study of moduli spaces of curves and Gromov–Witten theory, providing a deep and unexpected connection between combinatorics, algebraic geometry, and theoretical physics. In this thesis, by utilizing the moduli space of relative maps constructed by Jun Li (2001, 2002) alongside virtual localization techniques, a simpler proof of the ELSV formula is obtained. To achieve this, we first develop the necessary mathematical machinery, including the theory of algebraic stacks, equivariant cohomology, and virtual localization. By synthesizing these foundational tools, the thesis provides a rigorous and self-contained exposition of the proof.

Sammanfattning

ELSV-formeln, som introducerades av Ekedahl, Lando, Shapiro och Vainshtein år 2001, är ett grundläggande resultat inom studiet av modulirum för kurvor och Gromov–Witten-teori. Den erbjuder en djup och oväntad koppling mellan kombinatorik, algebraisk geometri och teoretisk fysik. I denna avhandling erhålls ett enklare bevis av ELSV-formeln genom att använda det modulirum för relativa avbildningar som konstruerades av Jun Li (2001, 2002) tillsammans med virtuella lokaliseringstekniker. För att uppnå detta utvecklar vi först den nödvändiga matematiska apparaten, inklusive teorin för algebraiska stackar, ekvivariant kohomologi och virtuell lokalisering. Genom att syntetisera dessa grundläggande verktyg ger avhandlingen en rigorös och självbärande framställning av beviset.

Acknowledgement

First of all, I would like to thank my supervisor, Prof. Dan Petersen, for granting me the freedom to choose the topics of this thesis based on my own interests, and for his valuable suggestions regarding its structure. I also want to thank my co-supervisor, Prof. David Rydh, for reading my draft and providing detailed feedback. I am highly grateful for our discussions on perfect obstruction theory and cotangent complexes, and for his guidance, which allowed me a glimpse into the grand world of derived algebraic geometry.

I would also like to thank my friends in the master's programme. Sincere thanks go to Nami and Julian for teaching me chromatic homotopy theory, even though I understood almost nothing, and for their patience in tolerating my lengthy explanations of algebraic stacks and Chern–Weil forms. I am deeply grateful for the warm and supportive environment in the Department of Mathematics, and I feel incredibly fortunate to have formed so many supportive friendships during these past two years.

Beyond the mathematical community, I am sincerely grateful for the unconditional support of my friends; your care and encouragement have made me who I am today. I want to thank Zhirui, Naiying, and Kajsa for celebrating the New Year, the Lunar New Year, my birthday, and so many other moments with me. A special word of gratitude goes to Kajsa for her continuous support throughout these entire two years. I also thank Yolanda for always sharing her life in New York with me and for keeping in touch so regularly.

I also want to thank Alanni, Emily, and Yolande for their wonderful friendship. A short acknowledgment can hardly express the depth of my gratitude. Thank you to Emily and Yolande for the beautifully prepared Christmas night in Edinburgh, and to Alanni for exploring every bar along Daxue Road in Shanghai with me, as well as for all the fun chats and board games we shared. Most importantly, thank you for our enduring friendship over the past thirteen years.

Finally, I want to thank my family for their unwavering support of the path I have chosen. Without you, I would not be standing here today, looking toward the next chapter of my life. You are the ones who have given me the wings to fly.

Contents

Introduction

Initiated by Gromov's introduction of pseudoholomorphic curves in symplectic geometry, Gromov–Witten theory has experienced a period of explosive growth, establishing profound connections between symplectic geometry, algebraic geometry, and theoretical physics. In the algebraic setting, let $\overline{\mathcal{M}}_{g,n}(X,\beta)$ denote the moduli stack of stable maps, where X is a smooth projective variety over $\mathbb{C}$ and $\beta \in \mathrm{CH}_1(X)$ is a curve class. Given cohomology classes $\gamma_1,\dots,\gamma_n \in \mathrm{CH}^*(X)$, the corresponding Gromov–Witten invariant is defined by the intersection number

$$\langle \gamma_1,\dots,\gamma_n\rangle^X_{g,\beta} \coloneqq \int_{[\overline{\mathcal{M}}_{g,n}(X,\beta)]^{\mathrm{vir}}} \mathrm{ev}_1^*\gamma_1 \cup \cdots \cup \mathrm{ev}_n^*\gamma_n,$$

where $\mathrm{ev}_i\colon \overline{\mathcal{M}}_{g,n}(X,\beta) \to X$ are the canonical evaluation maps, and $[\overline{\mathcal{M}}_{g,n}(X,\beta)]^{\mathrm{vir}}$ is the virtual fundamental class induced by the canonical perfect obstruction theory on the moduli stack.

In general, computing the complete Gromov–Witten theory for an arbitrary space X is an exceptionally difficult problem. Even when X is assumed to be an algebraic curve, the calculations are highly non-trivial, as demonstrated by the foundational trilogy of papers by Okounkov and Pandharipande [**OP06b, OP06a, OP06c**].

However, by applying the powerful machinery of virtual localization, Graber and Pandharipande proved in [**GP99**, Corollary 1] that the full system of Gromov–Witten classes of $\mathbb{P}^r$ lies in the tautological ring of the Deligne–Mumford compactification $\overline{\mathcal{M}}_{g,n}$. This remarkable result implies that one can reduce the Gromov–Witten theory of $\mathbb{P}^r$ to the problem of understanding the full structure of the tautological ring $\mathrm{R}^*(\overline{\mathcal{M}}_{g,n})$.

Regarding the structure of the tautological ring, Faber's conjecture proposed that the tautological ring $\mathrm{R}^*(\overline{\mathcal{M}}_{g,n})$ is Gorenstein. Although this conjecture has since been proven false in general (notably by the counterexamples of Petersen and Tommasi [**PT14**]), it catalyzed a wave of profound research into the intersection theory of $\overline{\mathcal{M}}_{g,n}$. Within this program, top intersection numbers have drawn a great deal of attention. This is because, under the assumption of a perfect pairing, understanding all top intersection numbers would allow one to recover the full structure of the tautological ring. As a corollary, the Gromov–Witten theory of $\mathbb{P}^r$ would be fully understood.

More precisely, one considers the Hodge integrals

$$\langle \tau_{a_1}\dots\tau_{a_n}\lambda_{b_1}\dots\lambda_{b_g}\rangle_g \coloneqq \int_{[\overline{\mathcal{M}}_{g,n}]} \psi_1^{a_1}\cdots\psi_n^{a_n}\cdot\lambda_1^{b_1}\cdots\lambda_g^{b_g} \in \mathbb{Q},$$

where $\sum_{i=1}^n a_i + \sum_{k=1}^g k\cdot b_k = 3g-3+n$. A major goal in the field is to systematically evaluate these top intersection numbers.

The following theorem, known as the ELSV formula, provides a complete solution to this problem.

Theorem 0.0.1 (Ekedahl–Lando–Shapiro–Vainshtein, 2001)**.** *Suppose that $2g-2+n>0$ and $d \in \mathbb{Z}_{\geq 1}$. Let μ be a partition of d of length n. Then for any integer $g \geq 0$,*

$$H_{g,\mu} = \frac{(2g-2+n+d)!}{|\operatorname{Aut}(\mu)|} \prod_{i=1}^{n} \frac{\mu_i^{\mu_i}}{\mu_i!} \int_{[\overline{\mathcal{M}}_{g,n}]} \frac{\Lambda_g^\vee(1)}{\prod_{j=1}^{n}(1-\mu_j\psi_j)}.$$

Here, $H_{g,\mu}$ is the Hurwitz number counting the isomorphism classes of branched coverings $f\colon C \to \mathbb{P}^1$ from a connected curve C of genus g, with prescribed ramification profile μ over ∞ (and simple branching elsewhere), weighted by $|\operatorname{Aut}(f)|^{-1}$. Furthermore, $\Lambda_g^\vee(1) = \sum_{i=0}^{g}(-1)^i\lambda_i$ is the alternating sum of the λ-classes.

This theorem is extraordinarily powerful because it employs the Hurwitz number, a purely combinatorial invariant, to fully determine top intersection numbers on $\overline{\mathcal{M}}_{g,n}$, which otherwise presents a highly abstract geometric problem. Even more surprisingly, Okounkov and Pandharipande [**OP09**] demonstrated that the ELSV formula can be used to give a new proof of Witten's conjecture (Theorem 4.0.3). Originating in topological string theory, Witten's conjecture (first proven by Kontsevich) provided the first systematic method for evaluating intersection numbers of ψ-classes in history. This remarkable proof reveals that the ELSV formula establishes a profound bridge between physics, combinatorics, and algebraic geometry. Furthermore, the natural and deep connection between Hodge integrals and Gromov–Witten theory has been systematically explored by Faber and Pandharipande in their seminal paper [**FP00**].

Following the original announcement of the ELSV formula [**ELSV01**], Graber and Vakil provided an alternative proof utilizing virtual localization on the moduli stack of stable maps, $\overline{\mathcal{M}}_g(\mathbb{P}^1, d)$. This moduli stack naturally parametrizes all branched coverings $C \to \mathbb{P}^1$ of degree d. Adapting the treatment of the Lyashko–Looijenga map from the original ELSV proof, they considered the branch morphism $\operatorname{br}\colon \overline{\mathcal{M}}_g(\mathbb{P}^1, d) \to \operatorname{Sym}^b \mathbb{P}^1$ [**FP02**], where the total branching number $b = 2g-2+2d$ is determined by the Riemann–Hurwitz formula.

However, a significant technical difficulty in this approach is that the chosen absolute moduli stack is excessively large and insensitive to specific ramification profiles. Also, since the target space is rigidly fixed as $\mathbb{P}^1$, limits of stable maps can cause domain components to degenerate in ways that destroy the prescribed ramification over ∞. Consequently, the closure of the locus of maps with the desired ramification profile is ill-behaved within this absolute ambient space.

To circumvent this, Graber and Vakil constructed a closed substack $M \subset \overline{\mathcal{M}}_g(\mathbb{P}^1, d)$ by pulling back a specific linear system from $\operatorname{Sym}^b \mathbb{P}^1$ via the branch morphism, expressing the Hurwitz number as the integral

$$H_{g,\mu} = \int_{[M]} \operatorname{br}^*[p].$$

For a detailed explanation of these notations, we refer to (4.1.4).

Evaluating this right-hand side via virtual localization was a difficult task due to the poorly understood geometry and ill-behaved boundary of M. Acknowledging these difficulties at the end of their paper [**GV03**, Section 5], Graber and Vakil proposed that a simpler proof could be achieved using the theory of relative stable maps. By allowing the target space to degenerate into a chain of rational curves, relative stable maps naturally resolve these boundary issues and preserve the ramification profile at the limits. This approach would set the stage to adapt Graber and Pandharipande's localization calculations from [**GP99**, Section 4].

At the time of Graber and Vakil's writing, the moduli space of relative stable maps had been constructed in the symplectic category [**LR01**], but the foundational algebraic construction was only finalized shortly after by J. Li [**Li01, Li02**]. Armed with Li's algebraic framework, we can now fully realize Graber and Vakil's suggestion. The objective of this thesis is to execute this streamlined proof, bypassing the technical boundary difficulties encountered in the absolute setting.

We also remark that since a virtual localization formula in the symplectic category was subsequently established by Chen and Li in 2006 [**CL06**], it is reasonable that the parallel statement and proof could now be formulated purely within symplectic geometry.

We close the introduction with a brief discussion of the general structure of the thesis.

The first chapter provides an introduction to Hurwitz theory. Our primary objective is to introduce the Hurwitz number $H_{g,d}(\lambda_1, \ldots, \lambda_n)$ and establish its finiteness, showing that it depends exclusively on the topology of the target Riemann surface Y. This topological dependence relies on Riemann's existence theorem, which asserts that any finite unbranched topological covering of a punctured Riemann surface uniquely extends to a holomorphic map between compact Riemann surfaces.

To prove finiteness, the key idea involves translating the geometric problem of counting holomorphic maps into the algebraic problem of counting monodromy representations. This translation is a milestone in Hurwitz theory. It reduces a geometric problem of counting holomorphic maps into a manageable combinatorial problem, which is then further reframed using the powerful machinery of representation theory. The chapter concludes with a discussion of Burnside's character formula, representing the culmination of the classical theory of Hurwitz numbers.

The primary goal of the second chapter is to provide a comprehensive introduction to moduli theory. We begin with the notion of moduli functors, which historically motivated the development of stack theory, as the existence of non-trivial automorphisms obstructs these functors from being representable. Our discussion will adopt both global and local perspectives to illuminate the essential role that automorphisms play in moduli problems. This motivates the enlargement of the category of schemes, leading naturally to the study of pseudo-functors and categories fibered in groupoids, which are bridged by the Grothendieck construction. Utilizing descent theory, we define sheaves and stacks on a site and present several fundamental examples, including quotient stacks, stacks of quasi-coherent sheaves, and the moduli stack of curves of genus g.

To import the geometric tools from classical algebraic geometry, we subsequently define algebraic spaces, Deligne–Mumford stacks, and algebraic stacks—the fundamental geometric spaces of modern moduli theory. We systematically study the geometry of these spaces, aiming toward the proof of the following foundational theorem: for $g \geq 2$, the moduli stack $\mathcal{M}_g$ is a smooth Deligne–Mumford stack of finite type over $\operatorname{Spec} \mathbb{Z}$ of relative dimension $3g - 3$.

Along the way, we establish numerous important consequences in stack theory, such as characterizing Deligne–Mumford stacks and algebraic spaces via their diagonals and stabilizer groups. Following this, we introduce the Deligne–Mumford compactification, establishing that $\overline{\mathcal{M}}_g$ is a smooth, proper, and irreducible Deligne–Mumford stack of dimension $3g - 3$.

In Chapter three, we review equivariant cohomology and classical localization techniques in algebraic geometry, drawing upon [**AF24**] and [**EG98a, EG98b**]. This includes Totaro's celebrated idea of "finite-dimensional approximation," which elegantly sidesteps the issue that topological classifying spaces are typically infinite-dimensional.

We then survey the critical technical machinery of this thesis: virtual fundamental classes and virtual localization. To perform intersection theory in highly degenerate settings, we rely on the virtual fundamental class $[X]^{\mathrm{vir}}$ [**BF97**]. This class is constructed via a perfect obstruction theory $E_X^\bullet$, which is governed by a morphism to the cotangent complex $\mathbb{L}_X$. We conclude by detailing the generalization of classical Atiyah–Bott localization formula to the virtual setting [**GP99**]. This asserts that for an algebraic scheme X with a torus action and a $\mathbb{C}^*$-equivariant perfect obstruction theory, the calculation of equivariant integrals can be reduced to the fixed loci of the group action. More precisely,

$$[X]^{\mathrm{vir}} = \iota_* \sum \frac{[X_i]^{\mathrm{vir}}}{e(N_i^{\mathrm{vir}})} \in \mathrm{CH}_*^{\mathbb{C}^*}(X) \otimes \mathbb{Q}\left[t, \frac{1}{t}\right].$$

This localization formula remains one of our most powerful computational tools. Historically, extending this formula to Deligne–Mumford stacks required the technical assumption that X admits a global embedding into a smooth Deligne–Mumford stack. Since the moduli space of relative stable maps is a Deligne–Mumford stack, this consideration is highly relevant to our subsequent arguments. Fortunately, this restrictive embedding condition was proven to be redundant in [**CKL17**], nearly two decades after Graber and Pandharipande's original work. It is also worthwhile to note that the virtual localization formula has since been generalized even further into the setting of derived algebraic geometry [**AKL$^+$25**].

The final chapter synthesizes the machinery developed in the preceding chapters. We will expand upon the modern tools introduced earlier to execute the proof of the ELSV formula via virtual localization on the moduli space of relative stable maps.

We begin by analyzing the technical difficulties inherent in proving the ELSV formula via virtual localization on the classical moduli space of stable maps $\overline{\mathcal{M}}_g(\mathbb{P}^1, d)$ in details. Following this, we introduce the moduli space of relative stable maps, before proceeding to the explicit virtual localization computations. Although the global embedding condition for virtual localization is now known to be redundant, it is worth noting that Graber and Vakil [**GV05**] explicitly constructed a global embedding of this relative moduli space into a smooth Deligne–Mumford stack. Therefore, one can apply the classical virtual localization formula to this space without needing to rely on later theoretical generalizations.

CHAPTER 1

Hurwitz Theory

Hurwitz numbers were introduced by Adolf Hurwitz in the late nineteenth century. In his celebrated paper [**Hur01**], Hurwitz posed the following problem: given the Riemann sphere $\mathbb{P}^1$ and b distinct points on $\mathbb{P}^1$, determine the number of equivalence classes of degree d branched coverings of $\mathbb{P}^1$ with simple branching over these b points. This problem admits the following generalization:

Question 1.0.1. Let Y be a compact Riemann surface and let $B = \{b_1, \ldots, b_n\} \subset Y$ be a set of distinct points. How many holomorphic maps between compact Riemann surfaces $X \to Y$ of degree d exist such that they have specified ramification profiles over the points in B and are unramified elsewhere?

More precisely, we address the following issues:

(1) Is the number of such maps (up to isomorphism) finite?
(2) Does this number depend on the complex and topological structure of the target Riemann surface Y?
(3) Does this number depend on the specific configuration of the branch points b_i?

This question originated from Riemann's study of multi-valued complex analytic functions, which led to the initial development of Riemann surfaces. Hurwitz transformed this enumerative geometric problem into a combinatorial one involving group homomorphisms, specifically monodromy representations. The key insight is that the ramification profile at the preimages of a certain branch point p corresponds to the cycle type of the permutation associated with a small loop winding around p. Via the representation theory of the symmetric group S_d, this counting problem reduces to an algebraic computation, culminating in Burnside's formula.

In this chapter, we employ the Riemann Existence Theorem to resolve the dependencies discussed above. Specifically, we demonstrate that the count of coverings is determined only by the topological genus of the target surface Y, remaining invariant under deformations of the complex structure of Y or changes in the configuration of the branch points b_i. Following this foundational result, we formally define the Hurwitz number and compute standard examples directly from the definition. We then establish the correspondence between isomorphism classes of Hurwitz coverings and monodromy representations, which provides an affirmative answer to the question regarding the finiteness of these numbers. Finally, we briefly survey how Hurwitz connected this enumerative problem to the representation theory of the symmetric group, culminating in the classical Burnside formula.

We adopt the following conventions throughout this chapter:

- Unless explicitly stated otherwise, all Riemann surfaces are assumed to be **connected**.
- The global degree of a holomorphic map $f\colon X \to Y$ is defined if and only if the map is **proper**. We will implicitly assume this condition whenever discussing the degree. Note that any non-constant holomorphic map between compact Riemann surfaces is automatically proper.

- In this chapter, the term *covering map* refers to a *branched covering map*. Specifically, a map $f\colon X \to Y$ is called a covering if there exists a finite set $B \subset Y$ such that the restriction $f|_{X\setminus f^{-1}(B)}$ is an honest covering map in the standard topological sense. When we wish to refer to a covering with an empty branch locus, we will explicitly use the term *unbranched covering*.

1.1. Riemann's Existence Theorem and Hurwitz Numbers

To address the questions (2) and (3) above, we rely on the Riemann's Existence Theorem, which shows that any unbranched topological covering of a punctured Riemann surface determines a unique holomorphic map of compact Riemann surfaces.

Before proving the main theorem, we introduce a fundamental lemma that serves as a guiding principle in the study of branched coverings. This result establishes a rigorous correspondence between the topological, analytic, and algebraic structures.

Lemma 1.1.1. *Let Y be a compact Riemann surface. There are natural bijections between the following sets:*

$$\left\{\begin{matrix}\textit{Non-constant algebraic morphisms } X \to Y \\ \textit{of smooth projective curves over } \mathbb{C}\end{matrix}\right\} \longleftrightarrow \left\{\begin{matrix}\textit{Non-constant holomorphic maps } X \to Y \\ \textit{of compact Riemann surfaces}\end{matrix}\right\}$$
$$\longleftrightarrow \left\{\begin{matrix}\textit{Topological coverings } X \to Y \\ \textit{of compact topological surfaces}\end{matrix}\right\}.$$

Proof. This result generalizes [**Alp26**, Proposition 6.7.9], which treats the specific case $Y = \mathbb{P}^1$. It is immediate that every algebraic morphism between smooth complex projective curves is holomorphic. Moreover, every non-constant holomorphic map between compact Riemann surfaces exhibits the local behavior $z \mapsto z^k$, which identifies it as a topological covering.

Conversely, let $f\colon X_{\text{top}} \to Y$ be a topological covering, where Y is a fixed compact Riemann surface. We endow X_{top} with a unique complex structure that renders f holomorphic. On the unramified locus, the complex charts of Y are pulled back via the local inverses of f. In the neighborhood of a ramification point $p \in X_{\text{top}}$ with local degree k, we define a coordinate z such that f takes the form $z \mapsto z^k$ relative to a centered local coordinate on Y. This construction defines a compact Riemann surface X such that the map $f\colon X \to Y$ is holomorphic.

Finally, to establish that a holomorphic map $f\colon X \to Y$ between compact Riemann surfaces is algebraic, we recall that every compact Riemann surface admits an embedding into some projective space $\mathbb{P}^N$ by Riemann–Roch and therefore is an algebraic projective variety by Chow's theorem. By Chow's Theorem again, any global holomorphic map between projective algebraic varieties is necessarily algebraic.

Note that this argument relies on the equivalence between the category of compact Riemann surfaces and that of smooth projective algebraic curves. For details, see [**Har77**, Appendix B, Theorem 3.1] and [**GR84**, Theorem 9.5.1]. ☐

Theorem 1.1.2 (Riemann's Existence Theorem)**.** *Let $f^\circ\colon X^\circ \to Y \setminus \{b_1, \ldots, b_n\}$ be an unbranched topological covering of finite degree with X° being a connected topological surface and Y being a compact Riemann surface. Then there exists a unique (up to isomorphism) compact Riemann surface X, such that X° is a dense open subset of X, such that f° extends to $f\colon X \to Y$, a holomorphic map between X and Y.*

Proof. By the bijection established in Lemma 1.1.1, it suffices to complete X° into a compact topological surface and extend f° to a continuous map $f\colon X \to Y$.

Denote $B \coloneqq \{b_1, \dots, b_n\}$ and fix $b \in B$. Consider a coordinate chart φ centered around b and let $\Delta \coloneqq \varphi^{-1}\{|w| < 1\}$, the preimage of the unit disk under the coordinate function φ, so Δ is homeomorphic to an open disk. From the definition of topological covering maps, the map $f^\circ\colon (f^\circ)^{-1}(\Delta \setminus b) \to \Delta \setminus b$ is an unbranched covering map of degree d. Let $U_1^\circ, \dots, U_m^\circ$ be the connected components of the preimage $(f^\circ)^{-1}(\Delta \setminus b)$. Since $\Delta \setminus b$ is homotopically equivalent to the circle S^1, its fundamental group is $\mathbb{Z}$, so each $U_i^\circ, i = 1, \dots, m$ is homeomorphic to a punctured disk. Moreover, there exist positive integers $k_1, \dots, k_m$ and homeomorphism $\psi_i^\circ\colon U_i^\circ \to \{0 < |z| < 1\}$ such that $\varphi \circ f^\circ \circ (\psi_i^\circ)^{-1}$ is described by $z \mapsto z^{k_i}$ for each $i = 1, \dots, m$. Now for each $i = 1, \dots, m$ we add a point p_i to X° in such a way that ψ_i° extends to a homeomorphism $\psi_i\colon U_i^\circ \cup p_i \to \{|z| < 1\}$ such that $\psi_i(p_i) = 0$ and $\psi_i|_{U_i^\circ} = \psi_i^\circ$.

Repeating this process for all points in B, we obtain a new topological space X by adjoining finitely many points to X°. By construction, X is a topological surface, X° is dense in X, and f° extends to a continuous map $f\colon X \to Y$. Moreover, the constructed surface X is compact. This follows from the fact that Y is compact and f is a finite-degree covering, so f is proper and X is compact. □

Remark 1.1.3. The terminology "Riemann's Existence Theorem" is ambiguous, referring to distinct results depending on the context. Classically, it denotes the existence of meromorphic functions on a fixed compact Riemann surface with prescribed singularities. This foundational result, originating in Riemann's seminal papers [**Rie13, Rie57**], is today typically established as a consequence of the Riemann–Roch theorem, see [**Har77**, IV. Ex. 1.1–1.2].

In the context of this chapter (Theorem 1.1.2), the term refers to the converse problem: constructing a complex structure on a topological surface given the data of a branched covering.

Finally, in modern algebraic geometry, the theorem was vastly generalized by Grothendieck to establish an equivalence between the category of finite topological coverings of a complex algebraic variety and the category of finite étale coverings [**SGA03**, Exposé XII, Theorem 5.1].

Theorem 1.1.2 allows us to see that the number of the maps in Question 1.0.1 is independent of the complex structure of Y or of the configuration of the branch points. This leads us to the following definition of Hurwitz numbers. Recall that two holomorphic maps $f\colon X_1 \to Y$ and $g\colon X_2 \to Y$ of Riemann surfaces are isomorphic if there is an isomorphism $\phi\colon X_1 \to X_2$ such that $f = g \circ \phi$.

Definition 1.1.4 (Hurwitz Numbers)**.** Let Y be a compact Riemann surface of genus g, with fixed marked points $b_1, \dots, b_n \in Y$. Let $\lambda_1, \dots, \lambda_n$ be partitions of d. Define the Hurwitz number $H_{g,d}(\lambda_1, \dots, \lambda_n)$ as the weighted count of covers:

$$H_{g,d}(\lambda_1, \dots, \lambda_n) \coloneqq \sum_{[f]} \frac{1}{|\operatorname{Aut}(f)|},$$

where the sum runs over each isomorphism class of $f\colon X \to Y$ where

- f is a holomorphic map of Riemann surfaces;
- X is a compact Riemann surface;
- the branch locus of f is $\{b_1, \dots, b_n\}$;
- the ramification profile of f at each b_i is λ_i.

We call a map f satisfying the above four conditions a Hurwitz cover for the discrete data $g, d, \lambda_1, \dots, \lambda_n$.

Essentially, the Hurwitz number $H_{g,d}(\lambda_1, \ldots, \lambda_n)$ counts the distinct branched coverings matching the given topological data, weighted by the inverse size of their automorphism groups. From Theorem 1.1.2 we see that the Hurwitz number is uniquely determined once the discrete data g, d, λ_i are given. An application of Riemann–Hurwitz theorem ([**Har77**, IV. Corollary 2.4]) also determines the genus of X, which satisfies $2g_X - 2 = d(2g-2) + \big(nd - \sum_{i=1}^n l(\lambda_i)\big)$, where $l(\lambda_i)$ denotes the length of the partition λ_i.

Example 1.1.5 (Hyperelliptic coverings)**.** Recall that a compact Riemann surface X is called hyperelliptic if it admits a holomorphic map $f\colon X \to \mathbb{P}^1$ such that $\deg f = 2$. We also refer to the map f as a hyperelliptic covering. Since $\deg f = 2$, every ramification point in X must have ramification index 2. Consequently, by the Riemann–Hurwitz formula, the genus of X is completely determined by the number of branch points, and conversely.

We first claim that for each $g \in \mathbb{N}_{\geq 2}$, there exists a hyperelliptic curve of genus g. We reduce the problem to the algebraic category via Lemma 1.1.1. Consider a curve X of type $(g+1, 2)$ on the quadric surface $\mathbb{P}^1 \times \mathbb{P}^1$. By the adjunction formula, the genus of X is $(g+1)\cdot 2 - (g+1) - 2 + 1 = g$, see [**Har77**, V. Proposition 1.5 & Example 1.5.2]. Let $\pi_1\colon \mathbb{P}^1 \times \mathbb{P}^1 \to \mathbb{P}^1$ be the projection onto the first factor, and define $f \coloneqq \pi_1|_X$. Since the curve is of type $(g+1, 2)$, its intersection number with the general fiber of π_1 is 2, see [**Har77**, V. Example 1.4.3]. Thus, f is a finite morphism of degree 2 between algebraic curves.

Furthermore, every hyperelliptic curve of genus g admits a canonical algebraic description. The following calculation expands upon [**Har77**, V. Exc 3.5]. Let $a_1, \ldots, a_r$ be distinct complex numbers with $r \geq 5$, and define a plane curve $C \subset \mathbb{P}^2$ (in an affine chart) by the equation $y^2 = \prod_{i=1}^r (x - a_i)$. Let $\widetilde{C}$ be the normalization of C and consider the morphism $f\colon \widetilde{C} \to \mathbb{P}^1$ given by the composition of the normalization map and the projection $[x : y : z] \mapsto [x : z]$. This morphism is well-defined since $[0 : 1 : 0]$ is a singular point of C, and f is a finite morphism of degree 2.

To calculate the genus, we apply the Riemann–Hurwitz formula to see that $2g(\widetilde{C}) - 2 = \deg(f) \cdot (2g(\mathbb{P}^1) - 2) + B = -4 + B$, where B is the total branching index. The affine points $a_1, \ldots, a_r$ are clearly simple branch points. We must also check the behavior at infinity. Substituting $x = 1/t$, the equation becomes $y^2 = t^{-r} \prod_{i=1}^r (1 - a_i t)$.

- If r is even, the function y has a pole of order $r/2$ at $t = 0$. The map has two distinct poles at infinity, so it is unramified there. Therefore, $B = r$.
- If r is odd, the map has a single branch point at infinity with ramification index 2. Therefore, $B = r + 1$.

Solving for the genus, we obtain:

$$g(\widetilde{C}) = \begin{cases} (r-1)/2 & \text{if } r \text{ is odd;} \\ (r/2) - 1 & \text{if } r \text{ is even.} \end{cases}$$

Finally we calculate the automorphism group $\operatorname{Aut}(f)$ of the algebraic morphism $f\colon X \to \mathbb{P}^1$ with $\deg f = 2$. From the canonical algebraic description above it is clear that the only possible non-trivial automorphism is the map $(x, y) \mapsto (x, -y)$, so $\operatorname{Aut}(f) \cong \mathbb{Z}/2\mathbb{Z}$ and $|\operatorname{Aut}(f)| = 2$. To conclude, we obtain that for any $g \geq 2$,

$$H_{0,2}\big((2)^{2g+2}\big) = \frac{1}{2}.$$

1.2. The Monodromy Correspondence

So far, we have not established whether the Hurwitz numbers defined above are finite. In this subsection, we provide an affirmative answer to question (1) by constructing a natural bijection between the isomorphism classes of Hurwitz coverings and a specific class of monodromy representations.

Let $f\colon X \to Y$ be a holomorphic map of degree d between two Riemann surfaces, with branch locus $B = \{b_1, \dots, b_n\} \subset Y$. Fix a base point $y_0 \in Y \setminus B$. By the path lifting and homotopy lifting properties of covering spaces (cf. [**For81**, Chapter 1, Theorems 4.10, 4.14]), there exists a natural group action

$$\begin{aligned} \sigma\colon \pi_1(Y \setminus B, y_0) \times f^{-1}(y_0) &\longrightarrow f^{-1}(y_0) \\ ([\gamma], x) &\longmapsto \widetilde{\gamma}_x(1), \end{aligned}$$

where $\widetilde{\gamma}_x$ denotes the unique path in X obtained by lifting the loop γ via f with initial point $x \in f^{-1}(y_0)$. We refer to this as the **monodromy action** of $\pi_1(Y \setminus B, y_0)$ on the fiber $f^{-1}(y_0)$. Since X is assumed to be connected, this action is transitive.

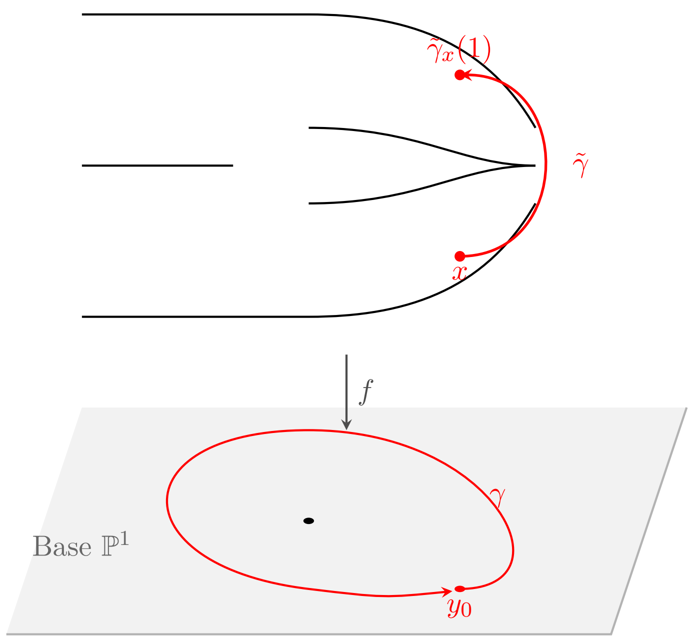

Figure 1. A monodromy action

From the transitivity of this group action, we see that the map $\sigma_\gamma\colon f^{-1}(y_0) \to f^{-1}(y_0)$ induced by a loop γ is a bijection. Therefore, if we identify the fiber $f^{-1}(y_0)$ with the set $[d] \coloneqq \{1, \dots, d\}$, then σ_γ corresponds to a permutation in the symmetric group S_d. This motivates the following definition.

Definition 1.2.1 (y_0-labeled maps)**.** A y_0-labeled map is a pair (f, L), where $f\colon X \to Y$ is a holomorphic map of degree d between two Riemann surfaces, and $L\colon f^{-1}(y_0) \xrightarrow{\sim} [d]$ is a bijection, called a labeling of the fiber over y_0. An isomorphism of y_0-labeled maps (f, L) and (g, L') is a biholomorphism φ between the domains such that $g \circ \varphi = f$ and $L = L' \circ \varphi$ on the fiber over y_0.

A y_0-labeled map (f, L) determines a group homomorphism $\Phi\colon \pi_1(Y \setminus B, y_0) \to S_d$, defined by $\Phi([\gamma]) = L \circ \sigma_\gamma \circ L^{-1}$. We call this homomorphism the **monodromy representation** associated with the y_0-labeled covering.

Crucially, this monodromy representation in fact encodes the data of the ramification profiles. We now illustrate precisely how the local behavior of f near a branch point determines the cycle structure of the monodromy.

Suppose that $f\colon X \to Y$ has ramification profile $\lambda = \{k_1, \ldots, k_m\}$ at a branch point $b \in B$. Let $x_1, \ldots, x_m$ be the points in the fiber $f^{-1}(b)$ with corresponding ramification indices k_j. To analyze the monodromy, we choose a loop γ in $Y \setminus B$ based at y_0 that "encircles b". Specifically, let $\gamma = \alpha \cdot \beta \cdot \alpha^{-1}$, where α is a path connecting y_0 to a point y near b, and β is a simple loop winding once counterclockwise around b inside a small disk D_b.

Inside D_b, the preimage $f^{-1}(D_b)$ consists of a disjoint union of neighborhoods $U_1, \ldots, U_m$, where U_j contains x_j. By the definition of ramification, we can choose local coordinates z on U_j and w on D_b, centered at x_j and b respectively, such that the map is given locally by $w = z^{k_j}$.

Consider the lifting of the loop β parameterized as $w(t) = \varepsilon e^{2\pi i t}$ for $t \in [0, 1]$ into the neighborhood U_j. The equation $z^{k_j} = \varepsilon e^{2\pi i t}$ yields k_j number of distinct solution branches:

$$z_\nu(t) = \varepsilon^{1/k_j} \exp\left(\frac{2\pi i(t + \nu)}{k_j}\right), \quad \nu = 0, \ldots, k_j - 1.$$

As t goes from 0 to 1, the argument increases by $2\pi/k_j$, mapping the sheet ν to $\nu + 1$ (modulo k_j). Consequently, the monodromy permutation σ_γ acts on the k_j points in $U_j \cap f^{-1}(y)$ as a cyclic permutation of length k_j, which is $(1 \to 2 \to \cdots \to k_j \to 1)$.

Since this phenomenon occurs simultaneously and independently in the disjoint neighborhoods $U_1, \ldots, U_m$, the total monodromy permutation σ_γ is the product of m disjoint cycles of lengths $k_1, \ldots, k_m$. Therefore, the cycle type of $\Phi([\gamma])$ corresponds exactly to the partition defined by the ramification profile at b.

In summary, we have established the following result characterizing the algebraic data of Hurwitz coverings.

Proposition 1.2.2. *Let Y be a Riemann surface of genus g with base point y_0 and branch locus $B = \{b_1, \ldots, b_n\} \subset Y$. Let $\lambda_1, \ldots, \lambda_n$ be partitions of a positive integer d.*

If $(f\colon X \to Y, L)$ is a y_0-labeled map of degree d with branch locus B and ramification profile λ_k at each b_k, then it induces a monodromy representation $\Phi\colon \pi_1(Y \setminus B, y_0) \to S_d$ satisfying the following conditions:

(1) Transitivity: The image $\operatorname{Im}\Phi$ acts transitively on the fiber $[d]$;
(2) Local Monodromy: For each $k \in \{1, \ldots, n\}$, if γ_k denotes the homotopy class of a simple positively oriented loop around b_k, the permutation $\Phi(\gamma_k)$ has cycle type λ_k.

We call the monodromy representation satisfying the above two conditions the monodromy representation of type $(g, d, \lambda_1, \ldots, \lambda_n)$.

Next, we address the converse problem: given a monodromy representation of type $(g, d, \lambda_1, \ldots, \lambda_n)$, does it uniquely determine a Hurwitz covering $f\colon X \to Y$? The answer is affirmative.

The strategy for the construction is rooted in the "gluing" of local sheets. Let $U := Y \setminus B$ be the dense open subset representing the unramified locus. Over this region, any degree d covering locally resembles d disjoint copies of U. The monodromy representation provides the precise "gluing instructions": it records how these d copies are permuted as one traverses loops in U.

This construction yields an unbranched topological covering $p\colon X^\circ \to U$. By Riemann's Existence Theorem 1.1.2, this covering extends uniquely to a holomorphic map $f\colon X \to Y$ between compact Riemann surfaces.

This motivates the use of the Galois correspondence for covering spaces to give a formal proof. While the geometric intuition of gluing sheets is vivid, the rigorous machinery of the fundamental group allows for a cleaner argument. Recall that for a connected manifold M, there is a bijection between the connected covering spaces of M and the subgroups of $\pi_1(M, y_0)$. Specifically, a subgroup $H \subseteq \pi_1(M, y_0)$ corresponds to the covering space $\widetilde{M}/H$, where $\widetilde{M}$ is the universal covering. For details see [**For81**, Chapter 1.5]. In our context, the monodromy representation determines a specific subgroup (stabilizer), which in turn determines the covering space X° as a quotient of the universal cover of $Y \setminus B$.

Proposition 1.2.3. *Let Y be a compact Riemann surface, $B \subset Y$ a finite set, and $y_0 \in U \coloneqq Y \setminus B$. Let $\Phi\colon \pi_1(U, y_0) \to S_d$ be a monodromy representation of type $(g, d, \lambda_1, \dots, \lambda_n)$. Then there exists a unique y_0-labeled map $(f\colon X \to Y, L)$ of degree d (up to labeled isomorphism) such that its associated monodromy representation is exactly Φ.*

Proof. The proof expands upon the ideas in [**Mir95**, pp. 88–89]. For a geometry-oriented proof using a gluing construction we refer to [**CM16**, Theorem 7.2.2].

Let $H \subset \pi_1(U, y_0)$ be the stabilizer of the index $1 \in [d]$ under the action defined by Φ:

$$H \coloneqq \{[\gamma] \in \pi_1(U, y_0) \mid \Phi([\gamma])(1) = 1\}.$$

Since Φ is of type $(g, d, \lambda_1, \dots, \lambda_n)$, its image is a transitive subgroup of S_d, so the index of H in $\pi_1(U, y_0)$ is equal to the size of the orbit of 1, which is d.

By the Galois correspondence for covering spaces, there exists a unique unbranched covering map $p^\circ\colon X^\circ \to U$ with a base point $x_0 \in (p^\circ)^{-1}(y_0)$ such that the induced homomorphism $(p^\circ)_*$ maps $\pi_1(X^\circ, x_0)$ isomorphically onto H. Since $[\pi_1(U, y_0) : H] = d$, this covering has degree d.

We define the labeling $L\colon (p^\circ)^{-1}(y_0) \to [d]$ as follows. For any point x in the fiber over y_0, choose a path $\tilde{\gamma}$ in X° from x_0 to x. Let $\gamma \coloneqq p^\circ \circ \tilde{\gamma}$ be the projection of this path to U. Define

$$L(x) \coloneqq \Phi([\gamma])(1).$$

To see that L is well-defined, suppose $\tilde{\eta}$ is another path from x_0 to x. Then $\tilde{\gamma} \cdot \tilde{\eta}^{-1}$ is a loop in X° based at x_0, so its projection $\gamma \cdot \eta^{-1}$ lies in the subgroup $(p^\circ)_*(\pi_1(X^\circ, x_0)) = H$. By the definition of H, $\Phi([\gamma \cdot \eta^{-1}])(1) = 1$, which implies $\Phi([\gamma])(1) = \Phi([\eta])(1)$. Thus $L(x)$ is independent of the choice of path.

Moreover, since the action of Φ is transitive, for any $k \in [d]$, there exists γ such that $\Phi([\gamma])(1) = k$. Lifting this γ to the path starting at x_0 yields a point x with $L(x) = k$. Thus L is surjective and, since the sets are finite of size d, a bijection.

By Riemann's Existence Theorem 1.1.2, the unbranched covering p° extends uniquely to a holomorphic map $f\colon X \to Y$ between compact Riemann surfaces, where X contains X° as a dense open subset. We now verify that the monodromy representation Ψ of the constructed labeled map (f, L) is exactly Φ. Let $\rho \in \pi_1(U, y_0)$. Under the monodromy action, ρ lifts to a path $\tilde{\rho}$ starting at $x \in f^{-1}(y_0)$ and ending at a point x'. If x is defined by a path $\tilde{\gamma}$ starting at x_0, i.e., $L(x) = \Phi([\gamma])(1)$, then x' is reached by the path $\tilde{\gamma} \cdot \tilde{\rho}$. Therefore

$$\Psi(\rho)(L(x)) = L(x') = \Phi([\gamma \cdot \rho])(1) = \Phi([\rho])(\Phi([\gamma])(1)) = \Phi([\rho])(L(x)).$$

Since this holds for all x, $\Psi = \Phi$. Finally, since the monodromy matches Φ, and Φ is of type $(g, d, \lambda_1, \dots, \lambda_n)$, the local monodromy around each branch point b_k has cycle type λ_k. As established previously, the cycle type of the local monodromy corresponds exactly to the ramification profile. Thus, f is a Hurwitz covering of the required type. $\square$

Combining Proposition 1.2.2 with Proposition 1.2.3 yields the bijection between isomorphism classes of y_0-labeled maps $X \to Y$ of type $(g, d, \lambda_1, \dots, \lambda_n)$ and monodromy representations $\pi_1(Y \setminus B, y_0) \to S_d$ of type $(g, d, \lambda_1, \dots, \lambda_n)$.

Theorem 1.2.4. *There is a bijection between the following two sets:*

$$\left\{ \begin{array}{l} \textit{Isomorphism classes of } (f\colon X \to Y, L) \\ y_0\textit{-labeled maps of compact Riemann surfaces,} \\ f \textit{ branches over } B = \{b_1, \dots, b_n\} \\ \textit{with ramification profile } \lambda_i \textit{ over } b_i. \end{array} \right\} \longleftrightarrow \left\{ \begin{array}{l} \textit{Monodromy representations} \\ \rho\colon \pi_1(Y \setminus B, y_0) \to S_d \\ \textit{of type } (g, d, \lambda_1, \dots, \lambda_n). \end{array} \right\}$$

As a corollary, for any given g, d and partitions $\lambda_1, \dots, \lambda_n$, the Hurwitz numbers $H_{g,d}(\lambda_1, \dots, \lambda_n)$ are always finite, since the fundamental group $\pi_1(Y \setminus B, y_0)$ is finitely generated when B is a finite set. This answers question (1). More precisely:

Corollary 1.2.5. *Let Y be a compact Riemann surface and let $B = \{b_1, \dots, b_n\} \subset Y$, and M be the set of monodromy representations $\pi_1(Y \setminus B, y_0) \to S_d$ of type $(g, d, \lambda_1, \dots, \lambda_n)$. Then*

$$H_{g,d}(\lambda_1, \dots, \lambda_n) = \frac{|M|}{d!}.$$

Proof. Given a Hurwitz covering $f\colon X \to Y$ of type $(g, d, \lambda_1, \dots, \lambda_n)$, there are $d!$ ways to label the fiber $f^{-1}(y_0)$. On the other hand, an automorphism $\varphi \in \operatorname{Aut}(f)$ induces an isomorphism of y_0-labeled maps $(f, L) \cong (f, L \circ \varphi^{-1})$. Since $L \circ \varphi^{-1} = L$ if and only if φ is trivial, these two facts imply that the number of isomorphism classes of y_0-labeled maps f is $d!/|\operatorname{Aut}(f)|$. Therefore,

$$H_{g,d}(\lambda_1, \dots, \lambda_n) = \sum_{[f]} \frac{1}{|\operatorname{Aut}(f)|} = \frac{1}{d!} \sum_{[f]} \frac{d!}{|\operatorname{Aut}(f)|} = \frac{|M|}{d!},$$

where the final equality comes from Theorem 1.2.4. $\square$

Example 1.2.6 (Hyperelliptic Covering Revisit)**.** We recalculate the Hurwitz number $H_{0,2}\big((2)^{2g+2}\big)$ utilizing the correspondence established in Theorem 1.2.4. It suffices to enumerate the monodromy representations $\Phi\colon \pi_1(\mathbb{P}^1 \setminus B, y_0) \to S_2$ of type $(0, 2, (2)^{2g+2})$.

The fundamental group $\pi_1(\mathbb{P}^1 \setminus B, y_0)$ is generated by loops $\gamma_1, \dots, \gamma_{2g+2}$ winding around the branched points, subject to the single relation $\prod_{i=1}^{2g+2} \gamma_i = e$. A representation of the specified type requires the image of each generator $\Phi(\gamma_i)$ to be a transposition. In the symmetric group S_2, the unique transposition is $\tau = (12)$. Therefore, we are forced to let $\Phi(\gamma_i) = \tau$ for each $i = 1, \dots, 2g+2$, and $\prod_{i=1}^{2g+2} \Phi(\gamma_i) = e$ since $2g+2$ is even. Consequently, there is exactly one valid monodromy representation, so $|M| = 1$. This yields

$$H_{0,2}\big((2)^{2g+2}\big) = \frac{1}{2}.$$

Compared with the geometric calculation in Example 1.1.5, this combinatorial calculation is significantly simpler.

Example 1.2.7. In this example, we calculate the Hurwitz number $H_{0,4}((3,1),(2,2)^2)$. First, we verify the genus using the Riemann–Hurwitz formula. The total branching index is $B = (3-1)+2\cdot(2-1+2-1) = 2+4 = 6$. Therefore, the genus of the source curve X satisfies:

$$2g_X - 2 = 4(2\cdot 0 - 2) + 6 = -8 + 6 = -2 \implies g_X = 0.$$

We now enumerate the monodromy representations of type $(0,4,(3,1),(2,2)^2)$. Let $\gamma_1, \gamma_2, \gamma_3$ be the generating loops winding around the three branch points, subject to the single relation $\gamma_1\gamma_2\gamma_3 = e$. A representation of the specified type requires that $\Phi(\gamma_1)$ is a 3-cycle and that $\Phi(\gamma_2), \Phi(\gamma_3)$ are products of two disjoint transpositions.

Let $K_4 = \{e, (12)(34), (13)(24), (14)(23)\}$ be the Klein group. Recall that K_4 consists exactly of the identity and the conjugacy class of type $(2,2)$, and that K_4 is a normal subgroup of S_4. Consider the canonical quotient homomorphism $\pi\colon S_4 \to S_4/K_4 \cong S_3$. We immediately obtain a contradiction that

$$e_{S_3} = \pi(\Phi(\gamma_1\gamma_2\gamma_3)) = \pi(\Phi(\gamma_1)) \neq e_{S_3}.$$

Consequently, no such monodromy representation exists. We conclude that

$$H_{0,4}((3,1),(2,2)^2) = 0.$$

Example 1.2.7 demonstrates that the Riemann–Hurwitz formula, while necessary, is not sufficient for the existence of a Hurwitz covering. This observation motivates the Hurwitz existence problem: the quest to determine necessary and sufficient conditions for a Hurwitz number to be non-zero. Despite extensive study, this problem remains open. The case where the target curve has genus $g > 0$ was fully resolved by Edmonds, Kulkarni, and Stong in [**EKS84**]. However, for the case of $\mathbb{P}^1$, a complete solution is still missing, although numerous partial results exist. In the same work [**EKS84**], it is conjectured that for prime degrees d, the Riemann–Hurwitz condition is indeed necessary and sufficient. Recent computational advances [**WLZ$^+$25**] have verified this conjecture for all degrees $d \leq 31$.

1.3. The Burnside Character Formula

For the sake of completeness, we conclude this chapter with a brief survey of how Hurwitz further reformulated the enumeration of covering maps using the representation theory of the symmetric group S_d. As these results are not required for subsequent chapters, we omit detailed proofs and refer the interested reader to [**CM16**, Chapter 9].

Let $\mathbb{C}[S_d]$ denote the group algebra of the symmetric group S_d, and let $\mathcal{Z}\mathbb{C}[S_d]$ denote its center. Recall that the conjugacy class sums form a natural basis for the center. Specifically, for any partition $\lambda \vdash d$, let $C_\lambda \in \mathbb{C}[S_d]$ be the sum of all permutations with cycle type λ. As a vector space, the center admits the decomposition

$$\mathcal{Z}\mathbb{C}[S_d] = \bigoplus_{\lambda \vdash d} \mathbb{C}C_\lambda.$$

We introduce the disconnected Hurwitz number $H^\bullet_{g,d}(\lambda_1, \ldots, \lambda_n)$, which counts Hurwitz coverings of a specified type without imposing the condition that the source curve is connected. Correspondingly, we define a disconnected monodromy representation by dropping the requirement that the image of the homomorphism is a transitive subgroup of S_d.

With these preparations, we consider the genus zero case $H^\bullet_{0,d}(\lambda_1, \ldots, \lambda_n)$, where the target curve is $\mathbb{P}^1$. By the correspondence in Theorem 1.2.4, enumerating these coverings is equivalent to counting disconnected monodromy representations $\Phi\colon \pi_1(\mathbb{P}^1 \setminus B, y_0) \to S_d$. The fundamental group is generated by simple loops γ_i around the branch points b_i $(i = 1, \dots, n)$, subject to the relation $\gamma_1 \cdots \gamma_n = e$. Thus, such a representation is determined by a choice of permutations $\sigma_1, \dots, \sigma_n \in S_d$ such that each σ_i has cycle type λ_i and their product satisfies $\sigma_n \cdots \sigma_1 = e$.

We can translate this counting problem directly into the group algebra. Consider the product of the class sums $C_{\lambda_n} \cdots C_{\lambda_1}$ in $\mathcal{Z}\mathbb{C}[S_d]$. Expanding this product linearly yields

$$C_{\lambda_n} \cdots C_{\lambda_1} = \left(\sum_{\tau_n \in C_{\lambda_n}} \tau_n \right) \cdots \left(\sum_{\tau_1 \in C_{\lambda_1}} \tau_1 \right) = \sum_{\tau_i \in C_{\lambda_i}} (\tau_n \cdots \tau_1).$$

A tuple $(\sigma_1, \dots, \sigma_n)$ constitutes a valid monodromy representation if and only if the product $\sigma_n \cdots \sigma_1$ equals the identity e. Therefore, the number of such representations is exactly the coefficient of the identity element in the expansion of the product $C_{\lambda_n} \cdots C_{\lambda_1}$.

Let $[e]\mathbf{x}$ denote the coefficient of the identity element e in an element $\mathbf{x} \in \mathbb{C}[S_d]$. We arrive at the elegant formula

$$H^{\bullet}_{0,d}(\lambda_1, \dots, \lambda_n) = \frac{1}{d!}[e]\left(C_{\lambda_n} \cdots C_{\lambda_1}\right).$$

When the target curve Y has higher genus $g > 0$, the logic of the calculation remains largely the same, though the structure of the fundamental group becomes more complicated. The group $\pi_1(Y \setminus B, y_0)$ is generated by the n simple loops γ_i around the branch points, plus an additional $2g$ "handle loops" $\alpha_1, \beta_1, \dots, \alpha_g, \beta_g$. These $n + 2g$ generators are subject to the single relation

$$\gamma_n \cdots \gamma_1 [\alpha_1, \beta_1] \cdots [\alpha_g, \beta_g] = e,$$

where $[\alpha, \beta] \coloneqq \alpha\beta\alpha^{-1}\beta^{-1}$ denotes the commutator.

To enumerate the corresponding monodromy representations, we must account for the contribution of these handles. Let $\mathfrak{K} \in \mathcal{Z}\mathbb{C}[S_d]$ be the specific element defined by

$$\mathfrak{K} \coloneqq \sum_{\lambda \vdash d} |\mathcal{Z}(\sigma_\lambda)| C_\lambda,$$

where σ_λ is any representative permutation of cycle type λ, and $|\mathcal{Z}(\sigma_\lambda)|$ is the size of its centralizer. This element $\mathfrak{K}$ effectively encodes the sum of all commutators in the group algebra. A similar combinatorial expansion then yields the formula

$$H^{\bullet}_{g,d}(\lambda_1, \dots, \lambda_n) = \frac{1}{d!}[e]\left(\mathfrak{K}^g \cdot C_{\lambda_n} \cdots C_{\lambda_1}\right).$$

Starting from this expression, one can apply the change of basis formula implied by Maschke's theorem, which decomposes the semisimple algebra $\mathcal{Z}\mathbb{C}[S_d]$ into a sum of matrix algebras indexed by irreducible representations. This leads directly to the celebrated result of Burnside:

Theorem 1.3.1 (Burnside Character Formula)**.** *Fix a positive integer d and partitions $\lambda_1, \dots, \lambda_n$ of d. Let $\mathrm{Irr}(S_d)$ denote the set of irreducible representations of S_d. Then the disconnected Hurwitz number is given by*

$$H^{\bullet}_{g,d}(\lambda_1, \dots, \lambda_n) = \sum_{\rho \in \mathrm{Irr}(S_d)} \left(\frac{\dim \rho}{d!} \right)^{2-2g} \prod_{j=1}^{n} \frac{|C_{\lambda_j}| \chi_\rho(\lambda_j)}{\dim \rho},$$

where χ_ρ is the character associated with the representation ρ, and $|C_{\lambda_j}|$ denotes the size of the conjugacy class in S_d corresponding to λ_j.

CHAPTER 2

An Introduction to Moduli Problems

The term "moduli" traces its origins to Bernhard Riemann. In his 1857 paper on Abelian functions [**Rie57**], Riemann introduced the concept of the "number of moduli" and deduced that this number is $3g-3$ for a curve of genus $g \geq 2$. He termed these parameters "moduli" (from the Latin *modulus*, meaning a small measure). However, the problem of constructing a geometric space whose points correspond naturally to these geometric objects remained open for decades. For instance, in the context of Riemann surfaces, does there exist a complex analytic space M_g whose points are in one-to-one correspondence with Riemann surfaces of genus g? This question was studied extensively in the late 19th and early 20th centuries, attracting outstanding mathematicians including Hilbert, Chow, Kodaira, and Spencer. In the 1940s, Oswald Teichmüller constructed what is now known as the Teichmüller space T_g, parameterizing complex structures on a topological surface of genus g. He showed that T_g is homeomorphic to a ball in $\mathbb{C}^{3g-3}$ and realized $M_g = T_g/\Gamma_g$ as the quotient of T_g by the action of the mapping class group Γ_g. See [**IT92**] for an introduction and [**Alp25**] for a comprehensive history of moduli problems.

Building on this pioneering work, the field underwent a paradigm shift when Alexander Grothendieck introduced a new perspective. Rather than viewing a space X as a set with topology and additional structure, Grothendieck proposed understanding a space X via its relationship to all other spaces. This philosophy can be stated rigorously in the language of functors and justified by Yoneda's lemma, which states that a scheme X is determined up to isomorphism by its functor of points $h_X \coloneqq \mathrm{Hom}(-, X)$. This perspective leads to the modern definition of the moduli functor of smooth algebraic curves.

Definition 2.0.1 (Moduli functor of smooth curves of genus g)**.**

(1) A family of smooth curves of genus g is a smooth, proper morphism $\mathcal{C} \to S$ of schemes such that for every $s \in S$, the fiber $\mathcal{C}_s$ is a connected smooth curve of genus g.

(2) Two families of smooth curves of genus g, denoted $\mathcal{C}_i \to S$, $i = 1, 2$, are isomorphic if there exists an isomorphism $\varphi\colon \mathcal{C}_1 \xrightarrow{\cong} \mathcal{C}_2$ such that the diagram

$$\begin{tikzcd} \mathcal{C}_1 \arrow[rr, "\varphi"] \arrow[dr] & & \mathcal{C}_2 \arrow[dl] \\ & S & \end{tikzcd}$$

commutes.

(3) The moduli functor of smooth curves of genus g is the functor

$$F_g\colon \mathsf{Sch}^{\mathrm{op}} \to \mathsf{Sets},$$

where for each $S \in \mathsf{Sch}$, $F_g(S)$ is the set of isomorphism classes of families of smooth curves of genus g over S. Moreover, given a morphism of schemes $S_1 \to S_2$, the induced map $F_g(S_2) \to F_g(S_1)$ is given by pulling back the family $\mathcal{C}_2 \to S_2$ along the base change $S_1 \to S_2$.

In the definition above, the requirement that the family be a smooth morphism implies that it is flat. This flatness ensures the "continuous" variation of the curves along the parameterizing scheme S. The power of this functorial language becomes especially clear if the functor F_g is represented by a scheme M, known as a fine moduli space. In such a case, specifying a family over a base B is equivalent to specifying a morphism $f\colon B \to M$, where $f(b)$ corresponds to the fiber over b.

We can generalize this notion from smooth curves to other moduli problems. Using the category of schemes Sch as a framework, a general moduli problem consists of two components: first, a class of objects equipped with a notion of a *family* over a base scheme B; and second, an equivalence relation between these families. A moduli problem is thus encoded by a functor

$$\mathcal{M}\colon \mathsf{Sch}^{\mathrm{op}} \to \mathsf{Sets}, \qquad T \longmapsto \mathcal{M}(T) = \{\text{families of objects over } T\}/\sim,$$

where for a morphism $T_1 \to T_2$, the induced map $\mathcal{M}(T_2) \to \mathcal{M}(T_1)$ is the pullback along the base change:

$$\begin{tikzcd} C_1 \arrow[r] \arrow[d] \arrow[dr, phantom, "\square"] & C_2 \arrow[d] \\ T_1 \arrow[r] & T_2. \end{tikzcd}$$

This chapter is structured as follows. We begin by examining examples where moduli functors fail to be representable, tracing the root cause to the presence of non-trivial automorphisms. We present two complementary perspectives on how these automorphisms explicitly obstruct representability. Motivated by these difficulties, we introduce the language of pseudo-functors and fibered categories, which leads naturally to the definition of stacks over a site equipped with a Grothendieck topology. To perform concrete geometry, we require the concept of an atlas to bridge the gap between schemes and stacks. This yields the definitions of algebraic spaces, Deligne–Mumford stacks, and algebraic stacks. We then discuss how classical scheme-theoretic properties extend to this setting, analyzing the behavior of diagonal morphisms to establish the characterizations of algebraic spaces and Deligne–Mumford stacks. With this general framework in place, we investigate the geometric properties of the moduli stack $\mathcal{M}_g$, the central object of this thesis. We subsequently explore the geometry of algebraic stacks, extending the notions of quasi-coherent sheaves, cohomology, and intersection theory. In particular, following [**Vis89**], we outline the generalization of classical intersection theory [**Ful98**] to the setting of Deligne–Mumford stacks. Finally, we conclude the chapter by introducing the Deligne–Mumford compactification $\overline{\mathcal{M}}_{g,n}$ and defining its tautological classes.

2.1. Automorphisms and the Obstruction to Representability

Yoneda's lemma is frequently described as "the hardest trivial thing in mathematics" due to its profound utility across diverse mathematical branches despite its elementary proof. One of the most significant corollaries of Yoneda's lemma is the existence of a *universal object* for representable functors. Specifically, if a functor $F\colon \mathsf{Sch}^{\mathrm{op}} \to \mathsf{Sets}$ is represented by a scheme M, there exists a universal element $\mathcal{U} \in F(M)$ such that for any scheme S, every object in $F(S)$ is uniquely obtained by pullback from $\mathcal{U}$ via a morphism $g\colon S \to M$. Consequently, if a moduli functor is representable by a scheme M, the study of all families of the specified objects reduces to the study of the geometry of M and the single universal family $\mathcal{U} \to M$. Philosophically, this significantly simplifies the classification problem.

Unfortunately, many natural moduli functors are far from being representable, and the essence of this failure lies in the existence of non-trivial automorphisms. We begin with an analysis of the global moduli functor itself, examining how an isotrivial family obstructs the representability of the global functor. Following this, we shift our focus from global moduli functors to "local moduli functors," specifically the deformation functor Def_X of a given scheme X. A deeper analysis of the deformation functor will further illuminate the role of automorphisms.

2.1.1. Global Obstructions: Isotrivial Families. We begin with a fundamental example to illustrate why some global moduli functors fail to be representable.

Example 2.1.1. Over a field k of characteristic zero, consider the moduli functor $F\colon \mathsf{Sch}^{\mathrm{op}} \to \mathsf{Sets}$ which assigns to each scheme S the set of isomorphism classes of vector bundles of rank n on S. We claim that F is not representable. To see this, suppose that F were represented by a scheme M. The set $F(\operatorname{Spec} k)$ consists of a single element: the isomorphism class of the trivial bundle k^n. This implies that the moduli space M has only one k-point. Consequently, for any reduced, connected scheme S, the only morphism $S \to M$ is the constant map.

On the other hand, for any scheme S, a vector bundle $\mathcal{E}$ on S corresponds to a morphism $f\colon S \to M$ such that $\mathcal{E} \cong f^*\mathcal{U}$, where $\mathcal{U}$ is the universal bundle on M. However, since f must be constant, it factors through the point $\operatorname{Spec} k \to M$. This implies that every vector bundle on S must be trivial. This is a contradiction, as non-trivial vector bundles exist (e.g., on $\mathbb{P}^1$).

The example above is fundamental because it reveals the mechanism of the failure. For a general moduli functor, we have the following proposition showing that automorphisms obstruct representability via the existence of isotrivial families, which are families that are trivial fiberwise but not globally.

Proposition 2.1.2. *Over a field k of characteristic zero, let $F\colon \mathsf{Sch}^{\mathrm{op}} \to \mathsf{Sets}$ be a moduli functor. Suppose there exists a family of objects $\mathcal{E} \in F(S)$ over a variety S such that:*

(1) The fibers $\mathcal{E}_s$ are isomorphic for all $s \in S(k)$; and
(2) The family $\mathcal{E}$ itself is not trivial, i.e., it is not isomorphic to the pullback of an object $E \in F(k)$ along the structure map $S \to \operatorname{Spec} k$.

Then F is not representable.

Proof. We proceed by contradiction. Suppose that F is represented by a scheme X over k. By the first condition, the isomorphism class of the fiber $E := \mathcal{E}_s$ is independent of $s \in S(k)$. By Yoneda's lemma, this class corresponds to a unique point $x \in X(k)$. Consequently, the classifying map $f\colon S \to X$ associated with $\mathcal{E}$ maps every k-point of S to x, so the map f factors through $\operatorname{Spec} k$:

$$\begin{tikzcd} S \arrow[rr, "f"] \arrow[dr] & & X \\ & \operatorname{Spec} k \arrow[ur, "x"'] & \end{tikzcd}$$

Here $S \to \operatorname{Spec} k$ is the structure map of S. This factorization implies that $\mathcal{E} \cong f^*\mathcal{U}_{\mathrm{univ}}$ is the pullback of the object over the point x, meaning $\mathcal{E}$ is a trivial family. This contradicts the second condition. □

Example 2.1.3. Consider the moduli functor of elliptic curves, denoted by $\mathcal{M}_{1,1}$. Let $k = \mathbb{C}$. Consider the family of elliptic curves $\mathcal{E} \to S = \mathbb{A}^1_k \setminus \{0\}$ defined by the equation:

$$\begin{tikzcd} \mathcal{E} = V(y^2z - x^3 + tz^3) \arrow[r, hook] \arrow[d] & S \times \mathbb{P}^2 \\ S = \mathbb{A}^1_k \setminus \{0\} & \end{tikzcd}$$

We verify that the fibers are isomorphic. For any closed point $t_0 \in S$, the fiber E_{t_0} is given by $y^2z = x^3 - t_0z^3$. The change of variables $x \mapsto \sqrt[3]{t_0}x'$ and $y \mapsto \sqrt{t_0}y'$ transforms the equation into $(y')^2z = (x')^3 - z^3$. Thus, every fiber is isomorphic to the standard curve $E_0\colon y^2z = x^3 - z^3$. This satisfies the first condition of Proposition 2.1.2.

Second, we show that the family $\mathcal{E}$ itself is not trivial over S, i.e., it is not isomorphic to the constant family $\mathcal{E}_{\text{triv}} = E_0 \times S$. To see this, we utilize the modular discriminant Δ. For a Weierstrass equation $y^2z = x^3 + Axz^2 + Bz^3$, the discriminant is given by $\Delta = -16(4A^3 + 27B^2)$. An isomorphism between two such families over S corresponds to a coordinate change of the form $x = u^2x', y = u^3y'$ for some unit $u \in \Gamma(S, \mathcal{O}_S^\times)$. Under this transformation, the coefficients scale as $A = u^4A'$ and $B = u^6B'$, which implies the discriminant transforms as

$$\Delta = -16(4(u^4A')^3 + 27(u^6B')^2) = u^{12} \cdot [-16(4(A')^3 + 27(B')^2)] = u^{12}\Delta'.$$

We now compute the discriminants for our two families:

- For the constant family $\mathcal{E}_{\text{triv}}$, we have $A = 0, B = -1$, so $\Delta_{\text{triv}} = -16(27(-1)^2) = -432$.
- For the twisted family $\mathcal{E}$, we have $A = 0, B = -t$, so $\Delta_{\mathcal{E}} = -16(27(-t)^2) = -432t^2$.

If an isomorphism existed, there must be a unit $u \in k[t, t^{-1}]^\times$ such that $\Delta_{\mathcal{E}} = u^{12}\Delta_{\text{triv}}$. Substitute our values and we see that $t^2 = u^{12}$. The units in $k[t, t^{-1}]$ are of the form ct^n for $c \in k^\times, n \in \mathbb{Z}$. Thus we require $t^2 = (ct^n)^{12} = c^{12}t^{12n}$. Comparing the exponents of t, we get $n = 1/6$. Since there is no integer solution for n, no such unit u exists.

This contradiction proves that the family $\mathcal{E}$ is not trivial over S. Since the family satisfies both conditions of Proposition 2.1.2, $\mathcal{M}_{1,1}$ is not a scheme.

2.1.2. Local Obstructions: Prorepresentability of the Deformation Functor. The examples discussed in the previous subsection, particularly Proposition 2.1.2, demonstrate that global automorphisms, manifesting as isotrivial families, prevent the global moduli functor from being representable. In this subsection, we shift our focus to the local, and specifically infinitesimal, structure of the moduli problem.

If global representability is obstructed, one might hope for a weaker condition: is the moduli functor at least "locally representable"? In other words, can we uniquely describe the infinitesimal deformations of a single fixed object X by a complete local ring? This question leads us to the study of the deformation functor Def_X. As we will see, the answer remains intimately tied to automorphisms. Just as global automorphisms obstruct the existence of a global universal family, infinitesimal automorphisms obstruct the pro-representability of the local deformation functor.

To formalize our discussion, we introduce the fundamental notions of deformation theory, adopting the notation from [**Ser06**] and [**FM98**].

Let X be an algebraic scheme over an algebraically closed field k (i.e., a separated, locally Noetherian scheme over k). A deformation of X parameterized by a connected scheme S with a distinguished k-point $s \in S(k)$ consists of a flat, surjective morphism $\pi\colon \mathcal{X} \to S$ together with an isomorphism $i\colon X \xrightarrow{\sim} \mathcal{X}_s$ fitting into the following Cartesian diagram:

$$\begin{tikzcd} X \arrow[r, "i"] \arrow[d] \arrow[dr, phantom, "\square"] & \mathcal{X} \arrow[d, "\pi"] \\ \operatorname{Spec} k \arrow[r, "s"] & S \end{tikzcd}$$

We often abuse notation by referring to the total space $\mathcal{X}$ or the map π as the deformation, suppressing the specific identification i of the central fiber. When $S = \operatorname{Spec} A$ for a local Noetherian k-algebra A with residue field k, we refer to $\mathcal{X}$ as a local deformation of X over A.

In this subsection, we restrict our attention to infinitesimal deformations. Let Art_k denote the category of local Artinian k-algebras with residue field k. For a fixed algebra $A \in \mathsf{Art}_k$, an isomorphism between two deformations $(\mathcal{X}, i)$ and $(\mathcal{Y}, j)$ over A is an isomorphism of schemes $\varphi\colon \mathcal{X} \to \mathcal{Y}$ over A that restricts to the identity on the central fiber X. This is encoded by the commutativity of the following diagram:

$$\begin{tikzcd} & X \arrow[dl, "i"'] \arrow[dr, "j"] & \\ \mathcal{X} \arrow[rr, "\varphi", "\sim"'] \arrow[dr] & & \mathcal{Y} \arrow[dl] \\ & \operatorname{Spec} A & \end{tikzcd}$$

Based on these definitions, we define the deformation functor of X as the covariant functor

$$\begin{aligned} \operatorname{Def}_X \colon \mathsf{Art}_k &\longrightarrow \mathsf{Sets} \\ A &\longmapsto \left\{ \begin{matrix} \text{Isomorphism classes of deformations} \\ \text{of } X \text{ over } A \end{matrix} \right\}. \end{aligned}$$

This is a prototypical example of a functor of Artin rings. Let $\widehat{\mathsf{Art}}_k$ denote the category of complete local Noetherian k-algebras with residue field k. Note that Art_k is a full subcategory of $\widehat{\mathsf{Art}}_k$. For any ring $R \in \widehat{\mathsf{Art}}_k$, we define the functor of points h_R restricted to Artinian algebras as

$$h_R \colon \mathsf{Art}_k \to \mathsf{Sets}, \quad A \longmapsto \operatorname{Hom}_{\widehat{\mathsf{Art}}_k}(R, A).$$

We say that a functor $F\colon \mathsf{Art}_k \to \mathsf{Sets}$ is pro-representable if there exists $R \in \widehat{\mathsf{Art}}_k$ such that F is naturally isomorphic to h_R.

The following consequence, which appears as an exercise in [**FM98**, Example 2.3], serves as an analogue of Yoneda's lemma in the calculus of functors of Artin rings and is fundamental to the study of obstruction theory.

Proposition 2.1.4. *Let $F\colon \mathsf{Art}_k \to \mathsf{Sets}$ be a functor of Artin rings. For any complete local ring $R \in \widehat{\mathsf{Art}}_k$, we define the completion of F at R by*

$$\widehat{F}(R) \coloneqq \varprojlim_n F(R/\mathfrak{m}_R^{n+1}).$$

There is a natural bijection between elements of $\widehat{F}(R)$ and the set of natural transformations from h_R to F:

$$\widehat{F}(R) \cong \operatorname{Nat}(h_R, F).$$

Proof. Let $\widehat{u} \in \widehat{F}(R)$. By definition, $\widehat{u}$ corresponds to a sequence of elements $\{u_n\}_{n \geq 0}$ with $u_n \in F(R/\mathfrak{m}_R^{n+1})$, satisfying the following compatibility condition: for every n, the natural projection $\pi_n\colon R/\mathfrak{m}_R^{n+1} \to R/\mathfrak{m}_R^n$ induces a map $F(\pi_n)$ such that $F(\pi_n)(u_n) = u_{n-1}$.

By the classical Yoneda lemma applied to the category Art_k, each element u_n defines a unique natural transformation $\Phi_n\colon h_{R/\mathfrak{m}_R^{n+1}} \to F$. The compatibility of the elements u_n implies the commutativity of the following diagram for all n:

$$\begin{tikzcd} h_{R/\mathfrak{m}_R^n} \arrow[rr] \arrow[dr, "\Phi_{n-1}"'] & & h_{R/\mathfrak{m}_R^{n+1}} \arrow[dl, "\Phi_n"] \\ & F & \end{tikzcd}$$

where the horizontal map is induced by the projection $R/\mathfrak{m}_R^{n+1} \to R/\mathfrak{m}_R^n$.

We now construct a natural transformation $\Phi\colon h_R \to F$. Let $A \in \mathsf{Art}_k$ and let $f \in h_R(A)$. Since A is an Artinian ring, its maximal ideal $\mathfrak{m}_A$ is nilpotent. Since f is a local homomorphism, it factors through $R/\mathfrak{m}_R^n$ for sufficiently large n. Let $\bar{f}_n \in h_{R/\mathfrak{m}_R^n}(A)$ be such a factorization. We define $\Phi_A(f) \coloneqq \Phi_n(\bar{f}_n) \in F(A)$. This value is independent of the choice of n due to the compatibility of the Φ_n's, which establishes a well-defined natural transformation $\Phi\colon h_R \to F$.

Conversely, given a natural transformation $\Psi\colon h_R \to F$, we obtain a sequence $\{u_n\}$ by evaluating Ψ on the canonical projections. Specifically, let $p_n\colon R \to R/\mathfrak{m}_R^{n+1}$ be the quotient map, which is an element of $h_R(R/\mathfrak{m}_R^{n+1})$. We define $u_n \coloneqq \Psi_{R/\mathfrak{m}_R^{n+1}}(p_n) \in F(R/\mathfrak{m}_R^{n+1})$. The naturality of Ψ ensures that this sequence is compatible, thereby defining an element $\widehat{u} \in \widehat{F}(R)$.

These two constructions are mutually inverse, establishing the bijection. □

Proposition 2.1.4 allows us to introduce the following definition [**Ser06**, Definition 2.2.6].

Definition 2.1.5. Let F be a functor of Artin rings. A formal element $\widehat{u} \in \widehat{F}(R)$ for some $R \in \widehat{\mathsf{Art}}_k$ is called *versal* if the natural transformation $h_R \to F$ defined by $\widehat{u}$ is smooth. Furthermore, $\widehat{u}$ is called semiuniversal if it is versal and the induced differential $h_R(k[\varepsilon]) \to F(k[\varepsilon])$ is bijective. Here $k[\varepsilon]$ denotes the dual number.

The following theorem, which is a special case of [**Ser06**, Theorem 2.4.1, Theorem 2.6.1], provides a characterization of when Def_X is pro-representable.

Theorem 2.1.6. *Let X be a nonsingular algebraic scheme. If $\dim_k \mathrm{H}^1(X, T_X) < \infty$, then Def_X has a semiuniversal element. Moreover, assuming Def_X has a semiuniversal element, the following conditions are equivalent:*

(1) The functor Def_X is pro-representable.
(2) For every square-zero extension $A' \to A$ in Art_k and every deformation $\mathcal{X}'$ of X over A', every automorphism of the induced deformation $\mathcal{X}' \times_{\operatorname{Spec} A'} \operatorname{Spec} A \to \operatorname{Spec} A$ is induced by an automorphism of $\mathcal{X}'$.

A closer analysis of the structure of these automorphisms reveals that if a projective algebraic scheme X has no nontrivial infinitesimal automorphisms, that is, $\mathrm{h}^0(X, T_X) = 0$, then Def_X is pro-representable (see [**Ser06**, Corollary 2.6.4]). On the other hand, the condition $\mathrm{h}^0(X, T_X) = 0$ implies that the automorphism group $\mathrm{Aut}(X)$ is finite. This once again highlights the subtle role that automorphisms play in the study of local moduli problems. Although the absence of nontrivial infinitesimal automorphisms is not a strictly necessary condition for the pro-representability of Def_X (for example, $X = \mathbb{P}^n$ is rigid but has $\mathrm{h}^0(X, T_X) > 0$), it still provides a strong and convenient criterion for determining pro-representability. Ultimately, this reminds us of the critical importance of incorporating automorphism data when studying moduli problems.

2.2. Stacks and First Examples

2.2.1. Pseudo-functors, Fibered Categories and Prestacks. From the discussions in Section 2.1, we see that in the study of moduli problems and moduli functors, the automorphism data of the objects is just as essential as the families of the objects themselves. Therefore, it is necessary to incorporate this automorphism information directly into the moduli functors. This leads to Grothendieck's philosophy: rather than merely specifying *when* two families are isomorphic, we must specify *how* they are isomorphic. More precisely, we must enhance the moduli functor to an assignment

$$F\colon \mathsf{Sch}^{\mathrm{op}} \longrightarrow \mathsf{Grpd}, \qquad S \longmapsto \mathrm{Fam}_S,$$

where $F(S)$ is now the groupoid of families over S.

However, we cannot simply regard F as a strict 2-functor. A technical issue naturally arises when defining pullbacks. Given morphisms of schemes $f\colon S \to T$ and $g\colon T \to U$, we can generally only expect a natural isomorphism between the composite pullback functors, rather than an honest equality. Because of the subtleties inherent in working with categories of structures (like modules or sheaves) via universal properties, we only obtain an invertible 2-morphism

$$\mu_{f,g}\colon (f^* \circ g^*) \xRightarrow{\cong} (g \circ f)^*.$$

Similarly, for the identity morphism $\mathrm{id}_S\colon S \to S$, the pullback functor id_S^* is canonically isomorphic, but rarely strictly identical, to the identity functor $\mathrm{id}_{F(S)}$. We use the following fundamental example to illustrate this phenomenon.

Example 2.2.1 (The pseudo-functor of quasi-coherent sheaves)**.** We consider the assignment $\underline{\mathrm{Qcoh}}(-)\colon \mathsf{Sch}^{\mathrm{op}} \to \mathsf{Grpd}$ which sends a scheme X to $\underline{\mathrm{Qcoh}}(X)$, the groupoid of quasi-coherent $\mathcal{O}_X$-modules. Recall that for any morphism of schemes $f\colon X \to Y$, the pullback is defined by the topological inverse image and the tensor product [**Har77**, II. 5]:

$$f^*\mathcal{G} \coloneqq \mathcal{O}_X \otimes_{f^{-1}\mathcal{O}_Y} f^{-1}\mathcal{G}.$$

Now, suppose we have a composition of scheme morphisms $X \xrightarrow{f} Y \xrightarrow{g} Z$, and let $\mathcal{F} \in \underline{\mathrm{Qcoh}}(Z)$. The pullback along the composition is defined as

$$(g \circ f)^*\mathcal{F} = \mathcal{O}_X \otimes_{(g\circ f)^{-1}\mathcal{O}_Z} (g \circ f)^{-1}\mathcal{F}.$$

On the other hand, applying the pullbacks consecutively yields

$$f^*(g^*\mathcal{F}) = \mathcal{O}_X \otimes_{f^{-1}\mathcal{O}_Y} f^{-1}(\mathcal{O}_Y \otimes_{g^{-1}\mathcal{O}_Z} g^{-1}\mathcal{F}).$$

Since the tensor product is defined via a universal property, the sheaves $(g \circ f)^*\mathcal{F}$ and $f^*(g^*\mathcal{F})$ are not strictly identical on the level of sets. However, the universal properties of the inverse image and the tensor product induce a canonical isomorphism

$$\alpha_{f,g}(\mathcal{F})\colon f^*(g^*\mathcal{F}) \xrightarrow{\cong} (g \circ f)^*\mathcal{F}.$$

A closer analysis shows that these isomorphisms $\alpha_{f,g}$ further satisfy a higher-level compatibility condition for any composable triple of morphisms. Recall that for any morphism $f\colon X \to Y$ between schemes, there is a canonical isomorphism of groups [**Har77**, II. 5]:

$$\Theta_f(\mathcal{F},\mathcal{G})\colon \mathrm{Hom}_{\mathcal{O}_Y}(\mathcal{G}, f_*\mathcal{F}) \xrightarrow{\cong} \mathrm{Hom}_{\mathcal{O}_X}(f^*\mathcal{G}, \mathcal{F}),$$

which is natural in $\mathcal{F} \in \underline{\mathrm{Qcoh}}(X)$ and $\mathcal{G} \in \underline{\mathrm{Qcoh}}(Y)$. In particular, let $\mathcal{F} = \mathcal{G}$ and $f = \mathrm{id}_X$; then by Yoneda's lemma, we obtain a canonical isomorphism $\varepsilon_X\colon \mathrm{id}_X^* \cong \mathrm{id}_{\mathrm{Qcoh}(X)}$. Now, let $X \xrightarrow{f} Y \xrightarrow{g} Z$ be a composition of morphisms and let $\mathcal{H} \in \mathrm{Qcoh}(Z)$. We have the following sequence of natural isomorphisms:

$$\begin{aligned}
\mathrm{Hom}_{\mathcal{O}_X}\left((gf)^*\mathcal{H}, -\right) &\cong \mathrm{Hom}_{\mathcal{O}_Z}\left(\mathcal{H}, (gf)_*-\right) \\
&= \mathrm{Hom}_{\mathcal{O}_Z}(\mathcal{H}, g_*f_*-) \\
&\cong \mathrm{Hom}_{\mathcal{O}_Y}(g^*\mathcal{H}, f_*-) \\
&\cong \mathrm{Hom}_{\mathcal{O}_X}(f^*g^*\mathcal{H}, -),
\end{aligned}$$

which is the composition of $\Theta_f(g^*\mathcal{H}) \circ \Theta_g(\mathcal{H}, f_*-) \circ \Theta_{gf}(\mathcal{H}, -)^{-1}$. By Yoneda's lemma, this induces an isomorphism $\alpha_{f,g}\colon f^*g^* \cong (gf)^*$, which aligns with the isomorphism deduced above via the universal property. From this categorical characterization, one can verify that

(1) For any morphism $f\colon X \to Y$ and $\mathcal{G} \in \underline{\mathrm{Qcoh}}(Y)$,

$$\alpha_{\mathrm{id}_X, f}(\mathcal{G}) = \varepsilon_X(f^*\mathcal{G}), \qquad \alpha_{f, \mathrm{id}_Y}(\mathcal{G}) = f^*\varepsilon_Y(\mathcal{G}).$$

(2) Given morphisms $U \xrightarrow{f} V \xrightarrow{g} W \xrightarrow{h} T$ and $\mathcal{J} \in \underline{\mathrm{Qcoh}}(T)$, the following diagram commutes:

$$\begin{tikzcd}
f^*g^*h^*\mathcal{J} \arrow[r, "\alpha_{f,g}(h^*\mathcal{J})"] \arrow[d, "f^*\alpha_{g,h}(\mathcal{J})"] & (gf)^*h^*\mathcal{J} \arrow[d, "\alpha_{gf,h}(\mathcal{J})"] \\
f^*(hg)^*\mathcal{J} \arrow[r, "\alpha_{f,hg}(\mathcal{J})"] & (hgf)^*\mathcal{J}
\end{tikzcd}$$

These coherence requirements appearing in the example above lead to the general notions of *pseudo-functors* and *categories fibered in groupoids*, which provide the precise mathematical language to handle these assignments rigorously.

Definition 2.2.2 (pseudo-functors)**.** A pseudo-functor Φ on a category $\mathcal{C}$ consists of the following data:

(1) for each object U of $\mathcal{C}$, a category ΦU;
(2) for each morphism $f\colon U \to V$, a functor $f^*\colon \Phi V \to \Phi U$;
(3) for each object U of $\mathcal{C}$, an isomorphism

$$\varepsilon_U\colon\ \mathrm{id}_U^* \xrightarrow{\cong} \mathrm{id}_{\Phi U};$$

(4) for each pair of composable morphisms $U \xrightarrow{f} V \xrightarrow{g} W$, an isomorphism

$$\alpha_{f,g}\colon f^*g^* \xrightarrow{\cong} (g \circ f)^*.$$

This data is subject to the following compatibility conditions:

(a) If $f\colon U \to V$ is a morphism in $\mathcal{C}$ and η is an object of ΦV, then

$$\alpha_{\mathrm{id}_U, f}(\eta) = \varepsilon_U(f^*\eta), \qquad \alpha_{f, \mathrm{id}_V}(\eta) = f^*\varepsilon_V(\eta).$$

(b) Given morphisms $U \xrightarrow{f} V \xrightarrow{g} W \xrightarrow{h} T$ and an object θ of ΦT, the following diagram commutes:

$$
\begin{tikzcd}
f^*g^*h^*\theta \arrow[r, "\alpha_{f,g}(h^*\theta)"] \arrow[d, "f^*\alpha_{g,h}(\theta)"] & (gf)^*h^*\theta \arrow[d, "\alpha_{gf,h}(\theta)"] \\
f^*(hg)^*\theta \arrow[r, "\alpha_{f,hg}(\theta)"] & (hgf)^*\theta
\end{tikzcd}
$$

The concept of a pseudo-functor, although it arises naturally from our previous discussion, can be highly technical and practically cumbersome when managing coherence conditions. This motivates an alternative approach: "assembling" all the values of a pseudo-functor into a unified categorical structure. In other words, a pseudo-functor $\Phi\colon \mathcal{C}^{\mathrm{op}} \to \mathsf{Cat}$ can be viewed as encoding a category fibered over $\mathcal{C}$, where the pullback and coherence conditions describe how objects vary over morphisms in the base.

Definition 2.2.3 (Fibered categories)**.** A *fibered category* over a base category $\mathcal{C}$ is a category $\mathcal{F}$ equipped with a functor $p\colon \mathcal{F} \to \mathcal{C}$ such that for every morphism $f\colon U \to V$ in $\mathcal{C}$ and every object $\eta \in \mathcal{F}$ lying over V (i.e., $p(\eta) = V$), there exists a *Cartesian morphism* $\phi\colon \xi \to \eta$ in $\mathcal{F}$ lifting f (so $p(\phi) = f$):

$$
\begin{tikzcd}
\xi \arrow[r, dashed, "\phi"] \arrow[d, dashed, mapsto, "p"'] & \eta \arrow[d, mapsto, "p"] \\
U \arrow[r, "f"'] & V
\end{tikzcd}
$$

Recall that $\phi\colon \xi \to \eta$ is Cartesian if it satisfies the following universal property: for any object $\gamma \in \mathcal{F}$ and any morphism $\psi\colon \gamma \to \eta$ such that $p(\psi) = f \circ h$ for some $h\colon p(\gamma) \to U$, there exists a unique morphism $\chi\colon \gamma \to \xi$ over h (meaning $p(\chi) = h$) such that $\psi = \phi \circ \chi$. We also say ξ is a pullback of η to U, denoted by $f^*\eta$.

This brings us to the first main result of this section: the concept of fibered categories over $\mathcal{C}$ is equivalent to the concept of pseudo-functors $\mathcal{C}^{\mathrm{op}} \to \mathsf{Cat}$. More precisely, given a category $\mathcal{C}$, there is a bijective correspondence:

$$
\left\{ \begin{array}{l} \text{fibered categories over } \mathcal{C} \\ \text{equipped with a cleavage} \end{array} \right\} \longleftrightarrow \left\{ \begin{array}{l} \text{pseudo-functors} \\ \Phi\colon \mathcal{C}^{\mathrm{op}} \to \mathsf{Cat} \end{array} \right\}.
$$

Let us briefly recall the motivation for introducing the concept of a *cleavage* for a fibered category $p\colon \mathcal{F} \to \mathcal{C}$. Let $f\colon U \to V$ be a morphism in $\mathcal{C}$. If we wish to define a pseudo-functor Φ from the data of p, the most natural starting point is to define $\Phi(U) \coloneqq \mathcal{F}(U)$, where the fiber $\mathcal{F}(U)$ is the subcategory of $\mathcal{F}$ consisting of objects A with $p(A) = U$, and morphisms $\varphi\colon A \to B$ such that $p(\varphi) = \mathrm{id}_U$.

Next, to define the pullback functor $\Phi(f)\colon \mathcal{F}(V) \to \mathcal{F}(U)$, a natural approach is to send $\eta \in \mathcal{F}(V)$ to $f^*\eta \in \mathcal{F}(U)$. However, the pullback $f^*\eta$ is not strictly unique; it is only uniquely determined up to a unique isomorphism. To make this assignment a well-defined functor, we need to choose a specific pullback $f^*\eta$ for every f and η. This motivates the definition of a cleavage: a *cleavage* of a fibered category $p\colon \mathcal{F} \to \mathcal{C}$ consists of a designated class K of Cartesian arrows in $\mathcal{F}$ such that for each morphism $f\colon U \to V$ in $\mathcal{C}$ and each object $\eta \in \mathcal{F}(V)$, there exists a unique morphism in K with target η mapping to f in $\mathcal{C}$.

Assuming the fibered category $p\colon \mathcal{F} \to \mathcal{C}$ is equipped with a cleavage, the functor $\Phi(f)$ becomes well-defined. One checks that the assignment $\Phi\colon \mathcal{C}^{\mathrm{op}} \to \mathsf{Cat}$ constructed in this manner satisfies the coherence conditions of a pseudo-functor.

The converse direction—constructing a fibered category from a pseudo-functor—is more involved. We will outline the construction here for the case where the pseudo-functor $\Phi\colon \mathcal{C}^{\mathrm{op}} \to \mathsf{Cat}$ is an honest functor, referring the reader to standard literature (e.g., [**FGI+05**, Section 3.1.3]) for the general pseudo-functorial case. This process is well known as the *Grothendieck construction*.

We construct the category $\mathcal{F}$ as follows. An object of $\mathcal{F}$ is defined as a pair (ξ, U), where U is an object in $\mathcal{C}$ and ξ is an object in the category $\Phi(U)$. A morphism $(a, f)\colon (\xi, U) \to (\eta, V)$ consists of a morphism $f\colon U \to V$ in $\mathcal{C}$ together with a morphism $a\colon \xi \to \Phi(f)(\eta)$ in $\Phi(U)$.

To make $\mathcal{F}$ into a category, we define composition as follows. Given two morphisms $(a, f)\colon (\xi, U) \to (\eta, V)$ and $(b, g)\colon (\eta, V) \to (\gamma, W)$, we apply the functor $\Phi(f)$ to the morphism b to obtain $\Phi(f)(b)\colon \Phi(f)(\eta) \to \Phi(f)(\Phi(g)(\gamma))$. Since Φ is an honest functor, $\Phi(f) \circ \Phi(g) = \Phi(g \circ f)$, giving us $\Phi(f)(b)\colon \Phi(f)(\eta) \to \Phi(g\circ f)(\gamma)$. We can then compose this with a in $\Phi(U)$ to define the composition rule:

$$(b, g) \circ (a, f) \coloneqq (\Phi(f)(b) \circ a,\ g \circ f).$$

The identity morphism for an object (ξ, U) is naturally given by $(\mathrm{id}_\xi, \mathrm{id}_U)$. Finally, the projection functor $p\colon \mathcal{F} \to \mathcal{C}$ simply forgets the fiber data, acting on objects as $p(\xi, U) = U$ and on morphisms as $p(a, f) = f$. One can directly verify that p makes $\mathcal{F}$ into a fibered category over $\mathcal{C}$.

The preceding discussion leads us to the formal definition of a prestack.

Definition 2.2.4 (Prestack)**.** A *prestack*, also called a *category fibered in groupoids*, over a base category $\mathcal{C}$ is a fibered category $p\colon \mathcal{F} \to \mathcal{C}$ such that for every object $U \in \mathcal{C}$, the fiber category $\mathcal{F}(U)$ is a groupoid.

Equivalently, by the correspondence established via the Grothendieck construction, a prestack corresponds to a pseudo-functor $\Phi\colon \mathcal{C}^{\mathrm{op}} \to \mathsf{Grpd}$.

We close this subsection by studying several fundamental examples that will serve as the primary objects of study in the succeeding sections.

Example 2.2.5 (The prestack of quasi-coherent sheaves)**.** We define the category of quasi-coherent sheaves $\underline{\mathrm{Qcoh}}$ over the base category Sch as follows. The objects of $\underline{\mathrm{Qcoh}}$ consist of pairs $(X, \mathcal{F})$, where $X \in \mathsf{Sch}$ and $\mathcal{F} \in \mathrm{Qcoh}(X)$. A morphism from $(X, \mathcal{F})$ to $(Y, \mathcal{G})$ consists of a pair (f, α), where $f\colon X \to Y$ is a morphism of schemes and $\alpha\colon \mathcal{F} \xrightarrow{\cong} f^*\mathcal{G}$ is an isomorphism of $\mathcal{O}_X$-modules. The composition of two morphisms $(f, \alpha)\colon (X, \mathcal{F}) \to (Y, \mathcal{G})$ and $(g, \beta)\colon (Y, \mathcal{G}) \to (Z, \mathcal{H})$ is given by the scheme morphism $g \circ f$ and the composite isomorphism $\mathcal{F} \xrightarrow{\alpha} f^*\mathcal{G} \xrightarrow{f^*\beta} f^*g^*\mathcal{H} \cong (g\circ f)^*\mathcal{H}$. The projection functor $p\colon \underline{\mathrm{Qcoh}} \to \mathsf{Sch}$ simply forgets the sheaf data, sending $(X, \mathcal{F})$ to X and (f, α) to f. Following the coherence properties established in Example 2.2.1, one easily verifies that $\underline{\mathrm{Qcoh}}$ is a category fibered in groupoids, and thus a prestack over Sch.

Example 2.2.6 (Quotient prestacks)**.** Let G be a smooth group scheme and X be a scheme equipped with a right action of G. We define the quotient prestack $[X/G]^{\mathrm{pre}}$ as a category over Sch as follows. The objects of $[X/G]^{\mathrm{pre}}$ consist of pairs (T, x), where $T \in \mathsf{Sch}$ and $x \in X(T)$ is a T-valued point of X. A morphism from (T, x) to (S, y) over a scheme morphism $u\colon T \to S$ is an element $g \in G(T)$ such that $x \cdot g = y \circ u$ in $X(T)$. The composition of morphisms $(u, g_1)\colon (T, x) \to (S, y)$ and $(v, g_2)\colon (S, y) \to (W, z)$ is given by $(v \circ u, g_1 \cdot (g_2 \circ u))$, which utilizes the group structure of $G(T)$. The projection map to Sch is defined by forgetting the geometric point. One can verify that this forms a prestack over Sch. Notice that for any fixed scheme $T \in \mathsf{Sch}$, the objects in the fiber over T are the set $X(T)$, and the morphisms over id_T correspond to $G(T)$.

Example 2.2.7 (The moduli prestack of genus g curves)**.** Let $g \geq 0$ be a non-negative integer. We define $\mathcal{M}_g^{\mathrm{pre}}$ as a category over Sch where the objects are pairs $(S, \pi\colon \mathcal{C} \to S)$, such that $S \in \mathsf{Sch}$ and π is a family of smooth curves of genus g (as defined in Definition 2.0.1). A morphism from $(S, \mathcal{C} \to S)$ to $(T, \mathcal{D} \to T)$ consists of a pair of morphisms $f\colon S \to T$ and $F\colon \mathcal{C} \to \mathcal{D}$ such that the following diagram is Cartesian:

$$\begin{tikzcd} \mathcal{C} \arrow[r, "F"] \arrow[d, "\pi_{\mathcal{C}}"'] \arrow[dr, phantom, "\square"] & \mathcal{D} \arrow[d, "\pi_{\mathcal{D}}"] \\ S \arrow[r, "f"'] & T \end{tikzcd}$$

The Cartesian condition ensures that $\mathcal{C}$ is isomorphic to the pullback $f^*\mathcal{D} = S \times_T \mathcal{D}$. The projection to Sch simply maps the family to its base scheme S. This gives a prestack over Sch.

2.2.2. Grothendieck Topology and Stacks. The notion of a prestack plays a role analogous to that of a presheaf in classical sheaf theory. Naturally, this raises the question of how to elevate a prestack to a *stack*. Just as a presheaf must satisfy a local-to-global gluing axiom to become a sheaf, a prestack must satisfy descent conditions allowing for the gluing of both objects and morphisms. However, to rigorously discuss the gluing of categorical and geometric data, we first need a framework to define what constitutes a "local covering" of an object within an arbitrary category. This requirement leads to the concept of a *Grothendieck topology*, which generalizes the classical topological notion of open covers.

Through this, we obtain a framework in which prestacks can be equipped with analogous gluing conditions. The resulting objects are called stacks. Therefore, morally, a stack may be regarded as a "sheaf of groupoids" on the base category $\mathcal{C}$.

Definition 2.2.8 (Grothendieck topology and sites)**.** Let $\mathcal{S}$ be a category. A *Grothendieck topology* on $\mathcal{S}$ consists of the following data: for each object $X \in \mathcal{S}$, there is a collection of sets of arrows $\{X_i \to X\}_{i \in I}$, called *coverings* of X and denoted by $\mathrm{Cov}(X)$, satisfying the following conditions:

(1) (*Isomorphisms*) If $f\colon Y \to X$ is an isomorphism, then the single-element set $\{f\} \in \mathrm{Cov}(X)$.
(2) (*Base change*) If $\{X_i \to X\}_{i \in I} \in \mathrm{Cov}(X)$ and $Y \to X$ is any arrow in $\mathcal{S}$, then the fiber products $X_i \times_X Y$ exist for all $i \in I$, and the collection of induced projections $\{X_i \times_X Y \to Y\}_{i \in I}$ belongs to $\mathrm{Cov}(Y)$.
(3) (*Composition*) If $\{X_i \to X\}_{i \in I} \in \mathrm{Cov}(X)$ and for each $i \in I$ we have a covering $\{X_{ij} \to X_i\}_{j \in J_i} \in \mathrm{Cov}(X_i)$, then the collection of composite arrows $\{X_{ij} \to X_i \to X\}_{i \in I, j \in J_i}$ belongs to $\mathrm{Cov}(X)$.

A category $\mathcal{S}$ equipped with a Grothendieck topology is called a *site*.

As examples, we list some of the most important and frequently used sites throughout this thesis.

Example 2.2.9 (The big topological site)**.** Consider the category Top of topological spaces. For any $X \in \mathsf{Top}$, a covering of X is a collection of open immersions $\{X_i \hookrightarrow X\}$ such that $X = \bigcup_i X_i$.

Example 2.2.10 (The big étale site)**.** Consider the category Sch. An étale covering of a scheme X is a collection of étale morphisms $\{T_i \to X\}$ such that the induced morphism $\bigsqcup T_i \to X$ is surjective. We denote the category Sch equipped with this topology by $\mathsf{Sch}_{\text{ét}}$.

Example 2.2.11 (The fppf and fpqc topologies)**.** Recall that a morphism of schemes $f\colon X \to Y$ is called *fppf* (fidèlement plate de présentation finie) if f is faithfully flat and locally of finite presentation, and is called *fpqc* (fidèlement plat et quasi-compacte) if f is faithfully flat and every quasi-compact open subset of Y is the image of a quasi-compact open subset of X. See [**FGI+05**, Proposition 2.33] for a finer characterization of fpqc morphisms. An fppf (resp. fpqc) covering of X is a jointly surjective collection of morphisms $\{X_i \to X\}$ such that $\bigsqcup X_i \to X$ is fppf (resp. fpqc).

Example 2.2.12 (The lisse-étale site)**.** Let $\mathcal{X}$ be a scheme. The lisse-étale site of $\mathcal{X}$, denoted by $\mathcal{X}_{\text{lis-ét}}$, has as its underlying category the objects $U \to \mathcal{X}$, where U is a scheme and the morphism is smooth. The morphisms in this category are $\mathcal{X}$-morphisms $U \to V$ (which are not required to be smooth themselves). A covering of an object $U \to \mathcal{X}$ is a collection of étale morphisms $\{U_i \to U\}$ over $\mathcal{X}$ such that the induced map $\bigsqcup U_i \to U$ is surjective. This topology is particularly vital for defining quasi-coherent sheaves on algebraic stacks.

Let $\mathcal{S}$ be a site. Having established the necessary categorical machinery, we can now formalize the concept of sheaves and their 2-categorical generalizations—stacks over a site—by extending the local-to-global gluing conditions of classical sheaf theory.

Definition 2.2.13 (Sheaves)**.** Let $\mathcal{S}$ be a site and let $F\colon \mathcal{S}^{\mathrm{op}} \to \mathsf{Sets}$ be a presheaf. We say F is a *sheaf* if for any covering $\{U_i \to U\}$ of an object U in $\mathcal{S}$, and any collection of local sections $a_i \in F(U_i)$ such that their pullbacks agree on the overlaps, meaning $p_1^* a_i = p_2^* a_j \in F(U_i \times_U U_j)$ for all i, j, there exists a unique global section $a \in F(U)$ whose pullback to $F(U_i)$ is a_i for all i. Here, $p_1\colon U_i \times_U U_j \to U_i$ and $p_2\colon U_i \times_U U_j \to U_j$ denote the canonical projections. This means the following sequence of sets is an equalizer diagram:

$$F(U) \longrightarrow \textstyle\prod_i F(U_i) \underset{p_2^*}{\overset{p_1^*}{\rightrightarrows}} \textstyle\prod_{i,j} F(U_i \times_U U_j).$$

Definition 2.2.14 (Stacks)**.** Let $\mathcal{F}$ be a prestack over a site $\mathcal{S}$. We say that $\mathcal{F}$ is a *stack* if the following two gluing conditions hold for every object $S \in \mathcal{S}$ and every covering $\{S_i \to S\}_{i \in I}$ in $\mathcal{S}$:

(1) (*Morphisms glue*) Let a and b be objects in the fiber category $\mathcal{F}(S)$. Suppose we are given morphisms $\varphi_i\colon a|_{S_i} \to b|_{S_i}$ in $\mathcal{F}(S_i)$ such that their restrictions to the double overlaps coincide: $\varphi_i|_{S_{ij}} = \varphi_j|_{S_{ij}}$ for all $i, j \in I$. Then there exists a unique global morphism $\varphi\colon a \to b$ in $\mathcal{F}(S)$ such that $\varphi|_{S_i} = \varphi_i$ for all i.

(2) (*Objects glue*) Let $a_i \in \mathcal{F}(S_i)$ be a collection of local objects for each $i \in I$. Suppose that for any $i, j \in I$, there exists an isomorphism $\alpha_{ij}\colon a_i|_{S_{ij}} \xrightarrow{\cong} a_j|_{S_{ij}}$ in $\mathcal{F}(S_{ij})$, satisfying the *cocycle condition* on the triple overlaps S_{ijk}:

$$\alpha_{jk}|_{S_{ijk}} \circ \alpha_{ij}|_{S_{ijk}} = \alpha_{ik}|_{S_{ijk}}.$$

Then this data guarantees the existence of a global object $a \in \mathcal{F}(S)$ and a collection of isomorphisms $\varphi_i\colon a|_{S_i} \xrightarrow{\cong} a_i$ in $\mathcal{F}(S_i)$, such that $\varphi_j|_{S_{ij}} = \alpha_{ij} \circ \varphi_i|_{S_{ij}}$ for all i, j.

Remark 2.2.15 (Category of descent data)**.** In standard literature, such as [**FGI+05**, Definition 4.6], stacks over a site $\mathcal{S}$ are defined using the language of descent data. Let $\mathcal{U} = \{U_i \to U\}$ be a covering in $\mathcal{S}$ and let $\mathcal{F}$ be a fibered category over $\mathcal{S}$. An *object with descent data* on $\mathcal{U}$, denoted $(\{\xi_i\}, \{\varphi_{ij}\})$, consists of a collection of objects $\xi_i \in \mathcal{F}(U_i)$ together with isomorphisms $\varphi_{ij}\colon p_2^*\xi_j \xrightarrow{\cong} p_1^*\xi_i$ in $\mathcal{F}(U_i \times_U U_j)$ satisfying the cocycle condition $p_{13}^*\varphi_{ik} = p_{12}^*\varphi_{ij} \circ p_{23}^*\varphi_{jk}$ in $\mathcal{F}(U_i \times_U U_j \times_U U_k)$. Morphisms between objects with descent data are defined via families of compatible morphisms in the respective fiber categories, forming the category of descent data $\mathcal{F}(\mathcal{U})$. The fibered category $\mathcal{F}$ is defined to be a stack over $\mathcal{S}$ if, for every covering $\mathcal{U}$, the natural pullback functor $\mathcal{F}(U) \to \mathcal{F}(\mathcal{U})$ is an equivalence of categories. This categorical equivalence elegantly unpacks into the two standard stack axioms: full faithfulness corresponds to the unique gluing of morphisms, while essential surjectivity corresponds to the gluing of local objects.

Formulating this definition rigorously requires verifying that the category of descent data is invariant up to equivalence under the choice of a cleavage for the fibered category. This involves subtle and lengthy categorical argument. To maintain our focus, we omit this verification. For a comprehensive treatment of these categorical nuances, see [**FGI+05**, Section 4.1.2].

Before moving forward, we address a fundamental question: equipped with which Grothendieck topology, denoted by τ, does every scheme in Sch_τ become a stack? By Yoneda's lemma, this is equivalent to asking under which topology the representable functor $h_X \coloneqq \mathrm{Hom}(-, X)$ is a sheaf. This is a foundational issue because, following our previous discussions, the 2-category of stacks should be regarded as a natural enlargement of the category of schemes.

It turns out that not every topology automatically makes every scheme a stack. For instance, consider the naive *flat topology*, whose coverings $\{U_i \to U\}$ are jointly surjective collections of flat morphisms. This topology is known to be ill-behaved; without additional quasi-compactness or finiteness conditions, not all representable functors are sheaves under this topology (see [**FGI+05**, Remark 2.56]).

Fortunately, the fpqc topology introduced in Example 2.2.11 provides the correct framework. A celebrated theorem of Grothendieck establishes that any representable functor on Sch is indeed a sheaf in the fpqc topology. As a direct corollary, since the fpqc topology is finer than both the fppf and étale topologies, every scheme also defines a stack in the fppf and étale topologies (Example 2.2.10).

The proof of the preceding theorem is very detailed in Vistoli's notes [**FGI+05**, Theorem 2.55]. We therefore refer the reader to this text for the full details. The key observation relies on the following proposition [**FGI+05**, Lemma 2.60], which serves as an indispensable tool for our subsequent study.

Proposition 2.2.16. *Let $F\colon \mathsf{Sch}^{\mathrm{op}} \to \mathsf{Sets}$ be a functor. Suppose that F satisfies the following two conditions:*

(1) F is a sheaf in the big Zariski topology;
(2) for any faithfully flat morphism $V \to U$ of affine schemes, there is an equalizer sequence:

$$F(U) \longrightarrow F(V) \underset{p_2^*}{\overset{p_1^*}{\rightrightarrows}} F(V \times_U V).$$

Then F is a sheaf in the fpqc topology, and consequently a sheaf in both the fppf and étale topologies.

Equipped with this technical lemma, the proof of Grothendieck's theorem is vastly simplified, as the representable functor h_X is evidently a sheaf in the big Zariski topology. The problem can be furthermore boiled down to the affine case where $X = \operatorname{Spec} A$. This reduces the verification of the second condition in Proposition 2.2.16 to a purely algebraic fact: for any faithfully flat ring homomorphism $f\colon A \to B$, there is an equalizer sequence

$$0 \longrightarrow A \xrightarrow{\ f\ } B \underset{b \mapsto b\otimes 1}{\overset{b\mapsto 1\otimes b}{\rightrightarrows}} B \otimes_A B.$$

Grothendieck's theorem can also be regarded as the beginning of *descent theory*, which essentially asks whether a given functor $F\colon \mathsf{Sch}^{\mathrm{op}} \to \mathsf{Sets}$ (or a prestack $\mathcal{F}\colon \mathsf{Sch}^{\mathrm{op}} \to \mathsf{Grpd}$) is a sheaf (or a stack) in a specified Grothendieck topology τ. We can now appreciate the true power of Proposition

2.2.16: when τ is the fpqc topology, it suffices to verify the gluing conditions in the classical algebraic geometry context (the big Zariski topology) and then check the affine case.

Remark 2.2.17. We can formally generalize Proposition 2.2.16 from presheaves to prestacks, but doing so requires categorifying the concept of an equalizer. In the 2-categorical setting, the equalizer of sets is replaced by a 2-limit. More precisely, for a covering $V \to U$, the 2-limit of the sequence of groupoids translates to the category of descent data (as discussed in Remark 2.2.15).

Therefore, the stack-theoretic analogue of Proposition 2.2.16 states that a prestack $\mathcal{F}\colon \mathsf{Sch}^{\mathrm{op}} \to \mathsf{Grpd}$ is an fpqc stack if and only if it is a Zariski stack and the canonical pullback functor $\mathcal{F}(U) \to \mathcal{F}(V \to U)$ is an equivalence of categories for all faithfully flat morphisms of affine schemes $V \to U$. Although this appears to be a straightforward generalization, the proof is technical and non-trivial (see [**FGI⁺05**, Lemma 4.25]).

As an application, we use this stack-theoretic analogue to prove the following foundational result.

Proposition 2.2.18 (Descent for quasi-coherent sheaves)**.** *The prestack of quasi-coherent sheaves* $\underline{\mathrm{Qcoh}}$ *(Example 2.2.5) is a stack over* Sch *in the fpqc topology.*

Proof. By the stack-theoretic generalization of Proposition 2.2.16 discussed above, it suffices to check two conditions that $\underline{\mathrm{Qcoh}}$ is a stack in the big Zariski topology, and that effective descent holds for faithfully flat morphisms of affine schemes. The first condition is satisfied by classical algebraic geometry: quasi-coherent sheaves and their morphisms glue over Zariski open covers.

For the second condition, we reduce the problem to commutative algebra. Let $f\colon A \to B$ be a faithfully flat ring homomorphism, corresponding to an fpqc cover $\operatorname{Spec} B \to \operatorname{Spec} A$. A quasi-coherent sheaf on $\operatorname{Spec} A$ is equivalent to an A-module M, and its pullback to $\operatorname{Spec} B$ is $M \otimes_A B$. A descent datum on $\operatorname{Spec} B$ relative to $\operatorname{Spec} A$ consists of a B-module N along with an isomorphism

$$\varphi\colon N \otimes_A B \xrightarrow{\cong} B \otimes_A N$$

(regarded as an isomorphism of $B \otimes_A B$-modules) satisfying the cocycle condition as $B \otimes_A B \otimes_A B$-modules.

Grothendieck's theorem on faithfully flat descent for modules guarantees that the category of such descent data is canonically equivalent to the category of A-modules. Specifically, given (N, φ), the A-module M is uniquely recovered as the equalizer of the two maps induced by the descent datum:

$$M \longrightarrow N \rightrightarrows N \otimes_A B.$$

Since $A \to B$ is faithfully flat, tensoring this exact sequence with B recovers the isomorphism $M \otimes_A B \cong N$, proving that the pullback functor is essentially surjective. Furthermore, the faithful flatness of $A \to B$ ensures the exactness of the Amitsur complex

$$0 \to M \to M \otimes_A B \rightrightarrows M \otimes_A B \otimes_A B$$

for any A-module M. This guarantees that morphisms of A-modules uniquely glue, so the pullback functor is fully faithful. Together, this proves that the pullback functor is an equivalence of categories. Concluding by the stack analogue of Proposition 2.2.16, $\underline{\mathrm{Qcoh}}$ is a stack in the fpqc topology. For the full details of the commutative algebra argument above, see [**FGI⁺05**, Section 4.2.1]. ☐

Corollary 2.2.19 (fpqc descent of open and closed subschemes)**.** *Let $f\colon S' \to S$ be an fpqc morphism of schemes. If $Z' \subseteq S'$ is a closed (resp. open) subscheme such that $p_1^{-1}(Z') = p_2^{-1}(Z')$ as subschemes of $S' \times_S S'$, then there exists a unique closed (resp. open) subscheme $Z \subseteq S$ such that $Z' = f^{-1}(Z)$.*

Proof. For the closed case, there is a canonical one-to-one correspondence between closed immersions $Z' \hookrightarrow S'$ and quasi-coherent ideal sheaves $\mathcal{I}_{Z'} \subseteq \mathcal{O}_{S'}$. The equality of the pullbacks of Z' over $S' \times_S S'$ guarantees that the descent data for the ideal sheaf $\mathcal{I}_{Z'}$ is effective. Therefore, by the fpqc descent of quasi-coherent sheaves (Proposition 2.2.18), $\mathcal{I}_{Z'}$ descends to a unique quasi-coherent ideal sheaf $\mathcal{I}_Z \subseteq \mathcal{O}_S$, which defines the desired closed subscheme $Z \subseteq S$.

For the case of an open subscheme, let $|Z'|$ denote the underlying topological space of Z', and define the set $Z = f(|Z'|) \subseteq S$. The condition $p_1^{-1}(Z') = p_2^{-1}(Z')$ guarantees that $f^{-1}(Z) = |Z'|$ set-theoretically. Since f is fpqc, it is a submersive map (i.e. f is surjective and S has the quotient topology) [**Alp26**, Appendix A.4.9]. By definition of the quotient topology, a subset of S is open if and only if its preimage in S' is open. Since $f^{-1}(Z) = |Z'|$ is open in S', it follows that Z is an open subset of S. Endowing Z with the canonical open subscheme structure yields an open subscheme whose pullback to S' has the underlying space $|Z'|$. Since an open subscheme is uniquely determined by its underlying topological space, this pullback is isomorphic to Z'. □

Until now, thanks to Proposition 2.2.16, the application of descent theory has largely felt like classical algebraic geometry combined with commutative algebra. However, to prove that a geometric moduli space is a stack, we also need to study the intrinsic geometry of the objects it parameterizes. Before we can prove that $\mathcal{M}_g^{\mathrm{pre}}$ is a stack, we need to collect some foundational geometric facts regarding families of smooth curves of genus $g \geq 2$.

Proposition 2.2.20 (Properties of families of smooth curves)**.** *Let $\pi\colon \mathcal{C} \to S$ be a family of smooth curves of genus $g \geq 2$. Then*

(1) $\pi_*\mathcal{O}_{\mathcal{C}} = \mathcal{O}_S$.

(2) The pushforward $\pi_*\left(\Omega_{\mathcal{C}/S}^{\otimes k}\right)$ is a vector bundle of rank

$$\mathrm{rk}(k) = \begin{cases} g, & k = 1; \\ (2k-1)(g-1), & k \geq 2. \end{cases}$$

(3) The higher direct image $\mathrm{R}^1\pi_*\Omega_{\mathcal{C}/S}^{\otimes k}$ is isomorphic to $\mathcal{O}_S$ when $k = 1$ and zero otherwise.

(4) For $k \geq 3$, the sheaf $\Omega_{\mathcal{C}/S}^{\otimes k}$ is relatively very ample. As a direct corollary, the morphism $\pi\colon \mathcal{C} \to S$ is projective.

The proof of this proposition is an application of cohomology and base change theorems (see [**Har77**, III. 12]). We refer the reader to [**Alp26**, Proposition A.6.11] for the complete details of the proof.

Proposition 2.2.21. *When $g \geq 2$, the prestack $\mathcal{M}_g^{\mathrm{pre}}$ (Example 2.2.7) is a stack over* Sch *in the fpqc topology (and consequently, in the fppf and étale topologies).*

Proof. By the stack-theoretic generalization of Proposition 2.2.16 and the fact that smooth curves and their morphisms glue in the Zariski topology, it suffices to show that the two stack axioms are satisfied with respect to a fpqc morphism of affine schemes $S' \to S$.

We first verify the first axiom (morphisms glue). Suppose $\mathcal{C} \to S$ and $\mathcal{D} \to S$ are two families of smooth curves of genus g. Let $F'\colon \mathcal{C}_{S'} \to \mathcal{D}_{S'}$ be an S'-morphism such that its two pullbacks to $S' \times_S S'$ agree. By Grothendieck's fpqc descent theorem for morphisms between schemes, F' descends to a unique global morphism $F\colon \mathcal{C} \to \mathcal{D}$ fitting into the following Cartesian diagram:

$$\begin{tikzcd}
\mathcal{C}_{S'} \arrow[r, "F'"] \arrow[d] \arrow[dr, phantom, "\square"] & \mathcal{D}_{S'} \arrow[d] \\
\mathcal{C} \arrow[r, dashed, "\exists! F"] & \mathcal{D}
\end{tikzcd}$$

Furthermore, if F' is an isomorphism, the descended morphism F is also an isomorphism. This relies on the fact that being an isomorphism is a property local on the target in the fpqc topology (see [**Sta26**, Tag 02YJ]). This proves the first stack axiom.

The verification of the second axiom (objects glue) is more involved. Let $\mathcal{C}' \to S'$ be a family of smooth curves equipped with descent data, meaning an isomorphism $\alpha\colon p_1^*\mathcal{C}' \xrightarrow{\cong} p_2^*\mathcal{C}'$ over $S' \times_S S'$ satisfying the cocycle condition on the triple fiber product. We prove there exists a family of smooth curves $\mathcal{C} \to S$ and an isomorphism $\varphi\colon \mathcal{C}_{S'} \xrightarrow{\cong} \mathcal{C}'$ compatible with α.

The key to this proof is the geometric fact that $E' := \pi_*\Omega^{\otimes 3}_{\mathcal{C}'/S'}$ is a vector bundle on S' of rank $5g-5$, and $\Omega^{\otimes 3}_{\mathcal{C}'/S'}$ is relatively very ample (Proposition 2.2.20). This provides a canonical closed immersion $\mathcal{C}' \hookrightarrow \mathbb{P}(E')$.

The isomorphism α induces an isomorphism of vector bundles $\beta\colon p_1^*E' \xrightarrow{\cong} p_2^*E'$ that also satisfies the cocycle condition. By the fpqc descent of quasi-coherent sheaves (Proposition 2.2.18), we obtain a quasi-coherent sheaf E on S such that its pullback to S' is isomorphic to E'. Next, we utilize the descent of fpqc local finiteness properties for quasi-coherent sheaves (see [**Sta26**, Tag 05AY]) to conclude that E is, in fact, a vector bundle on S.

Finally, because the descent data respects the closed immersion into the projective bundle, we can apply the fpqc descent of closed subschemes (Corollary 2.2.19). This guarantees the existence of a unique closed subscheme $\mathcal{C} \subset \mathbb{P}(E)$ over S yielding the following Cartesian diagrams:

$$\begin{tikzcd}
 & & & & \mathbb{P}(E') \arrow[r] \arrow[ddl] & \mathbb{P}(E) \arrow[ddl] \\
p_1^*\mathcal{C}' \arrow[rr, "\alpha"] \arrow[rrr, bend left] \arrow[dr] & & p_2^*\mathcal{C}' \arrow[r] \arrow[dl] & \mathcal{C}' \arrow[r, dashed] \arrow[d] \arrow[ur, hook] & \mathcal{C} \arrow[d, dashed] \arrow[ur, dashed, hook] & \\
 & S' \times_S S' \arrow[rr, shift left, "p_1"] \arrow[rr, shift right, "p_2"'] & & S' \arrow[r] & S &
\end{tikzcd}$$

Because $\mathcal{C}' \to S'$ is a family of smooth curves, and fpqc local properties like smoothness and properness descend along fpqc covers, we conclude that the descended object $\mathcal{C} \to S$ is the required family of smooth curves, since every geometric fiber of $\mathcal{C} \to S$ is identified with a geometric fiber of $\mathcal{C}' \to S'$. $\square$

Just as with the sheafification procedure in classical sheaf theory, we can construct a stackification procedure for prestacks defined via a universal property. More precisely, we have the following foundational result:

Proposition 2.2.22 (Stackification). *Let $\mathcal{X}$ be a prestack over a site $\mathcal{S}$. There exists a stack $\mathcal{X}^{\mathrm{st}}$, called the stackification of $\mathcal{X}$, along with a morphism of prestacks $\mathcal{X} \to \mathcal{X}^{\mathrm{st}}$ such that for every stack $\mathcal{Y}$ over $\mathcal{S}$, the induced functor*

$$\mathrm{Hom}(\mathcal{X}^{\mathrm{st}}, \mathcal{Y}) \xrightarrow{\simeq} \mathrm{Hom}(\mathcal{X}, \mathcal{Y})$$

is an equivalence of categories.

The construction is analogous to the classical sheafification process. We refer the reader to Alper's notes [**Alp26**, Theorem 3.5.18] for the full proof. Instead of presenting the abstract construction in its entirety, we will use a fundamental example—the construction of the quotient stack via the stackification of the quotient prestack (Example 2.2.6)—to illustrate how this procedure operates in practice.

Recall that given a base scheme T and a smooth group scheme G, a *principal G-bundle* (or a *G-torsor*) is a scheme P equipped with a smooth morphism $P \to T$ and a right action of G on P over T, such that the shear map

$$\begin{aligned} P \times G &\xrightarrow{\cong} P \times_T P \\ (p, g) &\longmapsto (p, p \cdot g) \end{aligned}$$

is an isomorphism.

Example 2.2.23 (Quotient stack in the fpqc topology)**.** In this example, we construct the quotient stack $[X/G]$ by taking the stackification of the quotient prestack $[X/G]^{\mathrm{pre}}$, where G is a smooth group scheme equipped with a right action on a scheme X. Recall from our earlier discussion that for each test scheme $T \in \mathsf{Sch}$, the fiber category $[X/G]^{\mathrm{pre}}(T)$ has the T-valued points $X(T)$ as its objects, and the group elements $G(T)$ acting on these points as its morphisms.

To stackify this prestack over the fpqc topology, we first verify that morphisms glue, and then formally adjoin objects that correspond to effective descent data.

First, we show that morphisms glue. Given two objects $x, y \in X(T)$ and an fpqc cover $\{T_i \to T\}$, a local morphism from $x|_{T_i}$ to $y|_{T_i}$ is an element $g_i \in G(T_i)$ such that $x|_{T_i} \cdot g_i = y|_{T_i}$. If these local morphisms agree on the overlaps T_{ij} (i.e., $g_i|_{T_{ij}} = g_j|_{T_{ij}}$), they glue to a unique global morphism $g \in G(T)$. This holds because G is a scheme, and by Grothendieck's fpqc descent theorem discussed previously, the representable functor $h_G = G(-)$ is a sheaf in the fpqc (and thus fppf) topology. Therefore, the prestack $[X/G]^{\mathrm{pre}}$ satisfies the first stack axiom.

Next, to satisfy the second stack axiom (objects glue), we formally adjoin effective descent data. Let $\{T_i \to T\}$ be an fpqc covering of a scheme T. A descent datum for $[X/G]^{\mathrm{pre}}$ with respect to this cover consists of local objects $x_i \in X(T_i)$ and local isomorphisms $g_{ij} \in G(T_i \times_T T_j)$ satisfying $x_j = x_i \cdot g_{ij}$ on the double overlaps T_{ij}, subject to the cocycle condition $g_{ik} = g_{ij} \cdot g_{jk}$ on the triple overlaps T_{ijk}.

The collection of group elements $\{g_{ij}\}$ satisfying the cocycle condition is precisely the Čech data required to construct a principal G-bundle $P \to T$ that is locally trivialized over the cover $\{T_i \to T\}$. Furthermore, the local maps $x_i \colon T_i \to X$ are compatible with the transition functions g_{ij} via the group action. This compatibility ensures that the local maps x_i glue together to form a globally defined, G-equivariant morphism $f \colon P \to X$.

Therefore, the stackification process forces the objects of the new category $[X/G](T)$ over T to be exactly these glued geometric structures. Specifically, an object of the quotient stack $[X/G]$ over a scheme T is a pair (P, f), where:

(1) $P \to T$ is a principal G-bundle over T.
(2) $f \colon P \to X$ is a G-equivariant morphism.

A morphism between two such objects (P_1, f_1) over T_1 and (P_2, f_2) over T_2 is a Cartesian diagram of principal G-bundles that is compatible with the equivariant maps to X.

The quotient stack framework can be generalized further if we drop the condition of equipping Sch exclusively with the fpqc topology. In fact, it is standard practice to *define* the objects of $[X/G]$ over T directly as the pairs (P, f) described in the previous example, regardless of the underlying topology. However, if one works in other topologies, Grothendieck's descent theorem may fail. In such cases, defining $[X/G]$ directly via principal bundles may no longer coincide with the stackification of the prestack $[X/G]^{\mathrm{pre}}$.

2.3. From Schemes to Stacks: Geometric Foundations

2.3.1. Algebraic Spaces, Deligne–Mumford Stacks, and Algebraic Stacks. In the preceding sections, we successfully generalized the concepts of presheaves and sheaves to prestacks and stacks. However, abstract stacks alone are purely categorical constructs. To draw a classical analogy: if we restrict our study only to arbitrary locally ringed spaces, we lack the geometric rigidity required to do meaningful algebraic geometry. It is only by requiring locally ringed spaces to be locally isomorphic to affine schemes that we obtain the rich geometry in the category of schemes. Similarly, to import the fruitful geometric machinery of classical schemes into the realm of stacks, we need to introduce the notion of an "atlas" or "chart" for a given stack. This allows us to bridge the abstract categorical world and the concrete geometric world.

Definition 2.3.1. Let $\mathsf{Sch}_{\text{ét}}$ be the category of schemes equipped with the big étale topology (Example 2.2.10).

(1) An *algebraic space* is a *sheaf* X on $\mathsf{Sch}_{\text{ét}}$ such that there exists a scheme U alongside a surjective, étale morphism $U \to X$ which is representable by schemes.
(2) A *Deligne–Mumford stack* is a stack $\mathcal{X}$ on $\mathsf{Sch}_{\text{ét}}$ such that there exists a scheme U alongside a surjective, étale, and representable morphism $U \to \mathcal{X}$.
(3) An *algebraic stack* (or *Artin stack*) is a stack $\mathcal{X}$ on $\mathsf{Sch}_{\text{ét}}$ such that there exists a scheme U alongside a surjective, smooth, and representable morphism $U \to \mathcal{X}$.

In all cases, the scheme U is referred to as an *atlas* (or presentation) for the stack.

We say a morphism of stacks $\mathcal{X} \to \mathcal{Y}$ over $\mathcal{S}$ is *representable by schemes* (resp. *representable*) if, for any morphism $T \to \mathcal{Y}$ from a scheme T, the fiber product $\mathcal{X} \times_{\mathcal{Y}} T$ is equivalent to a scheme (resp. an algebraic space).

If $\mathcal{P}$ is a property of morphisms of schemes that is stable under base change and local on the source in the étale (or smooth) topology, the representability condition allows us to transfer property $\mathcal{P}$ to a morphism whose target is an arbitrary stack. We will expand on the precise mechanics of transferring properties from the classical scheme world to the stacky world in the subsequent subsections.

We first examine the most fundamental example: the algebraicity of quotient stacks.

Theorem 2.3.2 (Algebraicity of quotient stacks)**.** *Let G be a smooth group scheme acting on a scheme X. The quotient stack $[X/G]$ is an algebraic stack, and the canonical projection $\pi\colon X \to [X/G]$ is a principal G-bundle.*

Furthermore, suppose G is a finite abstract group, viewed as a group scheme. If G acts freely on X, then the quotient stack $[X/G]$ is equivalent to an algebraic space, denoted by X/G.

Proof. To demonstrate that the canonical projection $\pi\colon X \to [X/G]$ is a principal G-bundle, consider an arbitrary scheme T and a morphism $T \to [X/G]$. By the definition of the quotient stack (see Example 2.2.23), this morphism is classified by a principal G-bundle $P \to T$ together with a G-equivariant map $P \to X$. This forms a Cartesian diagram:

$$\begin{tikzcd} P \arrow[r] \arrow[d] \arrow[dr, phantom, "\square"] & X \arrow[d, "\pi"] \\ T \arrow[r] & {[X/G]} \end{tikzcd}$$

By definition, a principal G-bundle $P \to T$ is smooth as a morphism. Moreover, since G is assumed to be a smooth group scheme, the bundle $P \to T$ is a smooth and surjective morphism of schemes. This proves that $\pi\colon X \to [X/G]$ is representable, smooth, and surjective, thus serving as a smooth atlas. Consequently, $[X/G]$ is an algebraic stack.

For the second statement, assume G is a finite group acting freely on X. A free action implies that the objects in the fibers of $[X/G]$ possess no non-trivial automorphisms. Therefore, the stack $[X/G]$ is equivalent to a sheaf, which we denote by X/G. Since G is a finite constant group scheme, the principal G-bundle $X \to X/G$ is finite, étale, and surjective. This provides an étale atlas for the sheaf X/G, proving it is an algebraic space. □

At this stage, it is appropriate to introduce a generalization of quotient stacks and provide a different, more concrete perspective on algebraic spaces and stacks: the language of equivalence relations and groupoids.

We start with the following observation. Let $\mathcal{X}$ be an algebraic stack and $p\colon U \to \mathcal{X}$ be a smooth presentation. If we assume that p is representable by schemes, then the fiber product $R \coloneqq U \times_{\mathcal{X}} U$ is a scheme. By the definition of fiber products in 2-categories, a T-valued point of R is a triple (u_1, u_2, α) where $\alpha\colon p(u_1) \xrightarrow{\cong} p(u_2)$ is an isomorphism in $\mathcal{X}(T)$. If we let $s = p_1\colon R \to U$ (the projection to the first factor) be the source morphism, and let $t = p_2\colon R \to U$ be the target morphism, this data defines a smooth groupoid in schemes $\mathcal{R} \coloneqq (R \rightrightarrows U)$. The composition $R \times_{s,U,t} R \to R$ is defined by composing the isomorphisms $(\alpha_2\colon p(u_2) \to p(u_3), \alpha_1\colon p(u_1) \to p(u_2)) \mapsto \alpha_2 \circ \alpha_1$. The identity morphism is given by the relative diagonal map $U \to U \times_{\mathcal{X}} U$, and the inverse map $R \to R$ is defined by swapping the factors and inverting the isomorphism α. If we drop the assumption that the smooth presentation is representable by schemes, this identical construction yields a smooth groupoid of algebraic spaces.

Therefore, any algebraic stack is associated with a smooth groupoid, dependent on the choice of the smooth presentation. What is remarkably powerful is that we can use this smooth groupoid of algebraic spaces to entirely recover the algebraic stack itself. To formalize this, we introduce the following generalization of quotient stacks.

Definition 2.3.3 (Quotients by groupoids)**.** Let $s, t\colon R \rightrightarrows U$ be a smooth groupoid of algebraic spaces. Define $[U/R]^{\mathrm{pre}}$ as the prestack whose objects over a scheme $T \in \mathsf{Sch}_{\text{ét}}$ are the objects $U(T)$, and whose morphisms over T are given by $R(T)$. We define the quotient stack $[U/R]$ to be the stackification of $[U/R]^{\mathrm{pre}}$ with respect to the big étale topology on $\mathsf{Sch}_{\text{ét}}$.

Following the argument in Example 2.2.23, we can explicitly describe the objects of this stackification by introducing principal $\mathcal{R}$-bundles.

Definition 2.3.4 (Principal $\mathcal{R}$-bundles)**.** Let $\mathcal{R} \coloneqq (R \rightrightarrows U)$ be a smooth groupoid of algebraic spaces.

(1) Let T be a scheme. A *principal $\mathcal{R}$-bundle* over T consists of an algebraic space P equipped with a smooth and surjective morphism $p\colon P \to T$, together with a morphism of groupoids

$[P \times_T P \rightrightarrows P] \longrightarrow [R \rightrightarrows U]$. This groupoid morphism consists of an anchor map $f\colon P \to U$ and a morphism $\tau\colon P \times_T P \to R$ such that the following natural square is Cartesian:

$$\begin{tikzcd} P \times_T P \arrow[r, "p_2"] \arrow[d, "\tau"'] \arrow[dr, phantom, "\square"] & P \arrow[d, "f"] \\ R \arrow[r, "s"'] & U \end{tikzcd}$$

(2) The quotient stack $[U/R]$ is a category fibered in groupoids over $\mathsf{Sch}_{\text{ét}}$. For any scheme T, the objects of the groupoid $[U/R](T)$ are diagrams:

$$\begin{tikzcd} P \arrow[r, "f"] \arrow[d, "p"'] & U \\ T & \end{tikzcd}$$

where $p\colon P \to T$ is a principal $\mathcal{R}$-bundle and f is its associated anchor map.

(3) The morphisms from $(P_1 \xrightarrow{p_1} T_1, f_1)$ to $(P_2 \xrightarrow{p_2} T_2, f_2)$ in this fibered category consist of a pair of scheme morphisms

$$u\colon T_1 \to T_2, \qquad \varphi\colon P_1 \to P_2$$

such that $p_2 \circ \varphi = u \circ p_1$ and $f_2 \circ \varphi = f_1$, and the resulting square

$$\begin{tikzcd} P_1 \arrow[r, "\varphi"] \arrow[d, "p_1"'] \arrow[dr, phantom, "\square"] & P_2 \arrow[d, "p_2"] \\ T_1 \arrow[r, "u"'] & T_2 \end{tikzcd}$$

is Cartesian (so φ is a pullback morphism of principal $\mathcal{R}$-bundles over u).

Remark 2.3.5 (Equivalence of definitions)**.** In the differential geometric setting [**BX11**, Definition 2.18], a principal bundle for a Lie groupoid $R \rightrightarrows U$ over a base manifold T is defined as a surjective submersion $P \to T$ equipped with a principal right groupoid action of R on P. The principality condition requires that for any two points $p, p' \in P$ in the same fiber over T, there exists a unique morphism $\gamma \in R$ such that $p \cdot \gamma = p'$.

The space of pairs (p, p') lying in the same fiber over the base is the fiber product $P \times_T P$. The space of valid action pairs (p, γ) is given by the fiber product $P \times_{f,U,s} R$, where $f\colon P \to U$ is the associated anchor map. Requiring the structural diagram of the bundle to be Cartesian is equivalent to demanding that the action map $P \times_{f,U,s} R \to P \times_T P$, defined by $(p, \gamma) \mapsto (p, p \cdot \gamma)$, is an isomorphism of algebraic spaces. Our definition above follows closely the version in [**Noo12**], and we will freely interchange these definitions throughout this thesis since they are equivalent.

Now we are able to establish the fundamental reconstruction theorem, linking the abstract stack definition back to the concrete groupoid presentation.

Theorem 2.3.6 (Reconstruction by groupoids)**.**

(1) Let $\mathcal{R} := (R \rightrightarrows U)$ be an étale (resp. smooth) groupoid of algebraic spaces. Then the quotient stack by the groupoid $[U/R]$ is a Deligne–Mumford stack (resp. algebraic stack), and the canonical projection $U \to [U/R]$ is an étale (resp. smooth) presentation.

(2) *Let $\mathcal{X}$ be a Deligne–Mumford stack (resp. algebraic stack) and $U \to \mathcal{X}$ be an étale (resp. smooth) presentation. Then there is an isomorphism of stacks $\mathcal{X} \cong [U/R]$, where $\mathcal{R} := (R \rightrightarrows U)$ is the étale (resp. smooth) groupoid of algebraic spaces defined by the fiber product $R := U \times_{\mathcal{X}} U$.*

Proof. For the first part, we will show that $U \to [U/R]$ is surjective, smooth and representable. Since $[U/R]$ is defined as the stackification of $[U/R]^{\mathrm{pre}}$, there exists an étale cover $T' \to T$ and a commutative diagram

$$\begin{tikzcd} T' \arrow[r] \arrow[d] & U \arrow[d] \\ T \arrow[r] & {[U/R]} \end{tikzcd}$$

Now consider the following commutative cube of stacks:

$$\begin{tikzcd}[row sep=small, column sep=small] & U_{T'} \arrow[rr] \arrow[dl] \arrow[dd] & & T' \arrow[dl] \arrow[dd] \\ R \arrow[rr, crossing over] \arrow[dd] & & U & \\ & U_T \arrow[rr] \arrow[dl] & & T \arrow[dl] \\ U \arrow[rr] & & {[U/R]} \arrow[from=uu, crossing over] & \end{tikzcd}$$

In this cube, the front, back, top, and bottom squares are cartesian. It suffices to show that the sheaf U_T is an algebraic space. Note that $U_{T'}$ is a scheme and $U_{T'} \to U_T$ is an étale surjective morphism, representable by schemes since it is the base change of $T' \to T$. Since $R \to U$ is surjective and étale (resp. smooth), so is $U_{T'} \to T'$. By étale descent, $U_T \to T$ is also, which concludes the first part.

For the second part, Our proof follows closely the proof in [**Sta26**, Tag 04T3]. We construct a canonical isomorphism of stacks $[U/R] \xrightarrow{\cong} \mathcal{X}$. The presentation $p\colon U \to \mathcal{X}$ naturally induces a map as follows. By the definition of the fiber product $R := U \times_{\mathcal{X}} U$, the two compositions $R \rightrightarrows U \xrightarrow{p} \mathcal{X}$ are canonically 2-isomorphic, which induces a functor $[U/R]^{\mathrm{pre}} \to \mathcal{X}$. This functor uniquely lifts to a functor from the stackification $F\colon [U/R] \to \mathcal{X}$. We claim that F is an isomorphism.

We first show that F induces an isomorphism of isomorphism sheaves $\underline{\mathrm{Isom}}_{[U/R]}(x, y) \to \underline{\mathrm{Isom}}_{\mathcal{X}}(F(x), F(y))$ for any test scheme T and any two objects $x, y \in [U/R](T)$. Since this is a local property on T, we may pass to a smooth (resp. étale) cover $T' \to T$ where x and y lift to objects $x', y' \in U(T')$. In the prestack $[U/R]^{\mathrm{pre}}$, the space of morphisms between x' and y' is by definition $R(T')$. However, by the construction of the groupoid $R = U \times_{\mathcal{X}} U$, we have $R(T') = U(T') \times_{\mathcal{X}(T')} U(T')$, which is the collection of isomorphisms between $p(x')$ and $p(y')$ in $\mathcal{X}(T')$. Therefore, the functor is fully faithful.

Let T be a scheme and let $t \in \mathcal{X}(T)$ be an object over T, corresponding to a morphism $t\colon T \to \mathcal{X}$. Consider the fiber product $T' := T \times_{\mathcal{X}} U$. Since $p\colon U \to \mathcal{X}$ is a smooth (resp. étale) surjective presentation, the projection $T' \to T$ is a smooth (resp. étale) surjective cover of schemes. The other projection $T' \to U$ defines an object in $U(T')$, which yields an object in $[U/R](T')$. By construction, the image of this object under F is canonically isomorphic to the pullback of t to T'. Since $[U/R]$

and $\mathcal{X}$ are both stacks, effective descent guarantees that this object descends back to an object in $[U/R](T)$ that maps identically to t. Therefore, F is essentially surjective. $\square$

We have now established a surjective mapping from the collection of smooth groupoids of algebraic spaces to the collection of algebraic stacks, mapping $\mathcal{R} := (R \rightrightarrows U)$ to the quotient stack $[U/R]$. The question naturally arises whether this induces a true equivalence between these collections. Trivially, they are not: choosing different smooth presentations for the same stack yields different groupoids. To rectify this, we introduce an equivalence relation between groupoids that accounts for changes in presentation. This is known as *Morita equivalence*. Motivated by the above discussion, two groupoids are Morita equivalent if and only if their corresponding fibered categories of principal bundles over any test scheme are equivalent. The following definition generalizes the definition in [**BX11**, Definition 2.24]

Definition 2.3.7 (Morita morphisms and Morita equivalence)**.**

(1) Let $\mathcal{R}_i := (R_i \rightrightarrows U_i)$ for $i = 1, 2$ be smooth (resp. étale) groupoids of algebraic spaces. A morphism of groupoids $\varphi\colon \mathcal{R}_1 \to \mathcal{R}_2$ (consisting of an anchor map $f\colon U_1 \to U_2$ and a compatible map $R_1 \to R_2$) is called a *Morita morphism* (or an equivalence of groupoids) if it satisfies two conditions:
 - The following commutative diagram is Cartesian:
$$\begin{tikzcd} R_1 \arrow[r, "(s_1,t_1)"] \arrow[d] \arrow[dr, phantom, "\square"] & U_1 \times U_1 \arrow[d, "f\times f"] \\ R_2 \arrow[r, "(s_2,t_2)"] & U_2 \times U_2 \end{tikzcd}$$
 - The natural morphism $t_2 \circ \mathrm{pr}_2\colon U_1 \times_{f,U_2,s_2} R_2 \to U_2$ is a smooth (resp. étale) and surjective morphism of algebraic spaces.

(2) Two smooth (resp. étale) groupoids of algebraic spaces $\mathcal{R}_1$ and $\mathcal{R}_2$ are called *Morita equivalent* if there exists a third smooth (resp. étale) groupoid of algebraic spaces $\mathcal{R}_3$ and Morita morphisms $\mathcal{R}_3 \to \mathcal{R}_1$ and $\mathcal{R}_3 \to \mathcal{R}_2$.

It turns out that there is a bijection between Morita equivalence classes of smooth (resp. étale) groupoids of algebraic spaces and isomorphism classes of algebraic stacks (resp. Deligne–Mumford stacks). The following theorem, although well-known among experts, is rarely explicitly detailed in standard algebraic geometry literature. We adapt the statement and proof of this equivalence theorem from the differential category, following [**BX11**, Theorem 2.26].

Theorem 2.3.8. *Let $\mathcal{R}_i := (R_i \rightrightarrows U_i)$ for $i = 1, 2$ be smooth (resp. étale) groupoids of algebraic spaces, and let $\mathcal{X}_i := [U_i/R_i]$ be their associated algebraic stacks (resp. Deligne–Mumford stacks). The following conditions are equivalent:*

(1) The stacks $\mathcal{X}_1$ and $\mathcal{X}_2$ are isomorphic.

(2) The groupoids $\mathcal{R}_1$ and $\mathcal{R}_2$ are Morita equivalent.

(3) There exists an algebraic space T equipped with two smooth (resp. étale) surjective morphisms $\pi_1\colon T \to U_1$ and $\pi_2\colon T \to U_2$, along with commuting left and right actions of R_1 and R_2 respectively, such that T is simultaneously a principal $\mathcal{R}_1$-bundle over U_2 (via π_2) and a principal $\mathcal{R}_2$-bundle over U_1 (via π_1). We refer to such an algebraic space T as a principal $\mathcal{R}_1 - \mathcal{R}_2$-bibundle.

Sketch of Proof. We prove the equivalence by establishing (1) $\implies$ (3) $\implies$ (2) $\implies$ (3) $\implies$ (1).

(1) $\implies$ (3): Suppose the stacks $\mathcal{X}_1 \coloneqq [U_1/R_1]$ and $\mathcal{X}_2 \coloneqq [U_2/R_2]$ are isomorphic. Let us identify them via this isomorphism and call the common stack $\mathcal{X}$. We have two smooth (resp. étale) surjective presentations $U_1 \to \mathcal{X}$ and $U_2 \to \mathcal{X}$. Define T to be the fiber product $T \coloneqq U_1 \times_{\mathcal{X}} U_2$. Clearly T is an algebraic space, and the natural projections $\pi_1 \colon T \to U_1$ and $\pi_2 \colon T \to U_2$ are smooth (resp. étale) and surjective. One can naturally pullback the groupoid actions of R_1 on U_1 and R_2 on U_2 to T. Because these actions commute, T becomes a principal $\mathcal{R}_1$-bundle over U_2 and a principal $\mathcal{R}_2$-bundle over U_1, making T a principal bibundle.

(3) $\implies$ (2): Choose a principal $\mathcal{R}_1 - \mathcal{R}_2$-bibundle T. We construct a third groupoid $\mathcal{R}_T$ that maps via Morita morphisms to both $\mathcal{R}_1$ and $\mathcal{R}_2$. Let the object space of $\mathcal{R}_T$ be T, and let the morphism space be $R_T \coloneqq R_1 \times_{s_1,U_1,\pi_1} T \times_{\pi_2,U_2,t_2} R_2$. There is a canonical way to define the composition making $\mathcal{R}_T \coloneqq (R_T \rightrightarrows T)$ a smooth (resp. étale) groupoid of algebraic spaces. The natural projection maps to $\mathcal{R}_1$ and $\mathcal{R}_2$ are Morita morphisms because their anchor maps (π_1 and π_2) are the smooth (resp. étale) surjective projections of the bibundle, and the principal bundle structures guarantee the required Cartesian conditions.

(2) $\implies$ (3): We establish this in two steps. First, if $\varphi \colon \mathcal{R}_1 \to \mathcal{R}_2$ is a Morita morphism with anchor $f \colon U_1 \to U_2$, then the algebraic space $Q \coloneqq U_1 \times_{f,U_2,s_2} R_2$ naturally inherits the structure of a principal $\mathcal{R}_1 - \mathcal{R}_2$-bibundle. Second, if Q is a principal $\mathcal{R}_1 - \mathcal{R}_2$-bibundle and Q' is a principal $\mathcal{R}_2 - \mathcal{R}_3$-bibundle, we form their quotient $(Q \times_{U_2} Q')/R_2$. Since the action of R_2 is free, the action groupoid $(Q \times_{U_2} Q') \times_{U_2} R_2 \rightrightarrows Q \times_{U_2} Q'$ is a smooth (resp. étale) equivalence relation. By a result on quotients of algebraic spaces ([**Alp26**, Corollary 5.5.12]), the quotient by a smooth equivalence relation exists as an algebraic space, and it naturally forms a principal $\mathcal{R}_1 - \mathcal{R}_3$-bibundle. Since Morita equivalence provides Morita morphisms $\mathcal{R}_3 \to \mathcal{R}_1$ and $\mathcal{R}_3 \to \mathcal{R}_2$, we obtain bibundles $\mathcal{R}_1 - \mathcal{R}_3$ and $\mathcal{R}_3 - \mathcal{R}_2$, whose contracted product yields the desired $\mathcal{R}_1 - \mathcal{R}_2$-bibundle.

(3) $\implies$ (1): We prove that the stacks are isomorphic by showing that their corresponding fibered categories of principal bundles over $\mathsf{Sch}_{\text{ét}}$ are equivalent. Let S be an arbitrary scheme and $P \to S$ be a principal $\mathcal{R}_1$-bundle over S, with anchor map $h \colon P \to U_1$. We define a new space $P' \coloneqq (T \times_{\pi_1,U_1,h} P)/R_1$. As before, since T is a principal $\mathcal{R}_1$-bundle over U_2, the action of R_1 on the fiber product is free, defining a smooth equivalence relation. By [**Alp26**, Corollary 5.5.12], the quotient P' exists as an algebraic space. The space P' canonically inherits the structure of a principal $\mathcal{R}_2$-bundle over S, with the anchor map to U_2 induced by π_2. Furthermore, a morphism of principal $\mathcal{R}_1$-bundles naturally induces a morphism of the contracted principal $\mathcal{R}_2$-bundles. This construction provides a functor from the category of principal $\mathcal{R}_1$-bundles to the category of principal $\mathcal{R}_2$-bundles. Utilizing the symmetric structure of the bibundle to construct a quasi-inverse, one easily checks this functor is an equivalence of categories. Thus, the stacks $\mathcal{X}_1$ and $\mathcal{X}_2$ are isomorphic. $\square$

The outline of this proof closely follows the analogous result for Lie groupoids and differentiable stacks. However, the underlying technical details diverge significantly. The existence of quotient spaces is more rigid in algebraic geometry than in differential geometry. Whereas taking the quotient of a manifold by a free and proper Lie group action yields a smooth manifold, guaranteeing that a quotient remains within the category of algebraic spaces requires heavier descent machinery. Consequently, our algebraic proof cannot be entirely self-contained; it needs to invoke the foundational result that the quotient of an algebraic space by a smooth (resp. étale) equivalence relation is still an algebraic space.

The bijection between Morita equivalence classes of groupoids and isomorphism classes of stacks provides a powerful computational perspective. Rather than working exclusively with abstract categories fibered in groupoids, one can utilize a groupoid presentation $\mathcal{R} \coloneqq (R \rightrightarrows U)$ to define and compute algebraic invariants. As we will explore in the final section of this chapter, this groupoid machinery is the essential vehicle for developing cohomology theory for stacks, including the singular and de Rham cohomology for topological and differentiable stacks.

Having established this theoretical framework for algebraic stacks, we now return to our investigation of fundamental geometric examples. Equipped with these tools, the next candidate to examine is the moduli stack of smooth curves of genus g, denoted $\mathcal{M}_g$ (see Proposition 2.2.21).

Theorem 2.3.9 (Algebraicity of $\mathcal{M}_g$)**.** *If $g \geq 2$, then $\mathcal{M}_g$ is an algebraic stack over $\operatorname{Spec}\mathbb{Z}$. Moreover, there is an isomorphism $\mathcal{M}_g \cong [H'/\operatorname{PGL}_{5g-5}]$, where H' is a locally closed subscheme of a suitable Hilbert scheme H.*

Sketch of Proof. We outline the core ideas behind this theorem and refer the reader to [**Alp26**, Theorem 4.1.17] for the details. Recall that for a family of smooth curves $\pi\colon \mathcal{C} \to S$, the tri-canonical sheaf $\Omega^{\otimes 3}_{\mathcal{C}/S}$ is very ample relative to S. If $g \geq 2$, the pushforward $E \coloneqq \pi_*\left(\Omega^{\otimes 3}_{\mathcal{C}/S}\right)$ is a vector bundle of rank $5g-5$ (Proposition 2.2.20). Consequently, this very ample sheaf defines a closed immersion $\mathcal{C} \hookrightarrow \mathbb{P}(E)$ over S.

We can compute the Hilbert polynomial of a geometric fiber $\mathcal{C}_s \hookrightarrow \mathbb{P}^{5g-6}_{\kappa(s)}$ using the Riemann–Roch theorem:

$$P(n) \coloneqq \chi\left(\mathcal{O}_{\mathcal{C}_s}(n)\right) = \deg\left(\Omega^{\otimes 3n}_{\mathcal{C}_s}\right) + 1 - g = (6n-1)(g-1).$$

Therefore, we define the base Hilbert scheme as

$$H \coloneqq \operatorname{Hilb}^P\left(\mathbb{P}^{5g-6}_{\mathbb{Z}}\right).$$

The analysis above shows that there is a subscheme $H' \subset H$ parameterizing smooth families of tri-canonically embedded curves, which will provide our smooth presentation $H' \to \mathcal{M}_g$. More precisely, if $\mathcal{U} \hookrightarrow \mathbb{P}^{5g-6} \times H$ is the universal family over H and $p\colon \mathcal{U} \to H$ is the projection, then H' is the locally closed subscheme of H uniquely characterized by the following properties:

(1) For each $h \in H'$, the fiber $\mathcal{U}_h \to \operatorname{Spec}\kappa(h)$ is smooth and geometrically connected.
(2) The global section map $\mathrm{H}^0\left(\mathbb{P}^{5g-6}_{\kappa(h)}, \mathcal{O}(1)\right) \to \mathrm{H}^0\left(\mathcal{U}_h, \mathcal{O}_{\mathcal{U}_h}(1)\right)$ is an isomorphism for all $h \in H'$.
(3) The line bundles $\Omega^{\otimes 3}_{\mathcal{U}_{H'}/H'}$ and $\mathcal{O}_{\mathcal{U}_{H'}}(1)$ differ only by the pullback of a line bundle from H'.
(4) (*Universal Property*) If $T \to H$ is a morphism of schemes such that the above three properties hold for the pullback family $\mathcal{U}_T \to T$, then the morphism $T \to H$ factors uniquely through H'.

These properties ensure that for all $h \in H'$, the curve $\mathcal{U}_h \hookrightarrow \mathbb{P}^{5g-6}_{\kappa(h)}$ is embedded exactly by the complete linear series of $\Omega^{\otimes 3}_{\mathcal{U}_h/\kappa(h)}$.

Having characterized the parameter space H', we obtain a morphism $H' \to \mathcal{M}_g$ by forgetting the embedding. The locally closed subscheme H' is invariant under the natural action of $\operatorname{PGL}_{5g-5}$, as $\operatorname{PGL}_{5g-5} = \underline{\operatorname{Aut}}\left(\mathbb{P}^{5g-6}\right)$ acts by changing the coordinates of the ambient projective space. Therefore, the morphism $H' \to \mathcal{M}_g$ descends to a morphism from the quotient stack $[H'/\operatorname{PGL}_{5g-5}] \to \mathcal{M}_g$ by the universal property of stackification (Proposition 2.2.22).

To prove this forgetful morphism is an isomorphism, we need to show that the fiber categories $\mathcal{M}_g(T)$ and $[H'/\operatorname{PGL}_{5g-5}](T)$ are equivalent for every scheme T. We do this by explicitly constructing the inverse functor. Let $\pi\colon \mathcal{C} \to T$ be a family of smooth curves in $\mathcal{M}_g(T)$. By Proposition 2.2.20, the pushforward $E = \pi_*\Omega^{\otimes 3}_{\mathcal{C}/T}$ is a vector bundle of rank $5g-5$, and we have a closed immersion

$\mathcal{C} \hookrightarrow \mathbb{P}(E)$. To map this data into the quotient stack, we construct a principal PGL_{5g-5}-bundle over T along with an equivariant map to H' (Example 2.2.23). We construct this as the projective frame bundle of E, defined as $P := \underline{\mathrm{Isom}}(\mathcal{O}_T^{\oplus 5g-5}, E)/\mathbb{G}_m$ (See [**Alp26**, Exercise B.1.57] for more details). By definition, the canonical projection $q\colon P \to T$ is a PGL_{5g-5}-torsor. Pulling the family $\mathcal{C}$ back along q trivializes the projective bundle $q^*\mathbb{P}(E) \cong \mathbb{P}_P^{5g-6}$. This gives us a family of curves $\mathcal{C} \times_T P \to P$ equipped with a closed immersion into $\mathbb{P}_P^{5g-6}$. This explicit embedding is classified by a morphism $P \to H$, which factors through H' because our tri-canonical embedding satisfies all the geometric criteria, and from the construction it is clear that the resulting map $P \to H'$ is PGL_{5g-5}-equivariant. This data $(P, P \to H')$ defines an object in $[H'/\mathrm{PGL}_{5g-5}](T)$, concluding the construction of the inverse and proving the algebraicity of $\mathcal{M}_g$. □

We can actually go further from this model by introducing marked points on a family of smooth curves. An *n-pointed family of smooth curves* of genus g is a smooth and proper morphism $\pi\colon \mathcal{C} \to S$ of schemes equipped with n sections $\sigma_i\colon S \to \mathcal{C}$ such that for each geometric point $s\colon \operatorname{Spec} k \to S$, the fiber $\mathcal{C}_s$ is a smooth curve of genus g, and the points $\sigma_1(s), \ldots, \sigma_n(s) \in \mathcal{C}_s$ are pairwise distinct. Let $\mathcal{M}_{g,n}$ denote the prestack over $\mathsf{Sch}_{\text{ét}}$ parameterizing n-pointed families of smooth curves of genus g.

A natural question arises: under what conditions does the moduli prestack $\mathcal{M}_{g,n}$ form a stack, or more restrictively, an algebraic stack over $\mathsf{Sch}_{\text{ét}}$? In Theorem 2.3.9, we established the fundamental result for the unpointed, higher-genus case: when $g \geq 2$ and $n = 0$, $\mathcal{M}_g$ is indeed an algebraic stack. What remains is to systematically examine the cases of low genus, specifically $g = 0$ and $g = 1$. We begin by detailing the consequences for genus zero.

Example 2.3.10 (Moduli space of rational curves)**.** The structure of the moduli space $\mathcal{M}_{0,n}$ depends entirely on the stabilizer of n distinct points on $\mathbb{P}^1$ under the action of the automorphism group PGL_2. We have the following classification:

$$\mathcal{M}_{0,n} \cong \begin{cases} \mathsf{B}\mathrm{PGL}_2, & n = 0; \\ \mathsf{B}B, & n = 1; \\ \mathsf{B}\mathbb{G}_m, & n = 2; \\ \operatorname{Spec} \mathbb{Z}, & n = 3; \\ \left(\mathbb{P}^1 \setminus \{0, 1, \infty\}\right)^{n-3} \setminus \Delta, & n \geq 4. \end{cases}$$

Here, $B \subset \mathrm{PGL}_2$ is the two-dimensional Borel subgroup of upper triangular matrices (the stabilizer of a single point, say ∞), and Δ is the large diagonal closed subscheme where at least two of the $n - 3$ points coincide.

The cases for $n \geq 3$ can be deduced through a geometric argument utilizing the 3-transitive nature of the PGL_2-action. For $n = 3$, given any smooth rational curve with three distinct marked points (C, p_1, p_2, p_3), there exists a unique isomorphism to $(\mathbb{P}^1, 0, 1, \infty)$. Therefore $\mathcal{M}_{0,3}$ is isomorphic to $\operatorname{Spec} \mathbb{Z}$. For $n = 4$, the first three points of any smooth rational curve (C, p_1, p_2, p_3, p_4) can be uniquely mapped to $0, 1$, and ∞. The fourth point is then mapped to some coordinate $q \in \mathbb{P}^1 \setminus \{0, 1, \infty\}$. This coordinate q corresponds to the cross-ratio of the four points, providing an isomorphism between $\mathcal{M}_{0,4}$ and $\mathbb{P}^1 \setminus \{0, 1, \infty\}$. This argument generalizes to any $n \geq 4$, mapping the remaining $n - 3$ points into $\mathbb{P}^1 \setminus \{0, 1, \infty\}$ while avoiding the diagonals.

The genus 1 case, however, is much more involved due to the geometry of the canonical bundle.

Example 2.3.11 (Moduli space of elliptic curves)**.** We first consider the unpointed prestack $\mathcal{M}_1 = \mathcal{M}_{1,0}$ parameterizing smooth families of genus 1 curves. It turns out that $\mathcal{M}_1$ fails to be a stack in the étale (and fpqc) topology! The proof strategy used for Proposition 2.2.21 breaks down here since for a family of genus 1 curves $\mathcal{C} \to S$, the relative canonical bundle $\Omega_{\mathcal{C}/S}$ is trivial, so neither it nor its tensor powers can provide the relative very ample line bundle needed to embed the family into a projective space and make descent effective. In fact, Raynaud famously constructed an explicit étale cover $S' \to S$ and a family $\mathcal{C}' \to S'$ of smooth genus 1 curves (without a marked section) that does not descend to any scheme $\mathcal{C} \to S$ (it only descends to an algebraic space). For details on this ineffective descent, see [**Ray70**].

However, when we introduce a single marked point, setting $n = 1$, the prestack $\mathcal{M}_{1,1}$ (see also Example 2.1.3) behaves well. A marked point is equivalent to a section $e\colon S \to \mathcal{C}$ of the family. By Riemann–Roch, the line bundle $\mathcal{O}_{\mathcal{C}}(3e)$ is relatively very ample, and the higher direct images $\mathrm{R}^i\pi_*\mathcal{O}_{\mathcal{C}}(3e) = 0$ for $i > 0$. This line bundle induces a closed embedding of the family into a projective bundle $\mathbb{P}(\pi_*\mathcal{O}_{\mathcal{C}}(3e))$ over S. With this projective embedding secured, we can mimic the descent arguments of Proposition 2.2.21 and the Hilbert scheme construction in Theorem 2.3.9 to prove that $\mathcal{M}_{1,1}$ is an algebraic stack (see [**Ols16**, Theorem 13.1.2] for the full details). In fact, the presence of this global section allows us to bootstrap the algebraicity of $\mathcal{M}_{1,n}$ for all $n \geq 1$.

Based on the two examples above, we can state and prove the following consequence:

Theorem 2.3.12 (Algebraicity of $\mathcal{M}_{g,n}$)**.** *If $2g - 2 + n > 0$, then $\mathcal{M}_{g,n}$ is an algebraic stack over* $\operatorname{Spec}\mathbb{Z}$*.*

Proof. The proof proceeds by induction on the number of marked points n. As the base for our induction, we rely on the previously established algebraic stacks: $\mathcal{M}_g$ for $g \geq 2$ (Theorem 2.3.9), $\mathcal{M}_{1,1}$ (Examples 2.1.3, 2.3.11), and $\mathcal{M}_{0,3}$ (Example 2.3.10).

For the inductive step, assume that $\mathcal{M}_{g,n-1}$ is an algebraic stack. Consider the forgetful morphism $\pi_n\colon \mathcal{M}_{g,n} \to \mathcal{M}_{g,n-1}$ which forgets the n-th marked point. By the 2-Yoneda lemma, let $S \to \mathcal{M}_{g,n-1}$ correspond to a family of smooth curves $p\colon \mathcal{C} \to S$ equipped with $n-1$ disjoint sections $\sigma_1, \ldots, \sigma_{n-1}$. A lift of $S \to \mathcal{M}_{g,n-1}$ to $\mathcal{M}_{g,n}$ corresponds to choosing an additional n-th section $\sigma_n\colon S \to \mathcal{C}$ that is disjoint from the first $n-1$ sections. Therefore, the fiber product $S \times_{\mathcal{M}_{g,n-1}} \mathcal{M}_{g,n}$ is isomorphic to the open subscheme of the curve $\mathcal{C}$ obtained by removing the images of the first $n-1$ sections. This gives us the following Cartesian diagram:

$$\begin{tikzcd}
\mathcal{C} \setminus \bigcup_{i=1}^{n-1} \sigma_i(S) \arrow[r] \arrow[d] \arrow[dr, phantom, "\square"] & \mathcal{M}_{g,n} \arrow[d, "\pi_n"] \\
S \arrow[r] & \mathcal{M}_{g,n-1}
\end{tikzcd}$$

Since the morphism $p\colon \mathcal{C} \to S$ is a smooth and proper morphism of schemes, the open subscheme $\mathcal{C} \setminus \bigcup_{i=1}^{n-1} \sigma_i(S) \to S$ is smooth and surjective. This Cartesian square shows that the forgetful morphism π_n is representable by schemes, smooth, and surjective. Since the target $\mathcal{M}_{g,n-1}$ is an algebraic stack by our inductive hypothesis, we conclude that the source $\mathcal{M}_{g,n}$ is also an algebraic stack. $\square$

The next question is: what is the geometric origin of the condition $2g - 2 + n > 0$? We can in fact see this inequality by examining the automorphisms of the marked curve, motivated by our earlier discussion in Section 2.1.

Let C be a smooth curve. Deformation theory shows that the infinitesimal automorphisms of C (the first-order deformations of the identity) are parameterized by the vector space $\mathrm{H}^0(C, T_C)$ (see [**Ser06**, Lemma 1.2.6]). Any infinitesimal automorphism of the n-pointed curve $(C, p_1, \ldots, p_n)$ must fix each of the n marked points, which requires the corresponding global section to vanish at the effective divisor $D = p_1 + \cdots + p_n$. Consequently, the space of infinitesimal automorphisms of $(C, p_1, \ldots, p_n)$ is given by $\mathrm{H}^0\left(C, T_C\left(-\sum_{i=1}^n p_i\right)\right)$.

By Riemann–Roch, the degree of the tangent bundle is $\deg(T_C) = 2 - 2g$. Twisting by the divisor $-D$ yields the degree:

$$\deg\left(T_C\left(-\sum_{i=1}^n p_i\right)\right) = 2 - 2g - n.$$

The geometric meaning of the condition $2g - 2 + n > 0$ now becomes transparent: it is the algebraic requirement that the degree of this twisted tangent bundle is negative. Since a line bundle of negative degree on a projective curve admits no non-zero global sections, we have $\mathrm{h}^0\left(C, T_C\left(-\sum_{i=1}^n p_i\right)\right) = 0$.

This vanishing guarantees that the space of infinitesimal automorphisms is trivial, which means the global automorphism group of any smooth curve with n marked points satisfying this inequality is discrete and reduced (and in fact, finite). This deformation-theoretic insight will play an essential role in the subsequent subsections, where we establish that $\mathcal{M}_{g,n}$ is a Deligne–Mumford stack using the fundamental characterization: an algebraic stack is Deligne–Mumford if and only if its diagonal is unramified, which is equivalent to the automorphism group of any geometric point being discrete and reduced.

Up to this point, we have established the algebraicity of quotient stacks and the moduli stack of curves $\mathcal{M}_g$. A subsequent question is whether the stack of quasi-coherent sheaves $\underline{\mathrm{Qcoh}}$ (Example 2.2.5), forms an algebraic stack over $\mathsf{Sch}_{\text{ét}}$. Unfortunately, it does not. Proving that a stack is *not* algebraic is difficult if one merely attempts to rule out the existence of all possible smooth presentations. To prove this failure, we need to invoke the celebrated *Artin's Axioms for Algebraicity*, which provide necessary and sufficient criteria for a stack to be algebraic over a field k. Because $\underline{\mathrm{Qcoh}}$ is not limit-preserving (it does not commute with filtered colimits of rings) and its diagonal is non-representable, it fails Artin's axioms.

Conversely, if we bound our geometric data to a proper scheme X over a field k, the situation improves. The stack $\underline{\mathrm{Coh}}(X)$ over Sch/k, whose objects over a test scheme S are finitely presented quasi-coherent $\mathcal{O}_{X_S}$-modules that are flat over S, is indeed an algebraic stack locally of finite type over k. This foundational result is achieved by verifying that $\underline{\mathrm{Coh}}(X)$ satisfies all of Artin's axioms (see [**Alp26**, Theorems C.7.4 and C.7.7] and [**Sta26**, Tag 07SZ]).

2.3.2. Extending Scheme-Theoretic Properties to Algebraic Stacks. After introducing the basic definitions and studying the fundamental examples, we next ask how to bring the extensive toolkit developed in classical algebraic geometry into the newly developed stack world. The immediate task is to generalize the properties of schemes and their morphisms (such as those found in [**Har77**, Chapters 2, 3]) to morphisms $\mathcal{X} \to \mathcal{Y}$ of algebraic stacks. Once the relative case of morphisms is clarified, the absolute property of a single stack $\mathcal{X}$ follows immediately by evaluating the structural morphism $\mathcal{X} \to \operatorname{Spec} \mathbb{Z}$.

Case I: We start with morphisms $f \colon \mathcal{X} \to \mathcal{Y}$ that are representable by schemes. In this case, for any morphism $T \to \mathcal{Y}$ from a scheme, the base change $\mathcal{X} \times_{\mathcal{Y}} T \to T$ is a morphism between schemes. This reduction allows us to generalize any property $\mathcal{P}$ of scheme morphisms that is *stable under base change* and *local on the target* in the required topology (smooth or étale).

- **Properties generalized via base change:** Being an isomorphism, surjective, an open immersion, a closed immersion, an affine morphism, a quasi-affine morphism, proper, separated, finite, integral, and unramified... See [**Sta26**, Tag 02YJ] for more properties that are local in the fpqc topology on the target.

Case II: In the case where the morphism $f\colon \mathcal{X} \to \mathcal{Y}$ is not representable, we use the charts of the algebraic stacks to establish a bridge to schemes. Let $V \to \mathcal{Y}$ be a smooth chart of $\mathcal{Y}$. The fiber product $\mathcal{X} \times_{\mathcal{Y}} V$ is also an algebraic stack, so we can take a smooth chart $U \to \mathcal{X} \times_{\mathcal{Y}} V$. The composition $U \to \mathcal{X} \times_{\mathcal{Y}} V \to V$ is then a morphism between schemes. We say f has property $\mathcal{P}$ if this induced morphism $U \to V$ has property $\mathcal{P}$. This paradigm allows us to generalize properties that are *smooth local on both the source and target* and are *stable under composition and base change*. (If both stacks are Deligne–Mumford, we require the properties to be *étale local*).

- **Properties generalized via charts:** Smoothness, flatness, surjectivity, locally of finite presentation, locally of finite type (étaleness and unramifiedness in the case that both targets and sources are Deligne–Mumford stacks). We also refer to [**Sta26**, Tag 04YE] for more properties that are local in the smooth topology.

However, the method of "reducing to charts" possesses a limitation: it fails for global geometric properties that are not local on the source or the target. Chief among these are properties like separatedness and properness. Furthermore, purely global topological properties, such as being Noetherian or irreducible, similarly resist this local generalization. Because these properties cannot be defined simply by base changing to test schemes or passing to local charts, we are forced to develop new geometric tools to capture the global structure of a stack. We will approach this in two stages, beginning with the purely topological properties. By constructing the underlying topological space of an algebraic stack, we can generalize and study global topological properties.

Definition 2.3.13 (Topological space of an algebraic stack)**.** Let $\mathcal{X}$ be an algebraic stack. Define the topological space of $\mathcal{X}$, denoted by $|\mathcal{X}|$, as the set consisting of field-valued morphisms $x\colon \operatorname{Spec} K \to \mathcal{X}$, where two morphisms $x_1\colon \operatorname{Spec} K_1 \to \mathcal{X}$ and $x_2\colon \operatorname{Spec} K_2 \to \mathcal{X}$ are identified if there exists field extensions $K_1 \hookrightarrow K_3$ and $K_2 \hookrightarrow K_3$ such that $x_1|_{\operatorname{Spec} K_3}$ and $x_2|_{\operatorname{Spec} K_3}$ are isomorphic in $\mathcal{X}(K_3)$.

We say a subset $U \subseteq |\mathcal{X}|$ is open if there exists an open substack $\mathcal{U} \subseteq \mathcal{X}$ such that $U = |\mathcal{U}|$. A substack $\mathcal{U} \subseteq \mathcal{X}$ is called open if the induced morphism $\mathcal{U} \to \mathcal{X}$ is representable by schemes and is an open immersion.

Example 2.3.14 (The underlying topological space of classifying stacks)**.** Let k be a field and G be a smooth group scheme over k. Consider the classifying stack $\mathsf{B}G \coloneqq [\operatorname{Spec} k/G]$. The topological space $|\mathsf{B}G|$ is a single point. For the classifying stack $\mathcal{X} = \mathsf{B}G$, a morphism $x\colon \operatorname{Spec} K \to \mathsf{B}G$ from a field extension K/k corresponds to the data of a principal G-bundle $P \to \operatorname{Spec} K$. Thus, the points of $|\mathsf{B}G|$ are equivalence classes of principal G-bundles over various field extensions of k.

However, for any field extension K/k and any principal G-bundle $P \in \mathsf{B}G(\operatorname{Spec} K)$, there exists a further field extension K'/K such that the pullback of P to $\operatorname{Spec} K'$ lifts to an object in the prestack $[\operatorname{Spec} k/G]^{\mathrm{pre}}(\operatorname{Spec} K')$. Since it possesses exactly one object (the unique morphism $\operatorname{Spec} K' \to \operatorname{Spec} k$), this lift is necessarily the trivial principal G-bundle. Since every point in $|\mathsf{B}G|$ becomes isomorphic to the trivial bundle over some field extension, there is only one equivalence class.

Now we are able to generalize the topological properties of schemes to algebraic stacks.

Definition 2.3.15.

(1) An algebraic stack $\mathcal{X}$ is quasi-compact (resp. connected, irreducible) if the underlying topological space $|\mathcal{X}|$ is quasi-compact (resp. connected, irreducible).
(2) A morphism $\mathcal{X} \to \mathcal{Y}$ of algebraic stacks is quasi-compact if for every morphism $\operatorname{Spec} A \to \mathcal{Y}$, the fiber product $\mathcal{X} \times_{\mathcal{Y}} \operatorname{Spec} A$ is quasi-compact; it is of finite type if $\mathcal{X} \to \mathcal{Y}$ is locally of finite type and quasi-compact.

The topological approach discussed previously allows us to generalize purely global topological properties. However, to generalize geometric properties like separatedness and properness, it is natural to study the diagonal $\Delta \colon \mathcal{X} \to \mathcal{X} \times \mathcal{X}$ of the given stack, aligning with the classical scheme-theoretic definitions. It turns out that the diagonal of a stack is far more than just a tool for generalizing separatedness or properness; it encodes the "stackiness" and the infinitesimal automorphisms of the objects, which are the central structures we wish to study. Furthermore, analyzing the diagonal leads directly to the characterizations of Deligne–Mumford stacks and algebraic spaces, revealing the geometric essence of these different types of stacks.

To begin with, we first establish the foundational theorem regarding the representability of the diagonal.

Theorem 2.3.16 (Representability of the Diagonal)**.**

(1) The diagonal of an algebraic space is representable by schemes.
(2) The diagonal of an algebraic stack is representable.

Proof. Suppose that X is an algebraic space with an étale presentation $U \to X$. Let $R \coloneqq U \times_X U$, which is a scheme. Suppose that $T \to X \times X$ is a morphism from a scheme, then we need to show that the sheaf $Q_T \coloneqq X \times_{X \times X} T$ is a scheme. We consider the following cartesian cube.

$$\begin{tikzcd}
 & Q_{T'} \arrow[rr] \arrow[dl] \arrow[dd] & & T' \arrow[dl] \arrow[dd] \\
R \arrow[rr] \arrow[dd] & & U \times U \arrow[dd] & \\
 & Q_T \arrow[rr] \arrow[dl] & & T \arrow[dl] \\
X \arrow[rr, "\Delta"] & & X \times X &
\end{tikzcd}$$

Note that $U \to X$ is étale, surjective and representable by schemes by assumption, so is $U \times U \to X \times X$. The base change of $T \to X \times X$ by $U \times U \to X \times X$ is therefore a scheme T' which is surjective and étale over T, and $Q_{T'}$ is therefore a scheme. Also note that the morphism of schemes $R \to U \times U$ is locally quasi-finite and separated, so is $Q_{T'} \to T'$.

Now we need to apply a descent criterion for an fppf sheaf to be a scheme. Let $\mathcal{P}$ be one of the following properties: open/closed immersion, (quasi)-affine, or locally quasi-affine and separated. Let $X \to Y$ be a surjective smooth (resp. fppf) morphism of schemes. Let F be a sheaf on $(\mathsf{Sch}/Y)_{\text{ét}}$ (resp. $(\mathsf{Sch}/Y)_{\text{fppf}}$). If the fiber product $F_X \to X$ satisfies that F_X is a scheme and $F_X \to X$ has $\mathcal{P}$, then F is a scheme and $F \to Y$ has $\mathcal{P}$. For the proof we refer to [**Sta26**, Tag 02W5].

Apply this descent criterion and we conclude that Q_T is a scheme, which concludes the proof of the first part. For the second part, the proof is similar. Let $\mathcal{X}$ be an algebraic stack and $U \to \mathcal{X}$ be a smooth presentation, then $R \coloneqq U \times_{\mathcal{X}} U$ is an algebraic space. If $T \to \mathcal{X} \times \mathcal{X}$ is a morphism from a scheme, then similarly its base change along $U \times U \to \mathcal{X} \times \mathcal{X}$ yields an algebraic space

T_1 which is surjective and smooth over T. We furthermore choose an étale presentation $T_2 \to T_1$, then $T_2 \to T$ is a surjective and smooth morphism of schemes. We choose an étale cover $T' \to T$ such that $T_2 \to T$ admits a section ([**Sta26**, Tag 039P]). The composition $T' \to T_2 \to T_1 \to U \times U$ provides a morphism $T' \to U \times U$ of schemes. Therefore we obtain a similar commutative cube above except that the left and right squares are not necessarily cartesian.

Note that Q_T is always a sheaf since Q_T is identified with $\underline{\mathrm{Isom}}_{\mathcal{X}(T)}(a,a)$ with $a\colon T \to \mathcal{X}$, and this is always a sheaf as long as $\mathcal{X}$ is a stack. The morphism $Q_{T'} \to Q_T$ is an étale, surjective morphism representable by schemes since $T' \to T$ is. Choose an étale presentation $V \to Q_{T'}$ of the algebraic space $Q_{T'}$, and the composition yields an étale presentation of Q_T, showing that Q_T is an algebraic space. $\square$

Corollary 2.3.17 (Properties of representable morphisms)**.**

(1) If the diagonal of a stack $\mathcal{X}$ is representable (resp. representable by schemes), then every morphism $U \to \mathcal{X}$ from a scheme is representable (resp. representable by schemes).
(2) Every morphism from a scheme to an algebraic stack (resp. algebraic space) is representable (resp. representable by schemes).
(3) If $\mathcal{X} \to Y$ is a representable morphism of stacks over $\mathsf{Sch}_{\acute{e}t}$ and Y is an algebraic space, then $\mathcal{X}$ is an algebraic space. Consequently, the composition of two representable morphisms of stacks is representable.

Proof. For the first part, let T_1 and T_2 be arbitrary schemes mapping to $\mathcal{X}$ via morphisms a and b. The fiber product $T_1 \times_{\mathcal{X}} T_2$ is identified by the following Cartesian diagram

$$\begin{tikzcd} T_1 \times_{\mathcal{X}} T_2 \arrow[r] \arrow[d] \arrow[dr, phantom, "\square"] & T_1 \times T_2 \arrow[d, "a\times b"] \\ \mathcal{X} \arrow[r, "\Delta"'] & \mathcal{X} \times \mathcal{X} \end{tikzcd}$$

Since the base change of a representable morphism is representable, if Δ is representable (resp. representable by schemes), the fiber product $T_1 \times_{\mathcal{X}} T_2$ is an algebraic space (resp. a scheme). This shows that any morphism $T_1 \to \mathcal{X}$ is representable. The second part is an immediate consequence of the first part combined with our previous theorem establishing that the diagonal of an algebraic stack (resp. algebraic space) is representable (resp. representable by schemes).

For the third part, choose an étale presentation $V \to Y$ by a scheme. Because $\mathcal{X} \to Y$ is a representable morphism, the fiber product $\mathcal{X}_V := \mathcal{X} \times_Y V$ is an algebraic space. The diagonal of any algebraic space is a monomorphism (though not necessarily an immersion). By the étale descent of properties of morphisms, the diagonal $\mathcal{X} \to \mathcal{X} \times_Y \mathcal{X}$ is also a monomorphism. This monomorphism property guarantees that $\mathcal{X}$ is a sheaf. This can be seen via the following Cartesian diagram for any test scheme S and objects $a, b \in \mathcal{X}(S)$:

$$\begin{tikzcd} \underline{\mathrm{Isom}}_{\mathcal{X}(S)}(a,b) \arrow[r] \arrow[d] \arrow[dr, phantom, "\square"] & S \arrow[d, "{(a,b)}"] \\ \mathcal{X} \arrow[r, "\Delta"'] & \mathcal{X} \times_Y \mathcal{X} \end{tikzcd}$$

Since Δ is a monomorphism, its pullback is a monomorphism. Therefore, the isomorphism sheaf is either empty or a singleton, ensuring that $\mathcal{X}$ is a sheaf. To conclude that this sheaf is an algebraic space, it suffices to construct an étale presentation. Since $\mathcal{X}_V$ is an algebraic space, we can choose an étale presentation $U \to \mathcal{X}_V$ by a scheme. The composition $U \to \mathcal{X}_V \to \mathcal{X}$ is smooth (resp. étale) and surjective, providing the desired atlas.

Finally, the fact that compositions of representable morphisms are representable follows directly by applying this logic: if $\mathcal{Z} \to \mathcal{X}$ and $\mathcal{X} \to \mathcal{Y}$ are representable morphisms of stacks, base changing to any scheme $T \to \mathcal{Y}$ yields the sequence $\mathcal{Z}_T \to \mathcal{X}_T \to T$. Here, $\mathcal{X}_T$ is an algebraic space and $\mathcal{Z}_T \to \mathcal{X}_T$ is representable, forcing the fiber product $\mathcal{Z}_T$ to be an algebraic space. $\square$

As an application of the power of the diagonal, we first use this machinery to prove that taking the quotient of an étale equivalence relation of schemes $R \rightrightarrows U$ always yields an algebraic space U/R with an étale presentation $U \to U/R$. This completes the generalization regarding the algebraicity of quotient stacks discussed in Theorem 2.3.2.

Proposition 2.3.18 (Quotients of étale equivalence relations)**.** *Let $R \rightrightarrows U$ be an étale equivalence relation of schemes. Then the quotient sheaf U/R is an algebraic space, and the canonical projection $U \to U/R$ is an étale presentation.*

Proof. We already know that U/R is a sheaf and that the projection $U \to U/R$ is surjective, étale, and representable. Therefore, to conclude that U/R is an algebraic space, it suffices to prove that the morphism $U \to U/R$ is representable by schemes. By Proposition 2.3.17, this reduces to proving that the diagonal of U/R is representable by schemes.

Let $T \to (U/R) \times (U/R)$ be an arbitrary morphism from a test scheme T, and let Q_T be the fiber product $(U/R) \times_{(U/R)\times(U/R)} T$. We consider the following Cartesian cube constructed by base changing along the cover $U \times U \to (U/R) \times (U/R)$:

$$\begin{tikzcd}
& Q_{T'} \arrow[rr] \arrow[dl] \arrow[dd] & & T' \arrow[dl] \arrow[dd] \\
R \arrow[rr] \arrow[dd] & & U \times U \arrow[dd] & \\
& Q_T \arrow[rr] \arrow[dl] & & T \arrow[dl] \\
U/R \arrow[rr, "\Delta"'] & & (U/R) \times (U/R) &
\end{tikzcd}$$

Since $U \times U \to (U/R) \times (U/R)$ is étale and surjective, the base-changed morphism $T' \to T$ is an étale surjective cover of schemes. From the top face of the cube, we have $Q_{T'} \cong R \times_{U \times U} T'$. Since R and T' are schemes, $Q_{T'}$ is a scheme.

As established previously, since R is an equivalence relation, the structural morphism $R \to U \times U$ is a monomorphism and is locally quasi-finite and separated. Since these properties are stable under base change, the morphism $Q_{T'} \to T'$ is also locally quasi-finite and separated. By the descent criterion for an fppf sheaf to be a scheme (see the proof of Theorem 2.3.16), the fact that $Q_{T'}$ is a scheme implies that Q_T is also a scheme. Therefore, the diagonal of U/R is representable by schemes, completing the proof. $\square$

This is also an appropriate moment to introduce the stabilizer groups and the inertia stack of an algebraic stack $\mathcal{X}$. Since the diagonal of a stack governs its automorphisms, these objects are essential for encoding the internal geometry of $\mathcal{X}$.

Definition 2.3.19 (Stabilizers and inertia stacks)**.** Let $\mathcal{X}$ be an algebraic stack.

(1) Let $x\colon \operatorname{Spec} k \to \mathcal{X}$ be a field-valued point. The *stabilizer group* of x is defined as the group algebraic space $G_x := \underline{\operatorname{Aut}}_k(x)$ over $\operatorname{Spec} k$. It is identified with the following Cartesian fiber product:

$$\begin{tikzcd} G_x \arrow[r] \arrow[d] \arrow[dr, phantom, "\square"] & \operatorname{Spec} k \arrow[d, "{(x,x)}"] \\ \mathcal{X} \arrow[r, "\Delta"'] & \mathcal{X} \times \mathcal{X} \end{tikzcd}$$

(2) As the point $x \in \mathcal{X}$ varies, the stabilizer groups G_x assemble into a global family, which leads to the definition of the *inertia stack*. The inertia stack of $\mathcal{X}$, denoted $I_{\mathcal{X}}$, is defined as the fiber product of the diagonal with itself:

$$\begin{tikzcd} I_{\mathcal{X}} \arrow[r] \arrow[d] \arrow[dr, phantom, "\square"] & \mathcal{X} \arrow[d, "\Delta"] \\ \mathcal{X} \arrow[r, "\Delta"'] & \mathcal{X} \times \mathcal{X} \end{tikzcd}$$

Having taken a brief detour to witness the power of the diagonal, we now return to our primary objective: generalizing separation and properness properties for algebraic stacks. We define these geometric properties in terms of the behavior of the diagonal.

Definition 2.3.20 (Separatedness and properness)**.**

(1) A morphism of algebraic stacks $\mathcal{X} \to \mathcal{Y}$ is said to have an affine (resp. quasi-affine, separated) diagonal if the relative diagonal morphism $\Delta\colon \mathcal{X} \to \mathcal{X} \times_{\mathcal{Y}} \mathcal{X}$ is affine (resp. quasi-affine, separated).

(2) A morphism of algebraic stacks $\mathcal{X} \to \mathcal{Y}$ is quasi-separated if both its diagonal $\mathcal{X} \to \mathcal{X} \times_{\mathcal{Y}} \mathcal{X}$ and its second diagonal $\mathcal{X} \to \mathcal{X} \times_{\mathcal{X} \times_{\mathcal{Y}} \mathcal{X}} \mathcal{X}$ are quasi-compact.

(3) A morphism of algebraic stacks $\mathcal{X} \to \mathcal{Y}$ is of finite presentation if it is locally of finite presentation, quasi-compact, and quasi-separated.

(4) A morphism of algebraic stacks $\mathcal{X} \to \mathcal{Y}$ is universally closed if, for every morphism of algebraic stacks $\mathcal{Y}' \to \mathcal{Y}$, the base-changed morphism $\mathcal{X} \times_{\mathcal{Y}} \mathcal{Y}' \to \mathcal{Y}'$ induces a closed map on the underlying topological spaces.

(5) A representable morphism of algebraic stacks $\mathcal{X} \to \mathcal{Y}$ is separated if its diagonal, which is representable by schemes (see Corollary 2.3.17), is proper. It is proper if it is universally closed, separated, and of finite type.

(6) An arbitrary morphism of algebraic stacks $\mathcal{X} \to \mathcal{Y}$ is separated if its diagonal is proper. A morphism of algebraic stacks $\mathcal{X} \to \mathcal{Y}$ is proper if it is universally closed, separated, and of finite type.

(7) An algebraic stack $\mathcal{X}$ is noetherian if it is locally noetherian, quasi-separated and quasi-compact.

This foundation also enables us to define quasi-finiteness and finiteness for morphisms of algebraic stacks. Recall that a morphism of schemes is locally quasi-finite if it is locally of finite type and every topological fiber is discrete.

Definition 2.3.21 (Quasi-finiteness and finiteness)**.**

(1) A morphism of algebraic stacks $\mathcal{X} \to \mathcal{Y}$ is locally quasi-finite if it is locally of finite type, its relative diagonal $\mathcal{X} \to \mathcal{X} \times_{\mathcal{Y}} \mathcal{X}$ is locally quasi-finite, and for every morphism $\operatorname{Spec} k \to \mathcal{Y}$ from a field, the underlying topological space of the fiber $|\mathcal{X} \times_{\mathcal{Y}} \operatorname{Spec} k|$ is discrete.

(2) A morphism of algebraic stacks $\mathcal{X} \to \mathcal{Y}$ is quasi-finite if it is locally quasi-finite and quasi-compact.
(3) A representable morphism of algebraic stacks $\mathcal{X} \to \mathcal{Y}$ is finite if it is proper and quasi-finite.

Until now, we have established the general philosophies for generalizing the properties of schemes to algebraic stacks. We now transition to extending the basic geometric invariants and associated objects of schemes to the stack-theoretic setting. We will next develop the notions of dimension, tangent spaces, and residual gerbes for algebraic stacks.

2.3.3. Dimension, Tangent Spaces and Residual Gerbes. We first address the dimension of algebraic stacks. Recall that the dimension of a scheme X is the Krull dimension of its underlying topological space, while the local dimension $\dim_x X$ at a point $x \in X$ is the minimum dimension of open subsets containing x. If X is of finite type over a field and x is a closed point, this coincides with $\dim \mathcal{O}_{X,x}$.

A naive attempt to extend this concept would be to define the dimension of an algebraic stack as $\dim \mathcal{X} \coloneqq \dim |\mathcal{X}|$, the Krull dimension of its underlying topological space (Definition 2.3.13). However, this definition proves to be inadequate. As observed in Example 2.3.14, the underlying topological space of the classifying stack $\mathsf{B}G$ is a single point. Under this topological definition, we would have $\dim \mathsf{B}G = 0$ for any smooth group scheme G. This completely fails to capture the dimension of G. More fundamentally, it discards all information regarding the automorphisms of the objects parameterized by the stack, which is an unacceptable loss that violates the core philosophy of stack theory. Therefore, we need to resort to a finer definition. The following definition follows [**Sta26**, Tag 0AFN].

Definition 2.3.22 (Dimension of Algebraic Stacks)**.**

(1) Let X be a noetherian (See Definition 2.3.20) algebraic space and $x \in |X|$. The dimension of X at x is defined as
$$\dim_x X \coloneqq \dim_u U \in \mathbb{N} \cup \{\infty\},$$
where $U \to X$ is an étale presentation and $u \in U$ is a preimage of $x \in |X|$.
(2) Let $\mathcal{X}$ be a noetherian algebraic stack and $x \in |\mathcal{X}|$. Let $U \to \mathcal{X}$ be a smooth presentation and $s, t\colon R \rightrightarrows U$ be the associated smooth groupoid of algebraic spaces (See Definition 2.3.3). The dimension of $\mathcal{X}$ at x is defined as
$$\dim_x \mathcal{X} \coloneqq \dim_u U - \dim_{e(u)} R_u \in \mathbb{Z} \cup \{\infty\},$$
where R_u is the fiber of the source map $s\colon R \to U$ over u, and $e\colon U \to R$ denotes the identity morphism of the groupoid.
(3) Let $\mathcal{X}$ be a noetherian algebraic stack. The dimension of $\mathcal{X}$ is defined as
$$\dim \mathcal{X} \coloneqq \sup_{x \in |\mathcal{X}|} \dim_x \mathcal{X} \in \mathbb{Z} \cup \{\infty\}.$$

Note that $\dim_{e(u)} R_u$ is the relative dimension of the smooth presentation $U \to \mathcal{X}$. Formulating the dimension using the groupoid fiber expresses it intrinsically.

Example 2.3.23 (Dimension of quotient stacks)**.** Let G be a smooth group scheme over a field k acting on an irreducible scheme U of finite type over k. Then
$$\dim[U/G] = \dim U - \dim G.$$

In particular, we have $\dim \mathsf{B}G = -\dim G$. Therefore, the dimension of an algebraic stack can be negative!

To show this is well-defined, we verify that the local dimension $\dim_x \mathcal{X}$ is independent of the choice of the smooth presentation $U \to \mathcal{X}$ and the preimage $u \in U$ of x. We sketch the argument here and refer the reader to [**Sta26**, Tag 0AFM] for the full details.

Let $U \to \mathcal{X}$ be a smooth presentation and let $u \in U$ be a preimage of x with residue field $\kappa(u)$. (Since U is a scheme, the notion of a residue field is well-defined.) The fiber R_u is identified via the following commutative diagram, where the squares are Cartesian:

$$\begin{tikzcd} R_u \arrow[r] \arrow[d] & R \arrow[r] \arrow[d] & U \arrow[d] \\ \operatorname{Spec} \kappa(u) \arrow[r] & U \arrow[r] & \mathcal{X} \end{tikzcd}$$

This guarantees that R_u is a smooth algebraic space over $\kappa(u)$. Now, let $U' \to \mathcal{X}$ be another smooth presentation associated with a smooth groupoid $R' \rightrightarrows U'$, and let $u' \in U'$ be a preimage of x. By replacing U' with the fiber product $U \times_{\mathcal{X}} U'$, we may assume without loss of generality that there is a smooth morphism of schemes $U' \to U$ mapping u' to u. This yields a larger commutative diagram where every square is Cartesian

$$\begin{tikzcd} R'_{u'} \arrow[r] \arrow[d] & U''_u \arrow[r] \arrow[d] & U'' \arrow[r] \arrow[d] & U' \arrow[d] \\ R_u \times_{\operatorname{Spec} \kappa(u)} \operatorname{Spec} \kappa(u') \arrow[r] \arrow[d] & R_u \arrow[r] \arrow[d] & R \arrow[r] \arrow[d] & U \arrow[d] \\ \operatorname{Spec} \kappa(u') \arrow[r] & \operatorname{Spec} \kappa(u) \arrow[r] & U \arrow[r] & \mathcal{X} \end{tikzcd}$$

A careful diagram chase tracking the relative dimensions of these smooth, base-changed morphisms leads to the desired invariance. We omit the algebraic details here.

Next, we deal with the generalization of tangent spaces. Recall that for a scheme X over a field k, the tangent space of X at a k-point x is identified with the set of morphisms $\operatorname{Spec} k[\varepsilon] \to X$, where $k[\varepsilon] := k[t]/(t^2)$ is the ring of dual numbers. The k-vector space structure on this set is induced by the algebraic structure of the dual numbers: scalar multiplication by $c \in k$ is induced by pulling back along the ring map $k[\varepsilon] \to k[\varepsilon]$ defined by $\varepsilon \mapsto c\varepsilon$, while vector addition is induced by the map $k[\varepsilon] \to k[\varepsilon_1, \varepsilon_2]/(\varepsilon_1\varepsilon_2)$ defined by $\varepsilon \mapsto \varepsilon_1 + \varepsilon_2$.

In the 2-categorical setting of algebraic stacks, since we are working with categories fibered in groupoids, we need to explicitly track the 2-isomorphisms between the diagrams. A tangent vector is no longer just a morphism τ; it is a pair (τ, α). More precisely:

Definition 2.3.24 (Tangent spaces of algebraic stacks)**.** Let $\mathcal{X}$ be an algebraic stack over a field k, and let $x \colon \operatorname{Spec} k \to \mathcal{X}$ be a k-point. Set-theoretically, the tangent space of $\mathcal{X}$ at x, denoted $T_{\mathcal{X},x}$, is defined as the set of isomorphism classes of 2-commutative diagrams:

$$\begin{tikzcd} \operatorname{Spec} k \arrow[d, hook, "i"'] \arrow[dr, "x"] \arrow[d, phantom, shift left=4ex, "\alpha \nearrow"] & \\ \operatorname{Spec} k[\varepsilon] \arrow[r, "\tau"'] & \mathcal{X} \end{tikzcd}$$

where i is the canonical closed immersion induced by the quotient map $k[\varepsilon] \to k$, τ is a morphism, and $\alpha\colon \tau \circ i \Rightarrow x$ is a 2-isomorphism in $\mathcal{X}(\operatorname{Spec} k)$.

Thus far, we have not specified how the set $T_{\mathcal{X},x}$ inherits its k-vector space structure. We sketch the algebraic mechanics here and direct the reader to [**Alp26**, Proposition 4.5.10] for full details.

For a scalar $c \in k$ and a tangent vector $(\tau, \alpha) \in T_{\mathcal{X},x}$, scalar multiplication is defined as the composition $\operatorname{Spec} k[\varepsilon] \to \operatorname{Spec} k[\varepsilon] \xrightarrow{\tau} \mathcal{X}$, where the first morphism is induced by the k-algebra homomorphism $\varepsilon \mapsto c\varepsilon$. The vector addition, however, requires gluing two infinitesimal directions together, which introduces 2-categorical subtleties. Let (τ_1, α_1) and (τ_2, α_2) be two tangent vectors in $T_{\mathcal{X},x}$. We consider the fiber product of rings $k[\varepsilon_1] \times_k k[\varepsilon_2]$, which is isomorphic to $k[\varepsilon_1, \varepsilon_2]/(\varepsilon_1\varepsilon_2)$. Because $\mathcal{X}$ is an algebraic stack, by Artin's axioms for algebraicity [**Alp26**, Theorem C.7.4], it satisfies Schlessinger's (strong) homogeneity condition, which ensures that the natural functor

$$\mathcal{X}(k[\varepsilon_1] \times_k k[\varepsilon_2]) \xrightarrow{\cong} \mathcal{X}(k[\varepsilon_1]) \times_{\mathcal{X}(k)} \mathcal{X}(k[\varepsilon_2])$$

is an equivalence of categories. Under this equivalence, the pair $((\tau_1, \alpha_1), (\tau_2, \alpha_2))$ corresponds to a morphism $\tau_{1,2}\colon \operatorname{Spec}(k[\varepsilon_1] \times_k k[\varepsilon_2]) \to \mathcal{X}$, unique up to 2-isomorphism. We define the sum $(\tau_1, \alpha_1) + (\tau_2, \alpha_2)$ by composing $\tau_{1,2}$ with the closed immersion $\operatorname{Spec} k[\varepsilon] \to \operatorname{Spec}(k[\varepsilon_1] \times_k k[\varepsilon_2])$. This immersion is induced by the ring map $k[\varepsilon_1] \times_k k[\varepsilon_2] \to k[\varepsilon]$ defined by $(a + b\varepsilon_1, a + c\varepsilon_2) \mapsto a + (b + c)\varepsilon$. One checks that the scalar multiplication and addition defined above give $T_{\mathcal{X},x}$ the structure of a vector space.

Example 2.3.25 (Tangent space of $\mathcal{M}_g$)**.** Consider the moduli stack of curves of genus g, denoted $\mathcal{M}_g$, where $g \geq 2$ (Example 2.2.7). Let $[C] \in \mathcal{M}_g(k)$ be a k-valued point representing a smooth, connected, projective curve C over k. By Yoneda's lemma, the tangent space $T_{\mathcal{M}_g,[C]}$ classifies first-order infinitesimal deformations of C over $k[\varepsilon]$. Since C is smooth, any such deformation is locally trivial, so the isomorphism classes of deformations are parameterized by the cohomology group $\mathrm{H}^1(C, T_C)$, where T_C is the tangent bundle of C.

The degree of the tangent bundle is $\deg T_C = 2 - 2g$. Since $g \geq 2$, this degree is negative, which forces $\mathrm{H}^0(C, T_C) = 0$. Applying the Riemann–Roch theorem to T_C, we compute the dimension of the tangent space:

$$\dim_k T_{\mathcal{M}_g,[C]} = \dim_k \mathrm{H}^1(C, T_C) = -\chi(T_C) = -(\deg T_C + 1 - g) = 3g - 3.$$

We next generalize the concept of residue fields to algebraic stacks. For a scheme X, every point $x \in X$ admits a monomorphism $\operatorname{Spec} \kappa(x) \to X$ from the spectrum of its residue field. However, for algebraic stacks, non-trivial stabilizer groups prevent field-valued points from being monomorphisms. For example, the following Cartesian diagram demonstrates that the presentation $\operatorname{Spec} k \to \mathsf{B}G$ is a monomorphism if and only if the group scheme G is trivial.

$$\begin{tikzcd} G \arrow[r] \arrow[d] \arrow[dr, phantom, "\square"] & \operatorname{Spec} k \arrow[d] \\ \operatorname{Spec} k \arrow[r] & \mathsf{B}G \end{tikzcd}$$

This obstruction motivates a stack-theoretic replacement for the residue field. Note that the following definition specifies the structure of a residual gerbe, but does not inherently assume or guarantee its existence for an arbitrary algebraic stack.

Definition 2.3.26 (Residual gerbes)**.**

(1) Let $\mathcal{X}$ be an algebraic stack and let $x \in |\mathcal{X}|$ be a point in its underlying topological space. A residual gerbe of $\mathcal{X}$ at x is a reduced, locally Noetherian algebraic stack $\mathcal{G}_x$ equipped with a monomorphism $\mathcal{G}_x \hookrightarrow \mathcal{X}$, such that the topological space $|\mathcal{G}_x|$ consists of a single point mapping to x.
(2) A point $x \in |\mathcal{X}|$ is of finite type if there exists a representative $\operatorname{Spec} k \to \mathcal{X}$ of x that is locally of finite type.

The existence and uniqueness of residual gerbes for arbitrary algebraic stacks are technically demanding. We begin with the simpler case: residual gerbes of points of finite type. First, we need to establish that such finite type points exist on the stack.

Lemma 2.3.27 (Existence of finite type points)**.** *Let $\mathcal{X}$ be an algebraic stack. A point $x \in |\mathcal{X}|$ is of finite type if and only if there exists a scheme U, a closed point $u \in U$, and a smooth morphism $(U, u) \to (\mathcal{X}, x)$. Therefore, any algebraic stack has a finite type point.*

Proof. Assume $x \in |\mathcal{X}|$ is of finite type. By definition, there exists a representative $p\colon \operatorname{Spec} k \to \mathcal{X}$ mapping to x that is locally of finite type. Let $V \to \mathcal{X}$ be a smooth presentation by a scheme, and consider the fiber product $W \coloneqq V \times_{\mathcal{X}} \operatorname{Spec} k$. The projection $W \to V$ is locally of finite type, and the projection $W \to \operatorname{Spec} k$ is smooth. Pick a closed point $w \in W$. The field extension $\kappa(w)/k$ is finite by Zariski's lemma, ensuring the inclusion $\operatorname{Spec} \kappa(w) \to W$ is of finite type. The composition $\operatorname{Spec} \kappa(w) \to W \to V$ is therefore locally of finite type. Let $u \in V$ be the image of w, then u is a finite type point of the scheme V. Therefore, there exists an open subscheme $U \subset V$ containing u such that u is a closed point of U. The restriction $U \to \mathcal{X}$ remains smooth, yielding the required scheme U, closed point $u \in U$, and smooth morphism $U \to \mathcal{X}$.

Conversely, assume there exists a scheme U, a closed point $u \in U$, and a smooth morphism $U \to \mathcal{X}$ mapping u to x. Since u is a closed point of U, the inclusion $\operatorname{Spec} \kappa(u) \to U$ is a closed immersion, which is locally of finite type. The smooth morphism $U \to \mathcal{X}$ is locally of finite type, and the composition of locally of finite type morphisms is locally of finite type. Therefore, the representative $\operatorname{Spec} \kappa(u) \to U \to \mathcal{X}$ is locally of finite type.

To prove existence, let $\mathcal{X}$ be an algebraic stack. Choose a smooth presentation $V \to \mathcal{X}$ where V is an affine scheme. Since V is affine, it possesses a closed point v. By the equivalence established above, its image $x \in |\mathcal{X}|$ is a finite type point. □

Next, we study the existence and uniqueness of residual gerbes for a finite type point in a Noetherian algebraic stack. Under these finiteness conditions, the residual gerbe has an explicit structure.

Proposition 2.3.28 (Existence and uniqueness of residual gerbes in the Noetherian case)**.** *Let $\mathcal{X}$ be a Noetherian algebraic stack and let $x \in |\mathcal{X}|$ be a finite type point (Lemma 2.3.27). There exists a unique residual gerbe $\mathcal{G}_x$ at x. Furthermore, $\mathcal{G}_x$ satisfies the following properties:*

(1) The algebraic stack $\mathcal{G}_x$ is regular, and the morphism $\mathcal{G}_x \hookrightarrow \mathcal{X}$ is a locally closed immersion.
(2) If $\mathcal{X}$ is of finite type over a field k and $x \in \mathcal{X}(k)$ has a smooth stabilizer G_x, then $\mathcal{G}_x \cong \mathsf{B}G_x$.
(3) If $\mathcal{X}$ is a Noetherian algebraic space, then $\mathcal{G}_x \cong \operatorname{Spec} \kappa(x)$ for a field $\kappa(x)$, which is called the residue field of x.

Proof. For existence, we may replace $\mathcal{X}$ with the reduced induced substack structure on the closure $\overline{\{x\}}$. This allows us to assume $\mathcal{X}$ is reduced and $x \in |\mathcal{X}|$ is dense. Let $\operatorname{Spec} K \to \mathcal{X}$ be a representative of x of finite type. By generic flatness, this morphism is flat over a dense open substack $\mathcal{G}_x \subseteq \mathcal{X}$. Because x is dense, $\{x\}$ is the underlying topological space of $\mathcal{G}_x$. This open substack satisfies the conditions of a residual gerbe. The restricted morphism $\operatorname{Spec} K \to \mathcal{G}_x$ is fppf. Because $\operatorname{Spec} K$ is regular and regularity descends along fppf morphisms, $\mathcal{G}_x$ is a regular algebraic stack.

For uniqueness, let $\mathcal{G}$ and $\mathcal{G}'$ be two residual gerbes at x. The fiber product $\mathcal{G}'' := \mathcal{G} \times_{\mathcal{X}} \mathcal{G}'$ is a non-empty algebraic stack equipped with monomorphisms $\mathcal{G}'' \to \mathcal{G}$ and $\mathcal{G}'' \to \mathcal{G}'$. Pick a finite type point $\operatorname{Spec} F \to \mathcal{G}$, which is an fppf morphism by generic flatness. The base change $\mathcal{G}'' \times_{\mathcal{G}} \operatorname{Spec} F$ is a non-empty Noetherian algebraic space with a monomorphism to $\operatorname{Spec} F$. A reduced Noetherian algebraic space whose underlying topological space is a single point is isomorphic to the spectrum of a field. This implies the monomorphism $\mathcal{G}'' \times_{\mathcal{G}} \operatorname{Spec} F \to \operatorname{Spec} F$ is an isomorphism. By fppf descent, $\mathcal{G}'' \to \mathcal{G}$ is an isomorphism. By symmetry, $\mathcal{G}'' \to \mathcal{G}'$ is an isomorphism, proving uniqueness.

Assume $\mathcal{X}$ is of finite type over a field k and the stabilizer G_x is smooth. We define a monomorphism $(\mathsf{B}G_x)^{\mathrm{pre}} \to \mathcal{X}$. Over a k-scheme T, an object of $(\mathsf{B}G_x)^{\mathrm{pre}}(T)$ is the trivial principal G_x-bundle, which we map to the object $x_T \in \mathcal{X}(T)$ obtained by pulling back $x\colon \operatorname{Spec} k \to \mathcal{X}$ along $T \to \operatorname{Spec} k$. A morphism in $(\mathsf{B}G_x)^{\mathrm{pre}}(T)$ is an element $g \in G_x(T)$. By definition, $G_x(T) = \operatorname{Aut}_{\mathcal{X}(T)}(x_T)$. We map g to its corresponding automorphism of x_T. This defines the functor. Stackification yields a morphism $\mathsf{B}G_x \to \mathcal{X}$. Because G_x is a smooth group scheme, $\mathsf{B}G_x$ is an algebraic stack whose topological space is a point (Example 2.3.14). The morphism $\mathsf{B}G_x \to \mathcal{X}$ is a monomorphism because $\mathsf{B}G_x^{\mathrm{pre}} \to \mathcal{X}$ is. Therefore, $\mathsf{B}G_x$ is a residual gerbe. By uniqueness, $\mathcal{G}_x \cong \mathsf{B}G_x$.

The final statement follows because the stabilizer of any point in an algebraic space is trivial. □

Remark 2.3.29 (Non-Noetherian case)**.** The existence and uniqueness of residual gerbes for a non-Noetherian algebraic stack require additional technical conditions. If $\mathcal{X}$ is an algebraic stack and $x \in |\mathcal{X}|$ is of finite type, a residual gerbe at x exists. If furthermore the stabilizer group scheme of any representative of x is unramified, the residual gerbe exists and is unique (see [**Sta26**, Tag 06G3]). If the stabilizer is discrete and unramified, there exists an étale morphism $(U, u) \to (\mathcal{X}, x)$ from a scheme U where $u \in U$ is a closed point.

2.3.4. Characterization of Algebraic Spaces and Deligne–Mumford Stacks. As promised, we will next begin the study with the characterization of Deligne–Mumford stacks and algebraic spaces. After all the theory developed, it turns out that the "stackiness" of a space is entirely encoded in two interconnected places: its diagonal and its automorphism groups. This is exactly what we should expect since the entire philosophy behind stack theory is to incorporate automorphism data into the geometry. We summarize this dictionary in Table 1.

Type	Diagonal	Stabilizers
Algebraic Space	Monomorphism	Trivial
Deligne–Mumford Stack	Unramified	Discrete and reduced
Algebraic Stack	Representable	Arbitrary

Table 1. Characterization of Three Types of Stacks

We now state the theorems that establish these structural characterizations.

Theorem 2.3.30 (Characterization of Deligne–Mumford stacks)**.** *Let $\mathcal{X}$ be an algebraic stack. The following are equivalent:*

(1) *The stack $\mathcal{X}$ is a Deligne–Mumford stack.*
(2) *The relative diagonal $\Delta\colon \mathcal{X} \to \mathcal{X} \times \mathcal{X}$ is unramified.*
(3) *Every point of $\mathcal{X}$ has a discrete and reduced stabilizer group.*

Theorem 2.3.31 (Characterization of algebraic spaces)**.** *Let $\mathcal{X}$ be an algebraic stack whose diagonal is representable by schemes. The following are equivalent:*

(1) *The stack $\mathcal{X}$ is an algebraic space.*
(2) *The relative diagonal $\Delta\colon \mathcal{X} \to \mathcal{X} \times \mathcal{X}$ is a monomorphism.*
(3) *Every point of $\mathcal{X}$ has a trivial stabilizer group.*

Before we can prove these two main theorems, we need a powerful technical tool called the *minimal presentation*. To characterize a stack by its stabilizers, we need to be able to zoom in on a point and strip away all the "excess" dimension of our smooth presentations, cutting the covering space down until its dimension exactly matches the dimension of the stabilizer group. This slicing process yields what we call a minimal presentation. When the stabilizer is a discrete group, this minimal presentation drops to relative dimension zero, providing the étale covers required for Deligne–Mumford stacks.

Lemma 2.3.32 (Existence of minimal presentations)**.** *Let $\mathcal{X}$ be a Noetherian algebraic stack and let $x \in |\mathcal{X}|$ be a finite type point with a smooth stabilizer G_x. There exists a scheme U, a closed point $u \in U$, and a smooth morphism $(U,u) \to (\mathcal{X},x)$ of relative dimension $\dim G_x$ such that the diagram*

$$\begin{tikzcd} \operatorname{Spec}\kappa(u) \arrow[r, hook] \arrow[d] \arrow[dr, phantom, "\square"] & U \arrow[d] \\ \mathcal{G}_x \arrow[r, hook] & \mathcal{X} \end{tikzcd}$$

is Cartesian. In particular, if G_x is discrete and reduced, there exists an étale morphism $(U,u) \to (\mathcal{X},x)$ from a scheme.

Proof. Let $(V,v) \to (\mathcal{X},x)$ be a smooth morphism of relative dimension n from a scheme V such that $v \in V$ is a finite type point. By Proposition 2.3.28, the residual gerbe $\mathcal{G}_x$ exists, the inclusion $\mathcal{G}_x \hookrightarrow \mathcal{X}$ is a locally closed immersion, and $\mathcal{G}_x \cong \mathsf{B}G_x$ is a regular stack of dimension $-\dim G_x$. We form the Cartesian diagram

$$\begin{tikzcd} O(v) \arrow[r, hook] \arrow[d] \arrow[dr, phantom, "\square"] & V \arrow[d] \\ \mathcal{G}_x \arrow[r, hook] & \mathcal{X} \end{tikzcd}$$

Because $\mathcal{G}_x \hookrightarrow \mathcal{X}$ is a locally closed immersion, $O(v)$ is a locally closed subscheme of V. Because $V \to \mathcal{X}$ is smooth, $O(v) \to \mathcal{G}_x$ is smooth, ensuring $O(v)$ is a regular scheme of absolute dimension $c \coloneqq n - \dim G_x$.

Choose elements $f_1, \dots, f_c \in \mathcal{O}_{V,v}$ whose images in $\mathcal{O}_{O(v),v}$ form a regular sequence generating the maximal ideal at v. After replacing V with a sufficiently small open affine neighborhood of v, we may assume $f_i \in \Gamma(V, \mathcal{O}_V)$ for all $i = 1, \dots, c$. The closed subscheme $U \coloneqq V(f_1, \dots, f_c) \subset V$ intersects $O(v)$ transversely, satisfying $U \cap O(v) \cong \operatorname{Spec}\kappa(v)$. Setting $u = v$ in this subspace, the required Cartesian square is established.

We still need to verify that the restricted morphism $U \to \mathcal{X}$ is smooth at u. We first check flatness utilizing the fibral slicing criterion [**Sta26**, Tag 056X]. By smooth descent, it suffices to verify that $U \times_{\mathcal{X}} V \hookrightarrow V \times_{\mathcal{X}} V \xrightarrow{p_2} V$ is flat at a preimage u' of u under $p_1 \colon U \times_{\mathcal{X}} V \to U$. The base change $V \times_{\mathcal{X}} \operatorname{Spec} \kappa(u) \to O(u)$ is flat, and the pullbacks of $f_1, \ldots, f_c$ define a regular sequence in the local ring of $V \times_{\mathcal{X}} \operatorname{Spec} \kappa(u)$ at u'. Inductively applying the slicing criterion ensures $U \to \mathcal{X}$ is flat at u.

Because G_x is smooth, the morphism $\operatorname{Spec} \kappa(u) \to \mathcal{G}_x$ is smooth. For flat morphisms locally of finite presentation, smoothness can be checked on fibers. The fiber of $U \to \mathcal{X}$ over x is exactly $\operatorname{Spec} \kappa(u)$, so by the fibral criterion for smoothness, $U \to \mathcal{X}$ is smooth at u. The relative dimension of $U \to \mathcal{X}$ is $n - c = \dim G_x$. Replacing U with a suitably small open neighborhood of u completes the proof. ☐

With the technical machinery of minimal presentations established, we can now assemble the proofs of our main characterizations.

Proof of Theorem 2.3.30. We first establish (2) ⇔ (3). For any algebraic stack $\mathcal{X}$, the diagonal $\Delta \colon \mathcal{X} \to \mathcal{X} \times \mathcal{X}$ is representable and locally of finite type. The geometric fiber of Δ over a geometric point $(x, x) \in (\mathcal{X} \times \mathcal{X})(k)$ is isomorphic to the stabilizer group scheme G_x. A morphism locally of finite type is unramified if and only if its geometric fibers are discrete and reduced. Therefore, Δ is unramified if and only if every stabilizer group G_x is discrete and reduced.

To prove (1) ⇒ (2), assume $\mathcal{X}$ is a Deligne–Mumford stack. By definition, there exists an étale surjective presentation $V \to \mathcal{X}$ from a scheme. Consider the base change of Δ along the étale surjective morphism $V \times V \to \mathcal{X} \times \mathcal{X}$, which yields the canonical morphism $V \times_{\mathcal{X}} V \to V \times V$. Since $V \to \mathcal{X}$ is étale, the projection maps from $V \times_{\mathcal{X}} V$ to V are étale, so the morphism $V \times_{\mathcal{X}} V \to V \times V$ is unramified. Therefore, the base-changed diagonal $V \times_{\mathcal{X}} V \to V \times V$ is unramified, so is Δ.

To prove (3) ⇒ (1), assume every geometric point of $\mathcal{X}$ has a discrete and reduced stabilizer. By Lemma 2.3.32, for every finite type point $x \in |\mathcal{X}|$, there exists a scheme U_x and an étale morphism $U_x \to \mathcal{X}$ whose image contains x. Taking the disjoint union over all such points yields a morphism $U \coloneqq \coprod_x U_x \to \mathcal{X}$, which is an étale surjective presentation from a scheme. Therefore, $\mathcal{X}$ is a Deligne–Mumford stack. ☐

Proof of Theorem 2.3.31. Condition (2) asserts that the diagonal $\Delta \colon \mathcal{X} \to \mathcal{X} \times \mathcal{X}$ is a monomorphism, which is equivalent to $\mathcal{X}$ being a sheaf. Thus, the implication (1) ⇒ (2) follows directly from the definition of an algebraic space.

The equivalence (2) ⇔ (3) follows because a representable and locally of finite type morphism is a monomorphism if and only if every fiber is trivial or isomorphism. The geometric fiber of Δ at (x, x) is the stabilizer group G_x.

For (3) ⇒ (1), assume every stabilizer is trivial. By Lemma 2.3.32, there exists an étale, surjective morphism $U \to \mathcal{X}$ from a scheme U. Because we assumed the diagonal Δ is representable by schemes, any morphism from a scheme to $\mathcal{X}$ is representable by schemes. Therefore, $U \to \mathcal{X}$ is a representable, surjective, étale presentation. Because the stabilizers are trivial, Δ is a monomorphism, so $\mathcal{X}$ is a sheaf of sets. A sheaf of sets admitting a representable, surjective, étale presentation from a scheme is an algebraic space by definition. ☐

Remark 2.3.33 (Technical condition in Theorem 2.3.31)**.** Theorem 2.3.31 remains true even without the assumption that $\Delta_{\mathcal{X}}$ is representable by schemes. However, proving this requires Zariski's Main Theorem for Deligne–Mumford stacks, which asserts that a representable, quasi-finite, and separated morphism $f \colon \mathcal{X} \to \mathcal{Y}$ of Deligne–Mumford stacks factors as the composition of an open immersion $\mathcal{X} \to \underline{\operatorname{Spec}}_{\mathcal{Y}} f_* \mathcal{O}_{\mathcal{X}}$ and an affine morphism $\underline{\operatorname{Spec}}_{\mathcal{Y}} f_* \mathcal{O}_{\mathcal{X}} \to \mathcal{Y}$. As a corollary, any locally quasi-finite and separated morphism between algebraic spaces is automatically representable by schemes. Applying this directly to the diagonal $\Delta_{\mathcal{X}}$ removes the need for the representability condition. See [**Sta26**, Tag 05W7].

As a corollary, we can take a closer look at how a group action determines the geometry of its quotient stack.

Corollary 2.3.34. *Let G be a smooth group scheme acting on an algebraic space U. Then:*

(1) The quotient stack $[U/G]$ is a Deligne–Mumford stack if and only if every geometric point of U has a discrete and reduced stabilizer group, which occurs if and only if the action map $G \times U \to U \times U$ is unramified.

(2) The quotient stack $[U/G]$ is an algebraic space if and only if every geometric point of U has a trivial stabilizer group, which occurs if and only if the action map $G \times U \to U \times U$ is a monomorphism.

We close this subsection by a survey of the global properties of the moduli space of curves $\mathcal{M}_g$. We have already established that the moduli prestack of smooth curves of genus $g \geq 2$ is a stack (Proposition 2.2.21), that it is algebraic (Theorem 2.3.9), and we have computed the dimension of its tangent spaces (Example 2.3.25). We gather all the theoretical machinery we have developed to prove the following geometric properties of $\mathcal{M}_g$.

Theorem 2.3.35 (Properties of $\mathcal{M}_g$)**.** *For any integer $g \geq 2$, the moduli stack $\mathcal{M}_g$ is a smooth Deligne–Mumford stack of finite type over* $\operatorname{Spec} \mathbb{Z}$ *of dimension* $3g - 3$.

We verify these properties sequentially. First, we establish that $\mathcal{M}_g$ is of finite type over $\operatorname{Spec} \mathbb{Z}$. By Definition 2.3.15, a morphism of algebraic stacks is of finite type if it is locally of finite type and quasi-compact. This follows directly from the construction of the stack established in Theorem 2.3.9. We constructed the presentation $\mathcal{M}_g \cong [H' / \operatorname{PGL}_{5g-5}]$, where H' is a locally closed subscheme of a projective Hilbert scheme parameterizing tri-canonically embedded curves.

Next, we prove that $\mathcal{M}_g$ is a Deligne–Mumford stack. By Theorem 2.3.30, it suffices to show that for every smooth, connected, projective curve C over a field k, the automorphism group scheme $\operatorname{Aut}(C)$ is discrete and reduced. This condition is equivalent to showing that the dimension of its Lie algebra, $T_{\operatorname{Aut}(C),e}$, is zero. This Lie algebra classifies the infinitesimal automorphisms of the trivial first-order deformation of C over the dual numbers $k[\varepsilon]$. This vector space of infinitesimal automorphisms is isomorphic to the global sections of the tangent bundle, $\mathrm{H}^0(C, T_C)$. For any curve of genus $g \geq 2$, the degree of the tangent bundle is negative: $\deg T_C = 2 - 2g < 0$, so $\mathrm{H}^0(C, T_C) = 0$.

We are left to establish the smoothness and the dimension of $\mathcal{M}_g$. To achieve this, we utilize the fact that the classical infinitesimal lifting criterion for smoothness generalizes to algebraic stacks under suitable finiteness conditions (see [**Sta26**, Tag 0DP0]).

Theorem 2.3.36 (Infinitesimal lifting criterion for smoothness)**.** *Let $f\colon \mathcal{X} \to \mathcal{Y}$ be a locally of finite type morphism between locally Noetherian algebraic stacks with quasi-compact and separated diagonals. Consider a 2-commutative diagram*

$$\begin{tikzcd}
\operatorname{Spec} A_0 \arrow[r] \arrow[d, hook] & \mathcal{X} \arrow[d, "f"] \\
\operatorname{Spec} A \arrow[r] \arrow[ur, dashed] & \mathcal{Y}
\end{tikzcd}$$

of solid arrows, where $A \to A_0$ is a surjection of Artinian local rings with residue field k such that the kernel is isomorphic to k. Then:

(1) The morphism f is smooth if and only if there exists a lift for every such diagram.
(2) The morphism f is étale if and only if there exists a unique lift (up to a unique 2-isomorphism) for every such diagram.

A natural question arises: for a Noetherian, smooth Deligne–Mumford stack, is the dimension of the stack equal to the dimension of its tangent spaces, just as it is for smooth schemes? If this holds, our calculation in Example 2.3.25 instantly yields the dimension of $\mathcal{M}_g$ the moment we prove it is smooth. This geometric property holds in a broader context: the minimal presentation of an algebraic stack $\mathcal{X}$ is always miniversal, meaning it induces an isomorphism on the tangent space at the chosen preimage.

Proposition 2.3.37 (Minimality implies miniversality)**.** *Let $\mathcal{X}$ be a Noetherian algebraic stack and let $x \in |\mathcal{X}|$ be a finite type point with a smooth stabilizer G_x. Let $f\colon (U,u) \to (\mathcal{X},x)$ be a minimal presentation (i.e., a smooth morphism from a scheme such that $\mathcal{G}_x \times_{\mathcal{X}} U \cong \operatorname{Spec}\kappa(u)$). Then $U \to \mathcal{X}$ is miniversal at u, meaning the induced map on tangent spaces $T_{U,u} \to T_{\mathcal{X},f(u)}$ is an isomorphism of $\kappa(u)$-vector spaces.*

In particular, if $\mathcal{X}$ is smooth over a field k and $x \in \mathcal{X}(L)$ is a point with a smooth stabilizer over a finite extension L/k, then $\dim_x \mathcal{X} = \dim T_{\mathcal{X},x} - \dim G_x$.

Proof. The surjectivity of the differential $T_{U,u} \to T_{\mathcal{X},f(u)}$ is a direct translation of the infinitesimal lifting criterion (Theorem 2.3.36) applied to the smooth morphism $U \to \mathcal{X}$.

We prove injectivity. By the definition of a minimal presentation, we have the following Cartesian diagram:

$$\begin{tikzcd} \operatorname{Spec}\kappa(u) \arrow[r, hook] \arrow[d] \arrow[dr, phantom, "\square"] & U \arrow[d] \\ \mathcal{G}_x \arrow[r, hook] & \mathcal{X} \end{tikzcd}$$

Let $\tau\colon \operatorname{Spec}\kappa(u)[\varepsilon] \to U$ be an element of $T_{U,u}$ mapping to $0 \in T_{\mathcal{X},f(u)}$. The condition that τ maps to zero implies that its composition with $U \to \mathcal{X}$ is the trivial deformation in the stack, so it factors through the residual gerbe $\mathcal{G}_x$. Since the diagram is Cartesian, the morphism τ must factor through the fiber product $\operatorname{Spec}\kappa(u)$. Therefore, the infinitesimal deformation τ is trivial in U, proving $\tau = 0$. ☐

Proof of the smoothness of $\mathcal{M}_g$. We apply Theorem 2.3.36 to the structure morphism $\mathcal{M}_g \to \operatorname{Spec}\mathbb{Z}$. Let $[C] \in \mathcal{M}_g(k)$ be a moduli point representing a smooth curve C over k, and consider the lifting problem:

$$\begin{tikzcd} \operatorname{Spec} k \arrow[r] \arrow[rr, bend left, "{[C]}"] & \operatorname{Spec} A_0 \arrow[r] \arrow[d, hook] & \mathcal{M}_g \arrow[d, "f"] \\ & \operatorname{Spec} A \arrow[r] \arrow[ur, dashed] & \operatorname{Spec}\mathbb{Z} \end{tikzcd}$$

By Yoneda's lemma, the morphism $\operatorname{Spec} A_0 \to \mathcal{M}_g$ corresponds to a flat family of smooth curves $\mathcal{C}_0 \to \operatorname{Spec} A_0$. The lifting problem now translates into the classical deformation theory: does there exist a higher-order deformation $\mathcal{C}$ over A completing the following diagram?

$$\begin{tikzcd} C \arrow[r, hook] \arrow[d] & \mathcal{C}_0 \arrow[r, hook, dashed] \arrow[d] & \mathcal{C} \arrow[d, dashed] \\ \operatorname{Spec} k \arrow[r, hook] & \operatorname{Spec} A_0 \arrow[r, hook] & \operatorname{Spec} A \end{tikzcd}$$

For a smooth projective curve C, the obstructions to lifting the family $\mathcal{C}_0$ to A reside in the cohomology group $\mathrm{H}^2(C, T_C)$ (see [**Ser06**, Proposition 1.2.12]). Since C is a curve, Grothendieck's vanishing theorem shows that cohomology vanishes above the dimension of the space, forcing $\mathrm{H}^2(C, T_C) = 0$. The vanishing of the obstruction space guarantees the existence of the lift $\mathcal{C}$. Thus, the infinitesimal lifting criterion is satisfied, proving $\mathcal{M}_g$ is smooth over $\operatorname{Spec}\mathbb{Z}$.

Since we have proved that $\mathcal{M}_g$ is a Deligne–Mumford stack using Theorem 2.3.30, its stabilizers are discrete, so $\dim G_{[C]} = 0$. By Proposition 2.3.37, the relative dimension of the stack equals the dimension of its tangent space. From Example 2.3.25, we have $\dim T_{\mathcal{M}_g,[C]} = \dim \mathrm{H}^1(C, T_C) = 3g-3$. This concludes the proof of Theorem 2.3.35. □

2.3.5. Quasi-coherent sheaves, Cohomology, and Intersection Theory. In the preceding subsections, we have established the framework of algebraic stacks. In this subsection, we generalize the notion of quasi-coherent sheaves from schemes to algebraic stacks in order to study their geometry. We also discuss the definition of singular cohomology for an algebraic stack and construct the Chow ring of a smooth Deligne–Mumford stack. Since our primary concern is the moduli stack $\mathcal{M}_g$, which is a smooth Deligne–Mumford stack for $g \geq 2$, we will restrict ourselves to the Deligne–Mumford framework throughout this subsection. Until we define Chow rings, there is no significant difference between the theories for Deligne–Mumford stacks and general algebraic stacks, and we will explicitly highlight any differences when necessary.

To define quasi-coherent sheaves on a given Deligne–Mumford stack $\mathcal{X}$, we first define its small étale site.

Definition 2.3.38 (Small étale site)**.** Let $\mathcal{X}$ be a Deligne–Mumford stack. The small étale site of $\mathcal{X}$, denoted by $\mathcal{X}_{\text{ét}}$, is the category of schemes étale over $\mathcal{X}$. A covering of an $\mathcal{X}$-scheme U is a collection of étale morphisms $U_i \to U$ over $\mathcal{X}$ that is jointly surjective.

Equipped with this site, we can now discuss sheaves on $\mathcal{X}$. In this subsection, we only consider sheaves of abelian groups, viewed as functors $(\mathsf{Sch}/\mathcal{X})^{\text{op}}_{\text{ét}} \to \mathsf{Ab}$. The sections of a sheaf F over an étale $\mathcal{X}$-scheme U are denoted by $F(U \to \mathcal{X})$, or simply $F(U)$ if there is no danger of confusion. We denote by $\mathsf{Ab}(\mathcal{X}_{\text{ét}})$ the category of abelian sheaves on $\mathcal{X}_{\text{ét}}$. We first examine the following foundational examples.

Example 2.3.39 (Structure sheaf)**.** Let $\mathcal{X}$ be a Deligne–Mumford stack. The structure sheaf $\mathcal{O}_{\mathcal{X}}$ is defined by $\mathcal{O}_{\mathcal{X}}(U) \coloneqq \Gamma(U, \mathcal{O}_U)$.

Example 2.3.40 (Differentials)**.** Let $\mathcal{X}$ be a Deligne–Mumford stack over a scheme S. The relative sheaf of differentials $\Omega_{\mathcal{X}/S}$ is defined by $\Omega_{\mathcal{X}/S}(U) \coloneqq \Gamma(U, \Omega_{U/S})$.

Example 2.3.41 (Hodge bundle, first visit)**.** Consider the Deligne–Mumford stack $\mathcal{M}_g$ for $g \geq 2$, and let $U \to \mathcal{M}_g$ be an étale morphism, which is classified by a flat family of smooth curves $\pi\colon \mathcal{C} \to U$. We define the Hodge sheaf $\mathcal{H}$ by $\mathcal{H}(U \to \mathcal{M}_g) \coloneqq \Gamma(\mathcal{C}, \Omega_{\mathcal{C}/U})$. This is equivalent to setting $\mathcal{H} = \pi_*\Omega_{\mathcal{M}_{g,1}/\mathcal{M}_g}$. We will later prove that this is, in fact, a vector bundle of rank g.

We can also extend the definition of a sheaf on $\mathcal{X}_{\text{ét}}$ from étale $\mathcal{X}$-schemes to étale $\mathcal{X}$-Deligne–Mumford stacks by gluing. Let $\mathcal{U} \to \mathcal{X}$ be an étale morphism of Deligne–Mumford stacks. Choose an étale presentation $U \to \mathcal{U}$ by a scheme, and let $R \to U \times_{\mathcal{U}} U$ be an étale presentation of the fiber product by a scheme. We define the sections of F over $\mathcal{U}$ as the equalizer:

$$F(\mathcal{U} \to \mathcal{X}) := \mathrm{Eq}\left(F(U \to \mathcal{X}) \rightrightarrows F(R \to \mathcal{X})\right).$$

This definition is independent of the choice of presentations. To see this, suppose $U' \to \mathcal{U}$ and $R' \to U' \times_{\mathcal{U}} U'$ is another choice of presentations. By replacing U' with the fiber product $U \times_{\mathcal{U}} U'$ (and similarly for R'), we may assume without loss of generality that there exists a surjective étale morphism $U' \to U$ over $\mathcal{U}$. Since F is a sheaf on the étale site of schemes, the sheaf axiom applied to the étale covering $U' \to U$ guarantees that the equalizer computed via U is isomorphic to the equalizer computed via the refinement U'. Therefore, the sections $F(\mathcal{U} \to \mathcal{X})$ are well-defined. In particular, applying this to the identity morphism $\mathcal{X} \xrightarrow{\mathrm{id}} \mathcal{X}$, we define the global sections of the sheaf F as $\Gamma(\mathcal{X}, F) := F(\mathcal{X} \xrightarrow{\mathrm{id}} \mathcal{X})$.

We next turn to the generalization of $\mathcal{O}_{\mathcal{X}}$-modules, quasi-coherent sheaves, and vector bundles. The generalization of $\mathcal{O}_{\mathcal{X}}$-modules is completely natural: they are defined as module objects over the structure sheaf $\mathcal{O}_{\mathcal{X}}$ (Example 2.3.39).

Definition 2.3.42 ($\mathcal{O}_{\mathcal{X}}$-modules)**.** Let $\mathcal{X}$ be a Deligne–Mumford stack. An $\mathcal{O}_{\mathcal{X}}$-module is a sheaf F on $\mathcal{X}_{\text{ét}}$ such that for every étale $\mathcal{X}$-scheme $U \to \mathcal{X}$, $F(U)$ is a module over $\mathcal{O}_{\mathcal{X}}(U)$, and this module structure is compatible with the restriction maps along morphisms $V \to U$ of étale $\mathcal{X}$-schemes. We denote the category of $\mathcal{O}_{\mathcal{X}}$-modules by $\mathsf{Mod}(\mathcal{O}_{\mathcal{X}})$.

Exactly as in the classical scheme case [**Har77**, II. 5], given two $\mathcal{O}_{\mathcal{X}}$-modules F and G, we define their tensor product $F \otimes_{\mathcal{O}_{\mathcal{X}}} G$ as the sheafification of the presheaf given by $U \mapsto F(U) \otimes_{\mathcal{O}_{\mathcal{X}}(U)} G(U)$. The inner Hom sheaf, $\mathcal{H}om_{\mathcal{O}_{\mathcal{X}}}(F, G)$, is defined by assigning $U \mapsto \mathrm{Hom}_{\mathcal{O}_U}(F|_U, G|_U)$ for each étale $\mathcal{X}$-scheme U. For any morphism $f\colon \mathcal{X} \to \mathcal{Y}$ of Deligne–Mumford stacks, one defines the pushforward $f_*F(V \to \mathcal{Y}) := F(\mathcal{X} \times_{\mathcal{Y}} V \to \mathcal{X})$ and the pullback $f^*(-) := f^{-1}(-) \otimes_{f^{-1}\mathcal{O}_{\mathcal{Y}}} \mathcal{O}_{\mathcal{X}}$, where the inverse image $f^{-1}G$ is defined as the sheafification of the presheaf

$$(U \to \mathcal{X}) \mapsto \mathrm{colim}_{(V \to \mathcal{Y}, U \to \mathcal{X} \times_{\mathcal{Y}} V)}\, G(V \to \mathcal{Y}).$$

The pullback f^* remains left adjoint to f_*. Furthermore, $\mathsf{Mod}(\mathcal{O}_{\mathcal{X}})$ is an abelian category. The proof of these properties is identical to the classical proofs in [**Har77**].

The definition of a quasi-coherent sheaf is obtained by gluing local quasi-coherent sheaves (in the classical Zariski sense) together into a global sheaf on the site. We establish a notational convention here: for an étale $\mathcal{X}$-scheme U, we denote by $F|_{U_{\text{ét}}}$ its restriction to the small étale site of U, and by $F|_{U_{\mathrm{Zar}}}$ its restriction to the Zariski topology of U.

Definition 2.3.43 (Quasi-coherent sheaves)**.** Let $\mathcal{X}$ be a Deligne–Mumford stack. An $\mathcal{O}_{\mathcal{X}}$-module F is quasi-coherent if it satisfies two conditions:

(1) For every étale $\mathcal{X}$-scheme U, the restriction $F|_{U_{\mathrm{Zar}}}$ is a quasi-coherent $\mathcal{O}_{U_{\mathrm{Zar}}}$-module.
(2) For every étale morphism $f\colon U \to V$ of $\mathcal{X}$-schemes, the natural morphism $f^*(F|_{V_{\mathrm{Zar}}}) \to F|_{U_{\mathrm{Zar}}}$ is an isomorphism.

A quasi-coherent $\mathcal{O}_{\mathcal{X}}$-module F is a vector bundle if $F|_{V_{\mathrm{Zar}}}$ is a locally free sheaf of finite rank for every $V \to \mathcal{X}$. If $\mathcal{X}$ is locally Noetherian, then F is coherent if $F|_{V_{\mathrm{Zar}}}$ is a coherent sheaf for every $V \to \mathcal{X}$.

Clearly, the structure sheaf in Example 2.3.39 is quasi-coherent. The relative sheaf of differentials $\Omega_{\mathcal{X}/S}$ in Example 2.3.40 is also quasi-coherent. It suffices to verify the gluing condition in Definition

2.3.43. Let $f\colon U \to V$ be an étale morphism of étale $\mathcal{X}$-schemes. The exact sequence of Kähler differentials gives $f^*\Omega_{V/S} \to \Omega_{U/S} \to \Omega_{U/V} \to 0$. Since f is an étale morphism, it is both formally unramified and formally smooth, so the natural map $f^*\Omega_{V/S} \to \Omega_{U/S}$ is an isomorphism. Moreover, if the structure map $\mathcal{X} \to S$ is smooth, then $\Omega_{\mathcal{X}/S}$ is a vector bundle.

We also establish that the Hodge sheaf defined in Example 2.3.41 is a vector bundle of rank g. This follows directly from Proposition 2.2.20, which guarantees that the pushforward $\pi_*\Omega_{\mathcal{C}/U}$ is a locally free sheaf of rank g.

Example 2.3.44 (Quasi-coherent sheaves on $\mathsf{B}G$)**.** Let G be a finite abstract group viewed as a group scheme over a field k. One verifies that a quasi-coherent sheaf on $\mathsf{B}G$ corresponds to a representation V of G. To see this, note that giving a quasi-coherent sheaf on a quotient stack $[X/G]$ is equivalent to giving a G-equivariant quasi-coherent sheaf on X. Applying this to $X = \operatorname{Spec} k$, a quasi-coherent sheaf on $\operatorname{Spec} k$ is simply a k-vector space V. The G-equivariant structure on V is a k-linear action of the group G on V. Therefore, the category of quasi-coherent sheaves on $\mathsf{B}G$ is equivalent to the category of k-linear representations of G.

For the sake of completeness, we briefly discuss sheaf cohomology theory for Deligne–Mumford stacks. Fortunately, for algebraic spaces and Deligne–Mumford stacks, almost every foundational result from classical scheme theory (such as those in [**Har77**, Chapter III]) generalizes seamlessly. Mimicking the classical proofs (e.g., [**Har77**, Proposition III.2.2]), one can show that if $\mathcal{X}$ is a Deligne–Mumford stack, the abelian categories $\mathsf{Ab}(\mathcal{X}_{\text{ét}})$ and $\mathsf{Mod}(\mathcal{O}_{\mathcal{X}})$ have enough injectives. Furthermore, if $\mathcal{X}$ is quasi-separated, the category of quasi-coherent sheaves $\mathsf{Qcoh}(\mathcal{X})$ also has enough injectives. With these injectives in hand, one defines the sheaf cohomology group $\mathrm{H}^p(\mathcal{X}_{\text{ét}}, F)$ as the p-th right derived functor of the global sections functor $\Gamma(\mathcal{X}_{\text{ét}}, -)$. Similarly, one defines the Picard group $\operatorname{Pic}(\mathcal{X})$ of a Deligne–Mumford stack as the group of isomorphism classes of line bundles on $\mathcal{X}$, with the group operation given by the tensor product.

Remark 2.3.45 (Quasi-coherent sheaves on algebraic stacks)**.** If $\mathcal{X}$ is a general algebraic stack, the theory of quasi-coherent sheaves and sheaf cohomology can be generalized by replacing the small étale topology with the lisse-étale topology (Example 2.2.12). However, a technical issue arises: if $U \to V$ and $U' \to V$ are arbitrary morphisms of smooth $\mathcal{X}$-schemes, their fiber product $U \times_V U'$ is not necessarily smooth over $\mathcal{X}$. Consequently, the site lacks fiber products, so the inverse image functor f^{-1} does not commute with finite limits and is therefore not left exact. This lack of exactness makes functorial constructions in the lisse-étale site more technical.

We now briefly discuss the singular (co)homology of an algebraic stack, following the approach in [**Beh04**]. Recall that a stack $\mathcal{X}$ over the category Top (equipped with the big topological site, see Example 2.2.9) is called a topological stack if it admits a continuous, surjective morphism $p\colon X \to \mathcal{X}$ representable by topological spaces, where $X \in \mathsf{Top}$. See [**Noo12**].

Choose a presentation $U_0 \to \mathcal{X}$, which induces a topological groupoid $U_1 \coloneqq U_0 \times_{\mathcal{X}} U_0 \rightrightarrows U_0$. This topological groupoid induces a simplicial topological space:

$$U_\bullet : \qquad \cdots \; \substack{\longrightarrow\\[-2pt]\longrightarrow\\[-2pt]\longrightarrow\\[-2pt]\longrightarrow} \; U_2 \; \substack{\longrightarrow\\[-2pt]\longrightarrow\\[-2pt]\longrightarrow} \; U_1 \rightrightarrows U_0$$

where $U_p \coloneqq \underbrace{U_0 \times_{\mathcal{X}} \cdots \times_{\mathcal{X}} U_0}_{p+1 \text{ times}}$. The differential face maps $\partial_0, \ldots, \partial_p\colon U_p \to U_{p-1}$ are given by forgetting the $i+1$-th term, and the degeneracy maps $s_i\colon U_{p-1} \to U_p$ are given by inserting an identity morphism at the i-th position.

On the other hand, for each integer p, one associates a singular chain complex $C_\bullet(U_p)$ equipped with the standard topological differential $d\colon C_q(U_p) \to C_{q-1}(U_p)$. We define a simplicial differential $\partial \coloneqq \sum_{i=0}^{p}(-1)^i\partial_i\colon C_q(U_p) \to C_q(U_{p-1})$ induced by the face maps $\partial_i\colon U_p \to U_{p-1}$. Together, these operators define a bicomplex $C_\bullet(U_\bullet)$:

$$\begin{tikzcd}
\cdots \arrow[d] & \cdots \arrow[d] & \cdots \arrow[d] & \cdots \arrow[d] & \cdots \arrow[d] \\
\cdots \arrow[r] \arrow[d] & C_{q+1}(U_{p+1}) \arrow[r, "\partial"] \arrow[d, "d"] & C_{q+1}(U_p) \arrow[r, "\partial"] \arrow[d, "d"] & C_{q+1}(U_{p-1}) \arrow[r, "\partial"] \arrow[d, "d"] & \cdots \arrow[d] \\
\cdots \arrow[r] \arrow[d] & C_q(U_{p+1}) \arrow[r, "\partial"] \arrow[d, "d"] & C_q(U_p) \arrow[r, "\partial"] \arrow[d, "d"] & C_q(U_{p-1}) \arrow[r, "\partial"] \arrow[d, "d"] & \cdots \arrow[d] \\
\cdots \arrow[r] \arrow[d] & C_{q-1}(U_{p+1}) \arrow[r, "\partial"] \arrow[d] & C_{q-1}(U_p) \arrow[r, "\partial"] \arrow[d] & C_{q-1}(U_{p-1}) \arrow[r, "\partial"] \arrow[d] & \cdots \arrow[d] \\
\cdots & \cdots & \cdots & \cdots & \cdots
\end{tikzcd}$$

We define the singular chain complex of the topological stack, $C_\bullet(\mathcal{X})$, as the total complex of this bicomplex:

$$C_k(\mathcal{X}) \coloneqq \bigoplus_{p+q=k} C_q(U_p)$$

equipped with the total differential $\delta\colon C_k(\mathcal{X}) \to C_{k-1}(\mathcal{X})$ given by

$$\delta(\gamma) \coloneqq (-1)^{p+q}\partial\gamma + (-1)^q d\gamma, \quad \text{for } \gamma \in C_q(U_p).$$

A direct algebraic verification shows that $\delta^2 = 0$. Thus, we can define the singular homology groups of $\mathcal{X}$ with coefficients in an abelian group $A \in \mathsf{Ab}$ as

$$\mathrm{H}_n(\mathcal{X}, A) \coloneqq \mathrm{H}_n\left(C_\bullet(\mathcal{X}) \otimes_{\mathbb{Z}} A\right).$$

The singular cohomology is defined simply by dualizing the chain complex prior to taking cohomology

$$\mathrm{H}^n(\mathcal{X}, A) \coloneqq \mathrm{H}^n\left(C^\bullet(\mathcal{X}) \otimes_{\mathbb{Z}} A\right).$$

To ensure these groups are well-defined (i.e., independent of the choice of presentation $U_0 \to \mathcal{X}$), it suffices to prove that the definition is invariant under Morita equivalence (Theorem 2.3.8). This invariance is explicitly verified in [**Beh04**, Remark 1].

This theory admits a finer topological construction. In [**Noo12**], the classifying space of a topological stack is defined via a chart $U \to \mathcal{X}$ that is a universal weak equivalence. It is proven that for every topological stack $\mathcal{X}$, there exists a unique classifying space $\mathsf{B}\mathcal{X}$ up to weak homotopy equivalence. Moreover, this assignment satisfies a functoriality: there is a functor $\theta\colon \mathsf{TopSt} \to \mathsf{Ho}(\mathsf{Top})$ mapping $\mathcal{X}$ to $\mathsf{B}\mathcal{X}$ in the homotopy category of topological spaces. One can alternatively define the homology of the stack as $\mathrm{H}_*(\mathcal{X}) \coloneqq \mathrm{H}_*(\mathsf{B}\mathcal{X})$. The construction of the classifying space ensures that these two definitions of homology are naturally isomorphic, which inherently establishes the functoriality of $\mathrm{H}_*(-)$ for topological stacks.

We close this subsection with a brief development of the intersection theory for Deligne–Mumford stacks, following the foundational approach of Vistoli [**Vis89**]. We emphasize a standing convention for the remainder of this thesis: whenever we discuss intersection theory, we restrict our focus exclusively to Deligne–Mumford stacks. Furthermore, all stacks and schemes in this context are assumed to be of finite type over a fixed base field k, and all stacks are assumed to be Noetherian (Definition 2.3.20).

Recall that in the classical setting of Fulton [**Ful98**], the k-th Chow group of an algebraic scheme is defined as the free abelian group generated by its k-dimensional subvarieties, modulo rational equivalence. One defines proper pushforwards and flat pullbacks of cycles, and proves that these operations descend to the level of Chow groups. We define the Chow groups of a Deligne–Mumford stack in a parallel fashion. Let $\mathcal{X}$ be a Deligne–Mumford stack. A cycle of dimension k on $\mathcal{X}$ is an element of the free abelian group $Z_k(\mathcal{X})$ generated by all integral closed substacks of dimension k. The group of rational equivalences on $\mathcal{X}$ is generated by $W_k(\mathcal{X}) \coloneqq \bigoplus_{\mathcal{Y}} k(\mathcal{Y})^{\times}$, where the direct sum is taken over all integral closed substacks $\mathcal{Y}$ of $\mathcal{X}$ of dimension $k+1$, and $k(\mathcal{Y})^{\times}$ denotes the multiplicative group of the rational function field of $\mathcal{Y}$. We set the total graded groups as $Z_*(\mathcal{X}) \coloneqq \bigoplus_k Z_k(\mathcal{X})$ and $W_*(\mathcal{X}) \coloneqq \bigoplus_k W_k(\mathcal{X})$.

For an algebraic scheme X, there is a homomorphism $\partial_X \colon W_*(X) \to Z_*(X)$ mapping a non-zero rational function f on a subvariety to its associated Weil divisor $\operatorname{div}(f)$. For a proper morphism of schemes $f\colon X \to Y$, one defines a functorial pushforward homomorphism $f_*\colon Z_*(X) \to Z_*(Y)$. Crucially, this pushforward preserves rational equivalence, inducing a parallel map $f_*\colon W_*(X) \to W_*(Y)$ [**Ful98**, Proposition 1.4]. Conversely, for a flat morphism $f\colon X \to Y$, pulling back rational functions to the components of the flat pullback cycle induces a well-defined map $f^*\colon W_*(Y) \to W_*(X)$, which parallels the flat pullback of cycles $f^*\colon Z_*(Y) \to Z_*(X)$ [**Ful98**, Lemmas 1.7.1, 1.7.2]. These proper pushforward and flat pullback operations are functorial and respect the divisor map, yielding the following commutative diagrams:

$$\begin{tikzcd} W_*(X) \arrow[r, "f_*"] \arrow[d, "\partial_X"] & W_*(Y) \arrow[d, "\partial_Y"] \\ Z_*(X) \arrow[r, "f_*"] & Z_*(Y) \end{tikzcd} \qquad \begin{tikzcd} W_*(Y) \arrow[r, "f^*"] \arrow[d, "\partial_Y"] & W_*(X) \arrow[d, "\partial_X"] \\ Z_*(Y) \arrow[r, "f^*"] & Z_*(X) \end{tikzcd}$$

Since étale morphisms are flat, the flat pullback operations allow the assignments $U \mapsto Z_*(U)$ and $U \mapsto W_*(U)$ to satisfy the sheaf axiom on the small étale site of $\mathcal{X}$. Thus, they define sheaves $\mathcal{Z}_*$ and $\mathcal{W}_*$. The local divisor maps ∂_U glue into a sheaf morphism $\partial\colon \mathcal{W}_* \to \mathcal{Z}_*$. Evaluating these sheaves on global sections yields the global divisor homomorphism $\partial_{\mathcal{X}}\colon W_*(\mathcal{X}) \to Z_*(\mathcal{X})$. The Chow group of the stack $\mathcal{X}$, denoted by $\mathrm{CH}_*(\mathcal{X})$, is defined as the cokernel of $\partial_{\mathcal{X}}$. Finally, the rational Chow group is obtained by extending scalars to the rationals: $\mathrm{CH}_*(\mathcal{X})_{\mathbb{Q}} \coloneqq \mathrm{CH}_*(\mathcal{X}) \otimes_{\mathbb{Z}} \mathbb{Q}$.

We now explain why intersection theory in this thesis is restricted to the category of Deligne–Mumford stacks. In the setting of general algebraic stacks, naive cycle theory collapses because there are not enough closed substacks to support non-torsion line bundles. Consider the classifying stack $\mathsf{B}G$. For any smooth group scheme G, the underlying topological space is a single point (Example 2.3.14), so its only closed substacks are the empty stack and $\mathsf{B}G$ itself. Consequently, there are no closed substacks of codimension one, and the cycle group $Z_{-\dim G-1}(\mathsf{B}G)$ is zero. If G is a finite discrete group (making $\mathsf{B}G$ a Deligne–Mumford stack), every character of G has finite order. The corresponding line bundles are torsion, so their first Chern classes vanish when passing to the rational Chow group with $\mathbb{Q}$-coefficients. The cycle group being zero creates no contradiction. However, if G is an infinite smooth group scheme like $\mathbb{G}_m$ (making $\mathsf{B}\mathbb{G}_m$ an algebraic stack), it admits characters of infinite order. The stack $\mathsf{B}\mathbb{G}_m$ thus admits non-torsion line bundles (Example 2.3.44), which must have non-zero first Chern classes. Since the codimension-one cycle group is zero, these non-torsion line bundles cannot be represented by divisors. This geometric obstruction prevents a coherent definition of the first Chern class [**Ful98**, §2.5], causing naive intersection theory to fail on general algebraic stacks.

Our next goal is to establish whether the Chow group $\mathrm{CH}_*(\mathcal{X})$ of a Deligne–Mumford stack admits proper pushforwards and flat pullbacks. To achieve this, we first generalize the notion of the degree of a proper morphism $f\colon \mathcal{X} \to \mathcal{Y}$ between Deligne–Mumford stacks. Since stacks incorporate internal automorphisms, this generalized degree can be a rational number! Consequently, proper pushforwards for stacks will only be well-defined for rational Chow groups.

To define the degree of an arbitrary morphism of Deligne–Mumford stacks, we first establish the definition for representable morphisms. Let $f\colon \mathcal{X} \to \mathcal{Y}$ be a separated, dominant morphism of finite type between integral Deligne–Mumford stacks. If f is representable, we define

$$\deg(\mathcal{X}/\mathcal{Y}) \coloneqq \deg(\mathcal{X} \times_{\mathcal{Y}} V/V),$$

where $V \to \mathcal{Y}$ is an étale presentation by an integral scheme V. Since f is representable, the fiber product $\mathcal{X} \times_{\mathcal{Y}} V$ is an algebraic space. This degree is well-defined because every quasi-separated algebraic space contains a dense open subscheme [**Sta26**, Tag 06NN], allowing its rational function field and degree over V to be computed exactly as in the classical scheme-theoretic case. One easily verifies that this definition is independent of the choice of presentation: if V' is another presentation, we can assume without loss of generality that there is a surjective étale cover $V' \to V$ (by replacing V' with the fiber product $V \times_{\mathcal{Y}} V'$), which preserves the generic degree.

For an arbitrary separated, dominant morphism $f\colon \mathcal{X} \to \mathcal{Y}$, we choose an étale presentation $U \to \mathcal{X}$ from an integral scheme. Since the presentation morphism $U \to \mathcal{X}$ is representable, the composition $U \to \mathcal{Y}$ is also representable. We then define

$$\deg(\mathcal{X}/\mathcal{Y}) \coloneqq \frac{\deg(U/\mathcal{Y})}{\deg(U/\mathcal{X})}.$$

The well-definedness of this rational degree (its independence from the choice of U) is verified similarly. If one chooses another presentation $U' \to \mathcal{X}$, the fiber product $U'' \coloneqq U \times_{\mathcal{X}} U'$ is an algebraic space that dominates both U and U'. We then compute

$$\frac{\deg(U/\mathcal{Y})}{\deg(U/\mathcal{X})} = \frac{\deg(U''/U)\deg(U/\mathcal{Y})}{\deg(U''/U)\deg(U/\mathcal{X})} = \frac{\deg(U''/\mathcal{Y})}{\deg(U''/\mathcal{X})} = \frac{\deg(U'/\mathcal{Y})}{\deg(U'/\mathcal{X})}.$$

Note that this verification relies on the multiplicativity of the degree for representable morphisms of stacks, which is proven in [**Vis89**, Lemma 1.16].

With the generalized notion of degree established, we can formally define flat pullbacks and proper pushforwards. Let $f\colon \mathcal{X} \to \mathcal{Y}$ be a morphism of Deligne–Mumford stacks. If f is flat, we define the flat pullback $f^*\colon Z_*(\mathcal{Y}) \to Z_*(\mathcal{X})$ on the generators by $f^*[\mathcal{Y}'] \coloneqq [\mathcal{Y}' \times_{\mathcal{Y}} \mathcal{X}]$ for any integral closed substack $\mathcal{Y}' \subseteq \mathcal{Y}$. If f is proper, we define the proper pushforward on rational cycles, $f_*\colon Z_*(\mathcal{X})_{\mathbb{Q}} \to Z_*(\mathcal{Y})_{\mathbb{Q}}$, by extending the assignment linearly from the generators: $f_*([\mathcal{X}']) \coloneqq \deg(\mathcal{X}'/\mathcal{Y}')[\mathcal{Y}']$, where $\mathcal{Y}'$ is the image of the integral closed substack $\mathcal{X}'$. Just as in the classical proofs for schemes, one shows that both the flat pullback and the proper pushforward respect rational equivalence, thereby descending to well-defined operations on the rational Chow groups $\mathrm{CH}_*(\mathcal{X})_{\mathbb{Q}}$. See [**Vis89**, Proposition 3.7].

The most technical part of the intersection theory of Deligne–Mumford stacks is the construction of the generalized Gysin homomorphism. We recall the classical construction for regular embeddings of schemes. Let $i\colon X \to Y$ be a regular embedding of schemes of codimension d, and let $f\colon V \to Y$ be an arbitrary morphism of schemes with V being purely k-dimensional. Consider the fiber square:

$$\begin{tikzcd} W \arrow[r, "j"] \arrow[d, "g"] \arrow[dr, phantom, "\square"] & V \arrow[d, "f"] \\ X \arrow[r, "i"] & Y \end{tikzcd}$$

Let $N_X Y$ denote the normal bundle of the regular embedding, and let $N \coloneqq g^* N_X Y$ be the pullback of this normal bundle to W, which has rank d. Let $\pi\colon N \to W$ be the canonical projection. There is a closed embedding of the normal cone $C = C_W V$ into the vector bundle N. Because C is purely k-dimensional, it determines a k-cycle $[C]$ on N. We define the intersection product of V by X by intersecting $[C]$ with the zero section of N, yielding $X \cdot V \coloneqq s^*[C]$. Extending by linearity gives the generalized Gysin homomorphism:

$$\begin{aligned} i^! \colon Z_k(V) &\longrightarrow \mathrm{CH}_{k-d}(W); \\ \sum n_i [V_i] &\longmapsto \sum n_i (X \cdot V_i). \end{aligned}$$

The homomorphism $i^!$ factors as the composition:

$$Z_k(V) \xrightarrow{\sigma} Z_k(C_W V) \hookrightarrow \mathrm{CH}_k(N) \xrightarrow{s^*} \mathrm{CH}_{k-d}(W),$$

where σ is the specialization homomorphism mapping $[V_i] \mapsto [C_{V_i \cap X} V_i]$.

To prove that the generalized Gysin homomorphism descends to an operation on Chow groups, Fulton constructs the deformation space via blow-ups (deformation to the normal cone). This space provides a family of embeddings $X \to Y_t$ parametrized by $t \in \mathbb{P}^1$. For $t \neq \infty$, the embedding is the original $X \xrightarrow{i} Y$, and for $t = \infty$, it is the zero section of X in the normal cone $C_X Y$. This family interpolates between the given embedding and the zero section, ensuring the specialization homomorphism respects rational equivalence. For details, see [**Ful98**, Chapter 5].

This procedure yields the intersection product for a smooth scheme. When X is a smooth scheme, the diagonal $\Delta_X \colon X \to X \times X$ is a regular embedding. If V and W are subvarieties of X, one constructs their intersection product $V \cdot W$ by applying the generalized Gysin homomorphism to the following Cartesian diagram:

$$\begin{tikzcd} V \cap W \arrow[r] \arrow[d, hook] \arrow[dr, phantom, "\square"] & V \times W \arrow[d, hook] \\ X \arrow[r, "\Delta_X"] & X \times X \end{tikzcd}$$

For a smooth scheme X of dimension n, setting $\mathrm{CH}^d(X) \coloneqq \mathrm{CH}_{n-d}(X)$ equips $\mathrm{CH}^*(X)$ with the structure of a commutative ring under this intersection product, defining the Chow ring of X.

This highlights the primary difficulty in constructing the intersection product for Deligne–Mumford stacks. For a Deligne–Mumford stack $\mathcal{X}$, the diagonal $\mathcal{X} \to \mathcal{X} \times \mathcal{X}$ is only unramified (Theorem 2.3.30), not a closed immersion. Even when $\mathcal{X}$ is smooth, its diagonal is only a regular local embedding. On the other hand, blow-ups require a closed embedding; one cannot blow up a stack along a locus that is merely locally embedded. This geometric obstruction causes the classical deformation to the normal cone to fail at the first step.

To bypass this, Vistoli deviates from Fulton's framework. He proves that the generalized Gysin homomorphism for Deligne–Mumford stacks still descends to rational Chow groups, but his proof avoids the deformation to the normal cone entirely. We collect the main functorial consequences here and refer the reader to [**Vis89**] for the technical details.

Proposition 2.3.46 (Properties of Gysin homomorphisms)**.**

(1) Consider a Cartesian diagram

$$\begin{tikzcd} X' \arrow[r] \arrow[d, "p"] & Y' \arrow[d, "q"] \\ X \arrow[r] \arrow[d] & Y \arrow[d] \\ \mathcal{X} \arrow[r, "f"] & \mathcal{Y} \end{tikzcd}$$

where Y and Y' are schemes, and f is a regular local embedding of Deligne–Mumford stacks. If q is proper, then $f^!q_* = p_*f^!$. If q is flat, then $f^!q^* = p^*f^!$.

(2) Consider two regular local embeddings $f_i\colon \mathcal{X}_i \to \mathcal{Y}_i$ of Deligne–Mumford stacks for $i = 1, 2$. Let Y be a scheme, yielding the Cartesian diagram:

$$\begin{tikzcd} Z \arrow[r] \arrow[d] & X_2 \arrow[r] \arrow[d] & \mathcal{X}_2 \arrow[d, "f_2"] \\ X_1 \arrow[r] \arrow[d] & Y \arrow[r] \arrow[d] & \mathcal{Y}_2 \\ \mathcal{X}_1 \arrow[r, "f_1"] & \mathcal{Y}_1 & \end{tikzcd}$$

Then $f_1^! f_2^! = f_2^! f_1^!$.

We conclude this section by noting that while the classical framework of intersection theory does not readily generalize to algebraic stacks, the next chapter will utilize equivariant tools to develop intersection theory specifically for quotient stacks. The primary advantage of this equivariant approach is that it yields integer-valued intersection products on smooth Deligne–Mumford stacks. This provides the proper framework for retaining integral coefficients, bypassing the mandatory rational coefficients required in Vistoli's formulation. This equivariant foundation directly motivates Kresch's subsequent generalization of cycle groups and intersection theory to arbitrary Artin stacks [**Kre99**].

2.4. Deligne–Mumford Compactification $\overline{\mathcal{M}}_{g,n}$

In this final section, we apply the tools developed thus far to study the Deligne–Mumford compactification $\overline{\mathcal{M}}_{g,n}$ of the moduli stack of curves. Since this construction is extensively documented in standard texts such as [**HM98**, Chapter 3], [**Alp26**, Chapter 6], and [**ACG11**, Chapters X and XII], we provide a survey of the core ideas rather than exhaustive proofs. We conclude the section by defining the tautological classes on $\overline{\mathcal{M}}_{g,n}$, which form the geometric foundation of the ELSV formula.

We maintain the following convention throughout this section: whenever discussing the moduli stacks $\mathcal{M}_{g,n}$ and $\overline{\mathcal{M}}_{g,n}$, we assume the stability condition $2g - 2 + n > 0$ (see Theorem 2.3.12). Furthermore, we restrict our base field to the complex numbers $\mathbb{C}$. We will explicitly note when consequences diverge if the base field is not algebraically closed or of characteristic zero.

2.4.1. Geometry of Stable Curves. In the previous sections, we have thoroughly studied the moduli stack of g-genus curves $\mathcal{M}_g$ (Theorem 2.3.35). However, a drawback of $\mathcal{M}_g$—and more generally, $\mathcal{M}_{g,n}$—is that it is not a *proper* algebraic stack. For example, in the argument of Example 2.3.10, we saw that $\mathcal{M}_{0,n}$ fails to be a proper scheme when $n \geq 4$. This lack of properness makes performing intersection theory on $\mathcal{M}_{g,n}$ impossible, as we can no longer guarantee the existence of proper pushforwards. To remedy this, Mumford introduced the notion of stable curves to describe the boundary of the moduli stack. Together with Deligne [**DM69**], they successfully constructed a proper compactification of $\mathcal{M}_{g,n}$, which is known today as the Deligne–Mumford compactification.

To understand what the boundary of $\mathcal{M}_g$ should entail, we can examine the Legendre family of curves given by the homogeneous equations $\{y^2z = x(x-z)(x-\lambda z)\}_{\lambda \in \mathbb{A}^1 \setminus \{0,1\}}$. This defines a family of smooth genus 1 curves over the non-complete base scheme $S = \mathbb{A}^1_{\mathbb{C}} \setminus \{0,1\}$. The j-invariant of the fiber over λ is given by

$$j(\lambda) = 256\frac{(\lambda^2-\lambda+1)^3}{\lambda^2(\lambda-1)^2}.$$

This family cannot be extended to a family of smooth curves over a completion of S, as doing so would force the non-constant rational function $j(\lambda)$ to extend to a regular morphism from a projective curve to the affine line, which is impossible. As we approach the missing boundary points $\lambda \in \{0,1,\infty\}$, the smoothness of the curves inevitably breaks down. For instance, when $\lambda = 0$, the curve degenerates to $y^2z = x^2(x-z)$, acquiring a singular node at the projective point $[0:0:1]$.

This geometric breakdown indicates that the boundary of $\mathcal{M}_g$ must parameterize curves that are singular or reducible. Fortunately, the (Semi)-Stable Reduction Theorem (Proposition 2.4.17) guarantees that allowing curves with at worst nodal singularities is sufficient to close the space.

However, simply adjoining all nodal curves is too permissive and would prevent the compactified stack $\overline{\mathcal{M}}_g$ from remaining Deligne–Mumford, as arbitrary nodal curves can possess infinite automorphism groups. To ensure the moduli stack remains well-behaved, we restrict our boundary points to curves with finite automorphisms. This algebraic requirement motivates the definition of *stable curves*.

Definition 2.4.1 (Stable n-pointed curves)**.**

(1) An n-pointed curve is a curve C together with an ordered collection of $\mathbb{C}$-points $p_1, \dots, p_n \in C$, which we call the marked points. A point $q \in C$ is called *special* if q is either a node of C or one of the marked points p_i.

(2) An n-pointed connected, projective curve $(C, p_1, \dots, p_n)$ with at most nodal singularities is called *stable* if:
- The marked points $p_1, \dots, p_n \in C$ are distinct and lie in the smooth locus of C.
- The curve C is not an unpointed curve of arithmetic genus 1 (i.e., if $p_a(C) = 1$, then $n \geq 1$).
- Every smooth irreducible rational subcurve $E \cong \mathbb{P}^1 \subseteq C$ contains at least 3 special points.

If only the first two condition is satisfied, then we call the n-pointed curve $(C, p_1, \dots, p_n)$ *prestable*. If only the first condition is satisfied, then we call the n-pointed curve $(C, p_1, \dots, p_n)$ *quasi-stable*.

(3) A family of n-pointed stable curves over a scheme S is a proper, flat, and finitely presented morphism $\mathcal{C} \to S$ of schemes, together with n sections $\sigma_1, \dots, \sigma_n \colon S \to \mathcal{C}$, such that each geometric fiber $(\mathcal{C}_s, \sigma_1(s), \dots, \sigma_n(s))$ is an n-pointed stable curve.

Stable curves possess rich geometric properties, and there are several equivalent ways to characterize them. We start with the most fundamental characterization: an n-pointed prestable curve $(C, p_1, \dots, p_n)$ is stable if and only if its automorphism group $\mathrm{Aut}(C, p_1, \dots, p_n)$ is finite.

To understand this algebraically, we briefly recall the dualizing sheaf of a nodal curve. Because a nodal curve C is proper and a locally complete intersection (thanks to the local structure of a node), it possesses a dualizing line bundle ω_C. This comes equipped with a trace map $\mathrm{tr}\colon \mathrm{H}^1(C,\omega_C)\to\mathbb{C}$ such that for any coherent sheaf $\mathcal{F}$ over C, there is a perfect pairing

$$\mathrm{Hom}_{\mathcal{O}_C}(\mathcal{F},\omega_C)\times \mathrm{H}^1(C,\mathcal{F})\to \mathrm{H}^1(C,\omega_C)\xrightarrow{\mathrm{tr}}\mathbb{C}.$$

See [**Har77**, III. 7] for details. We next provide an explicit description of ω_C. Let $\Sigma\coloneqq C^{\mathrm{Sing}}$ be the singular locus of C and $U\coloneqq C\setminus\Sigma$, and let $\pi\colon\widetilde{C}\to C$ be the normalization of C. Let $\widetilde{U}$ and $\widetilde{\Sigma}$ be the preimages of U and Σ, respectively. This gives rise to the following commutative diagram:

$$\begin{tikzcd} \widetilde{U} \arrow[r, hook] \arrow[d] & \widetilde{C} \arrow[d, "\pi"] & \widetilde{\Sigma} \arrow[l, hook'] \arrow[d] \\ U \arrow[r, hook] & C & \Sigma \arrow[l, hook'] \end{tikzcd}$$

We fix an ordering of the singular locus $\Sigma=\{z_1,\dots,z_k\}$ and denote the preimages of each node as $\pi^{-1}(z_i)=\{p_i,q_i\}$. Note that $\omega_{\widetilde{C}}=\Omega_{\widetilde{C}}$ since the normalization $\widetilde{C}$ is a smooth curve. Let $D=\sum_{i=1}^k(p_i+q_i)$ be the corresponding Weil divisor on $\widetilde{C}$. We have a short exact sequence

$$0\to\Omega_{\widetilde{C}}\to\Omega_{\widetilde{C}}(D)\to\mathcal{O}_{\widetilde{\Sigma}}\to 0.$$

Since $\Omega_{\widetilde{C}}(D)|_{\widetilde{U}}=\Omega_{\widetilde{U}}$, we can interpret the global sections of $\Omega_{\widetilde{C}}(D)$ as rational sections of $\Omega_{\widetilde{C}}$ with at worst simple poles along the locus $\widetilde{\Sigma}$. Applying the global section functor, for any open set $\widetilde{V}\subseteq\widetilde{C}$, we obtain the exact sequence

$$0\to\Gamma(\widetilde{V},\Omega_{\widetilde{C}})\to\Gamma(\widetilde{V},\Omega_{\widetilde{C}}(\widetilde{\Sigma}))\xrightarrow{s\mapsto\mathrm{res}_y(s)}\bigoplus_{y\in\widetilde{V}\cap\widetilde{\Sigma}}\kappa(y).$$

Here, $\mathrm{res}_y(-)$ is the residue of a meromorphic section at a simple pole y. (For the existence of the residue map in the algebraic context, see [**Har77**, III. Theorem 7.14.1]).

We can now explicitly state the structure of the dualizing sheaf of a nodal curve C. The dualizing sheaf ω_C is the subsheaf $\omega_C\subseteq\pi_*\Omega_{\widetilde{C}}(D)$ whose sections over an open set $V\subseteq C$ consist of those rational differentials $s\in\Omega_{\widetilde{C}}(\widetilde{V})$ such that for every node $z_i\in V\cap\Sigma$ with preimages $p_i,q_i\in\widetilde{V}\cap\widetilde{\Sigma}$, the residues sum to zero:

$$\mathrm{res}_{p_i}(s)+\mathrm{res}_{q_i}(s)=0.$$

For a rigorous proof, see [**ACG11**, pp. 91–92]. The geometric intuition is clearest in the local affine case where $C=\operatorname{Spec}\mathbb{C}[x,y]/(xy)$. Here, the normalization is the disjoint union of two affine lines, $\widetilde{C}=\mathbb{A}^1\sqcup\mathbb{A}^1$. The global sections are $\Gamma(\widetilde{C},\Omega_{\widetilde{C}})=\Gamma(\mathbb{A}^1,\Omega_{\mathbb{A}^1})\times\Gamma(\mathbb{A}^1,\Omega_{\mathbb{A}^1})$, and every section of $\Gamma(C,\omega_C)$ takes the form

$$\left(f(x)\frac{\mathrm{d}x}{x},-g(y)\frac{\mathrm{d}y}{y}\right)$$

for polynomials $f,g\in\mathbb{C}[x]$ satisfying $f(0)=g(0)$. This condition is exactly what allows us to glue $(f,g)\in\Gamma(\widetilde{C},\mathcal{O}_{\widetilde{C}})$ into a global section $f(x)+g(y)-f(0)\in\Gamma(C,\mathcal{O}_C)$ at the node. This explicit calculation also demonstrates that ω_C is locally generated by $(\mathrm{d}x/x,-\mathrm{d}y/y)$, which is a line bundle.

Now we state and prove the fundamental characterizations of n-pointed stable curves.

Proposition 2.4.2. *Let $(C,p_1,\dots,p_n)$ be an n-pointed prestable curve. The following are equivalent*

(1) The prestable curve $(C,p_1,\dots,p_n)$ is stable.
(2) The automorphism group $\mathrm{Aut}(C,p_1,\dots,p_n)$ is finite.

(3) The twisted dualizing sheaf $\omega_C\left(\sum_{i=1}^n p_i\right)$ is ample.

Proof. To prove (1) $\Leftrightarrow$ (2), let $\pi\colon \widetilde{C} \to C$ be the normalization of C. Let $\widetilde{p_1}, \ldots, \widetilde{p_n} \in \widetilde{C}$ be the unique preimages of the marked points $p_i \in C$, and let $\widetilde{q_1}, \ldots, \widetilde{q_{2k}} \in \widetilde{C}$ be the preimages of the nodes of C. Any automorphism of C preserving the marked points lifts uniquely to an automorphism of $\widetilde{C}$ that preserves the set of marked points and permutes the nodal preimages. Since π is an isomorphism outside a finite set of nodes, an automorphism of C that lifts to the identity on $\widetilde{C}$ must be the identity on C. Therefore, $\mathrm{Aut}(C, p_1, \ldots, p_n)$ embeds injectively into the group of automorphisms of $\widetilde{C}$ preserving these special points. Consequently, the automorphism group of C is finite if and only if the automorphism group of every connected component of $\widetilde{C}$, fixing its special points, is finite.

Therefore, the problem reduces to the case where the curve C is smooth and connected. Following the calculation after Theorem 2.3.12, the automorphism group of C with m_C special points is finite if and only if $2g_C - 2 + m_C > 0$. This inequality is equivalent to the definition that the prestable curve C is stable.

To prove (1) $\Leftrightarrow$ (3), recall that the normalization map $\pi\colon \widetilde{C} \to C$ is a finite surjective morphism. A line bundle $\mathcal{L}$ on C is ample if and only if its pullback $\pi^*\mathcal{L}$ is ample on $\widetilde{C}$. Furthermore, a line bundle on a reduced scheme is ample if and only if its restriction to each of its irreducible components is ample (see [**Har77**, III. Ex 5.7]).

Let $T \subseteq \widetilde{C}$ be an irreducible component of the normalization. Because $\widetilde{C}$ is smooth, T is a smooth connected projective curve of genus g_T. The restriction of the pulled-back twisted dualizing sheaf to T is $\omega_T(D_T)$, where D_T is the divisor formed by the preimages of the nodes on T plus the marked points on T. The degree of this line bundle on T is $\deg(\omega_T(D_T)) = 2g_T - 2 + |D_T|$. By [**Har77**, IV. Corollary 3.3], a line bundle on a smooth projective curve is ample if and only if it has positive degree. As analyzed above, the condition $2g_T - 2 + |D_T| > 0$ for all irreducible components T is equivalent to the statement that the original curve C is stable. Therefore, the twisted dualizing sheaf is ample if and only if the original curve is stable. $\square$

In fact, we can establish much stronger positivity results for the twisted dualizing sheaf. This positivity is the essential geometric tool required to construct the compactified moduli stack $\overline{\mathcal{M}}_{g,n}$ using GIT, as it guarantees that stable curves can be uniformly embedded into a locally closed subscheme of a Hilbert scheme, which is a projective scheme. The proof of the following proposition can be found in [**ACG11**, Chapter X, Lemma 6.1].

Proposition 2.4.3. *Let $(C, p_1, \ldots, p_n)$ be an n-pointed stable curve and let $L := \omega_C\left(\sum_{i=1}^n p_i\right)$. Then the tensor power $L^{\otimes k}$ is very ample for all $k \geq 3$, and $\mathrm{H}^1(C, L^{\otimes k}) = 0$ for all $k \geq 2$.*

After the fundamental algebraic characterization, we introduce a useful geometric perspective: a prestable curve satisfying the global condition $2g - 2 + n > 0$ can be naturally "stabilized" by systematically contracting all of its rational tails and rational bridges.

Definition 2.4.4 (Rational tails and bridges)**.** Let $(C, p_1, \ldots, p_n)$ be a prestable curve over $\mathbb{C}$. An irreducible rational subcurve $E \cong \mathbb{P}^1 \subseteq C$ with a nonempty complement $E^c := \overline{C \setminus E}$ is called

- a *rational tail* if the intersection number $E \cdot E^c = 1$ and E contains no marked points;
- a *rational bridge* if $E \cdot E^c = 2$ and E contains no marked points, or if $E \cdot E^c = 1$ and E contains exactly one marked point.

Clearly, a prestable curve over $\mathbb{C}$ is stable if and only if it does not contain any rational tails or rational bridges. This is because tails and bridges are exactly the rational components containing

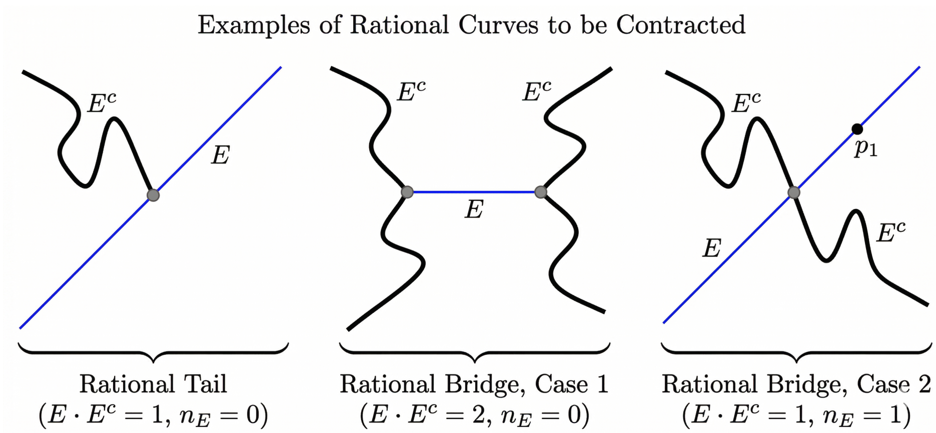

FIGURE 1. Rational Tails and Bridges

fewer than 3 special points. Moreover, for any prestable curve C satisfying $2g-2+n>0$, we can always stabilize the curve by contracting all such components.

Proposition 2.4.5. *Let $(C, p_1, \ldots, p_n)$ be an n-pointed prestable curve and $E \subseteq C$ be a rational tail or rational bridge. Then there is a canonical morphism $c\colon C \to C'$ contracting E to a point. Moreover, C' is identified with the pushout* $\operatorname{Spec}\Gamma(E, \mathcal{O}_E) \sqcup_E C$.

PROOF. In both cases, we construct C' via the following Ferrand pushout diagram:

$$\begin{tikzcd}
E \cap E^c \arrow[r] \arrow[d] & E^c \arrow[d] \\
E \arrow[r] \arrow[d] & C \arrow[d, "c"] \\
\operatorname{Spec}\Gamma(E, \mathcal{O}_E) \arrow[r] & C'
\end{tikzcd}$$

Here, both squares are pushout squares, and $c\colon C \to C'$ is defined as the unique map making the diagram commute. Since $E \cong \mathbb{P}^1$, its global sections ring $\Gamma(E, \mathcal{O}_E)$ is $\mathbb{C}$. Therefore, the bottom-left term $\operatorname{Spec}\Gamma(E, \mathcal{O}_E)$ is simply a point ($\operatorname{Spec}\mathbb{C}$), which achieves the contraction of the subcurve E. $\square$

As a corollary, if (C, p_i) is an n-pointed prestable curve of genus g such that $2g-2+n>0$, then there is a canonical morphism $c\colon C \to C^{\mathrm{st}}$, called the stable contraction, contracting all rational tails and bridges to points, such that $(C^{\mathrm{st}}, c(p_i))$ is an n-pointed stable curve of genus g. The morphism is obtained by iteratively applying Proposition 2.4.5 until no rational tails or bridges remain.

Finally, we introduce a third way to describe stable curves using the notion of dual graphs. This provides a powerful combinatorial characterization that elegantly encodes the geometry of nodal curves. Furthermore, dual graphs naturally induce a stratification of the compactified moduli stack $\overline{\mathcal{M}}_{g,n}$, which we will construct later.

A vertex-weighted, n-marked graph $\Gamma = (G, w, m)$ consists of a finite, connected, undirected multigraph G (allowing loops and multiple edges) with vertex set $V(G)$ and edge set $E(G)$, a weight function $w\colon V(G) \to \mathbb{N}$, and a marking function $m\colon \{1, \ldots, n\} \to V(G)$.

Definition 2.4.6 (Dual Graph)**.** The dual graph of an n-pointed prestable curve $(C, p_1, \ldots, p_n)$ is defined as $\Gamma = (G, w, m)$, where the vertices $v_i \in V(G)$ correspond to the irreducible components C_i of C, and the weight $w(v_i)$ is the geometric genus of the normalization $\widetilde{C}_i$. For each node q of C intersecting components C_i and C_j (where possibly $i = j$), there is an edge $e_q \in E(G)$ connecting v_i and v_j. The marking function $m(k) = v_i$ indicates that the marked point p_k lies on the smooth locus of the component C_i.

The power of dual graphs lies in their ability to completely capture the geometry of a prestable curve, particularly its arithmetic genus. Recall that for a connected, nodal, projective curve C over $\mathbb{C}$ with δ nodes and μ irreducible components $C_1, \ldots, C_\mu$, the arithmetic genus is given by

$$p_a(C) = \sum_{i=1}^{\mu} g(\widetilde{C}_i) + \delta - \mu + 1$$

where $\widetilde{C}_i$ is the smooth normalization of C_i. See [**Har77**, IV. Ex 1.8].

Translating this to the dual graph, the sum of the geometric genera $\sum_{i=1}^{\mu} g(\widetilde{C}_i)$ corresponds to the total weight $\sum_{v \in V(G)} w(v)$. The term $\delta - \mu + 1$ corresponds exactly to $|E(G)| - |V(G)| + 1$, which is the first Betti number $b_1(G)$ of the graph viewed as a 1-dimensional CW complex. This links the geometry of the prestable curve to the topology of its dual graph. We therefore define the genus of the dual graph as

$$g(\Gamma) \coloneqq \sum_{v \in V(G)} w(v) + b_1(G).$$

From the proof of Proposition 2.4.2, we also see that the prestable curve C is stable if and only if

$$2w(v) - 2 + \deg(v) + |m^{-1}(v)| > 0 \quad \text{for all } v \in V(G),$$

where $\deg(v)$ is the valence of the vertex (the number of incident half-edges, meaning loops contribute 2 to the degree). If this inequality holds for all vertices, we say the marked graph $\Gamma = (G, w, m)$ itself is *stable*.

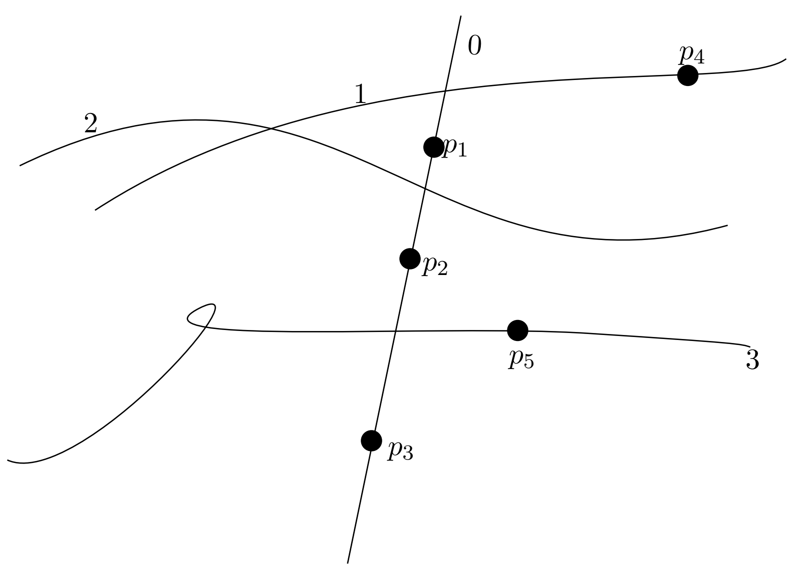

prestable curve

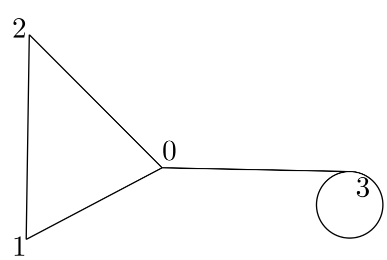

dual graph

2.4.2. The Moduli Space of n-Pointed Stable Curves. Having studied the basic properties of stable curves, we next generalize these concepts to families of n-pointed prestable curves and examine their deformation theory. This leads us to the construction of the moduli stack of n-pointed stable curves of genus g, denoted $\overline{\mathcal{M}}_{g,n}$. We will also explore the rich geometric properties admitted by this moduli stack. More precisely, let $\overline{\mathcal{M}}_{g,n}$ be the prestack of n-pointed families $(\mathcal{C} \to S, \sigma_1, \ldots, \sigma_n)$ of stable curves. The primary goal of this subsection is to outline the ideas behind the following fundamental theorem.

Theorem 2.4.7. *Assuming $2g - 2 + n > 0$, the moduli stack $\overline{\mathcal{M}}_{g,n}$ is a smooth, proper, and irreducible Deligne–Mumford stack of finite type whose dimension is $3g - 3 + n$.*

To establish this theorem, we first construct $\overline{\mathcal{M}}_{g,n}$ as an algebraic stack. There is a natural inclusion of prestacks $\overline{\mathcal{M}}_{g,n} \subseteq \mathcal{M}^{\text{all}}_{g,n}$, where $\mathcal{M}^{\text{all}}_{g,n}$ is the prestack parameterizing all curves (whose precise definition we introduce below). Our strategy is to prove that the prestack $\mathcal{M}^{\text{all}}_{g,n}$ is in fact an algebraic stack of locally finite type, and that the inclusion $\overline{\mathcal{M}}_{g,n} \subseteq \mathcal{M}^{\text{all}}_{g,n}$ is an open immersion. This immediately implies the algebraicity of $\overline{\mathcal{M}}_{g,n}$.

Definition 2.4.8 (Families of curves)**.**

(1) A family of n-pointed curves over a scheme S is a proper, flat, and finitely presented morphism $\mathcal{C} \to S$ of **algebraic spaces**, together with n sections $\sigma_1, \ldots, \sigma_n \colon S \to \mathcal{C}$, such that every geometric fiber $\mathcal{C}_s$ is a curve over $\operatorname{Spec} \kappa(s)$.

(2) Let $\mathcal{M}^{\text{all}}_{g,n}$ be the prestack over $\mathsf{Sch}_{\text{ét}}$ where an object over a scheme S is a family of curves $\mathcal{C} \to S$ of genus g with n sections $\sigma_1, \ldots, \sigma_n \colon S \to \mathcal{C}$. A morphism $(\mathcal{C}' \to S', \sigma'_1, \ldots, \sigma'_n) \to (\mathcal{C} \to S, \sigma_1, \ldots, \sigma_n)$ is the data of a Cartesian diagram:

$$\begin{tikzcd} \mathcal{C}' \arrow[r, "g"] \arrow[d] \arrow[dr, phantom, "\square"] & \mathcal{C} \arrow[d] \\ S' \arrow[r, "f"] \arrow[u, bend left, "\sigma'_i"] & S \arrow[u, bend left, "\sigma_i"] \end{tikzcd}$$

where the sections are compatible, meaning $g \circ \sigma'_i = \sigma_i \circ f$ for all $1 \leq i \leq n$.

Remark 2.4.9. In Definition 2.4.8, we allow the total space $\mathcal{C}$ to be an algebraic space rather than a scheme. This generalization is necessary because there exist families of prestable genus 0 curves and smooth genus 1 curves where the total space is an algebraic space but not a scheme (as mentioned in Example 2.3.11).

However, if $(\mathcal{C} \to S, \sigma_1, \ldots, \sigma_n)$ is a family of *stable* curves, the morphism $\mathcal{C} \to S$ is necessarily projective because the twisted relative dualizing sheaf $\omega_{\mathcal{C}/S}\left(\sum_{i=1}^n \sigma_i(S)\right)$ is relatively ample over S by Proposition 2.4.10. Consequently, for families of stable curves, $\mathcal{C}$ is always a scheme.

We first generalize the characterization of a single stable curve (Proposition 2.4.2) to a family of stable curves. Because each condition can be checked fiberwise, the proof of the statement for a single curve can be adapted to families without major changes. See [**ACG11**, Chapter X, Proposition 6.7] for more details.

Proposition 2.4.10 (Characterization of Families of Stable Curves)**.** *Let $(\mathcal{C} \to S, \sigma_i)$ be a family of n-pointed prestable curves of genus g over a scheme S. The following are equivalent:*

(1) $(\mathcal{C} \to S, \sigma_i)$ is a family of stable curves.

(2) The group scheme $\underline{\operatorname{Aut}}(\mathcal{C}/S, \sigma_1, \ldots, \sigma_n)$ is quasi-finite over S.

(3) The relative twisted dualizing sheaf $\omega_{\mathcal{C}/S}\left(\sum_{i=1}^n \sigma_i\right)$ is relatively ample over S.

Moreover, if $(\pi\colon \mathcal{C} \to S, \sigma_i)$ is a family of n-pointed stable curves of genus g and we let $L \coloneqq \omega_{\mathcal{C}/S}\left(\sum_{i=1}^n \sigma_i\right)$, then the tensor power $L^{\otimes k}$ is very ample relative to S, and the pushforward $\pi_*(L^{\otimes k})$ is a vector bundle of rank $(2k-1)(g-1)+kn$ for all $k \geq 3$.

Using this characterization, we can establish the fact that the inclusion $\overline{\mathcal{M}}_{g,n} \subseteq \mathcal{M}_{g,n}^{\text{all}}$ is an open immersion by proving that stability is an open condition.

Proposition 2.4.11 (Openness of Stability)**.** *Let $(\pi\colon \mathcal{C} \to S, \sigma_1, \ldots, \sigma_n)$ be a family of n-pointed curves. The locus of points $s \in S$ such that the geometric fiber $(\mathcal{C}_s, \sigma_i(s))$ is stable is an open subset of S.*

Sketch of Proof. We may assume that $\mathcal{C}$ is a scheme by replacing it with a dense open subscheme. We first claim that the locus of nodes is open. For a purely one-dimensional curve C, a closed point $p \in C$ is a node if and only if C is a local complete intersection at p and the singular locus $\mathrm{Sing}(C)$ is unramified at p. See [**Sta26**, Tag 0C4E] for details. Therefore, $p \in \mathcal{C}_s$ is a node if and only if the morphism $\mathcal{C} \to S$ is syntomic at p and $\mathrm{Sing}(\mathcal{C}/S) \to S$ is unramified at p. The claim then follows from the openness of both syntomicity and unramifiedness.

Next, we observe that the condition of being prestable is an open condition within a family of nodal curves. This simply follows from the fact that the locus in S where the sections $\sigma_i(s)$ are distinct and lie in the smooth locus of $\mathcal{C}_s$ is open.

As a final step, it suffices to verify that the stability condition itself is an open condition within an n-pointed prestable family. This is a direct consequence of the openness of ampleness. More precisely, the locus of points $s \in S$ such that the restriction $\omega_{\mathcal{C}/S}\left(\sum_{i=1}^n \sigma_i\right)\big|_{\mathcal{C}_s}$ is ample is an open locus. Applying Proposition 2.4.10 then yields the desired result. See [**Sta26**, Tag 0D3D] for the details regarding the openness of ampleness. $\square$

To establish the algebraicity of $\overline{\mathcal{M}}_{g,n}$, the only remaining step is to prove that the prestack $\mathcal{M}_{g,n}^{\text{all}}$ is an algebraic stack. This sufficiency follows from the openness of stability established in Proposition 2.4.11.

Proposition 2.4.12 (Algebraicity of $\mathcal{M}_{g,n}^{\text{all}}$)**.** *The prestack $\mathcal{M}_{g,n}^{\text{all}}$ is an algebraic stack locally of finite type. As a corollary, $\overline{\mathcal{M}}_{g,n}$ is an open algebraic substack of $\mathcal{M}_{g,n}^{\text{all}}$, which is also locally of finite type.*

Sketch of Proof. We first show that $\mathcal{M}_{g,n}^{\text{all}}$ is a stack over $\mathsf{Sch}_{\text{ét}}$. Let $\{S_i \to S\}$ be an étale cover of S, and let $(\mathcal{C}_i \to S_i, \sigma_{i,1}, \ldots, \sigma_{i,n})$ be a family of n-pointed curves over each S_i. Suppose we are given isomorphisms $\alpha_{ij}\colon \mathcal{C}_i|_{S_{ij}} \xrightarrow{\cong} \mathcal{C}_j|_{S_{ij}}$ that are compatible with the sections and satisfy the cocycle condition. The local spaces $\mathcal{C}_i$ glue into a global algebraic space $\mathcal{C}$ over S because the quotient of the étale equivalence relation $\bigsqcup_{i,j} \mathcal{C}_{ij} \rightrightarrows \bigsqcup_i \mathcal{C}_i$ is an algebraic space $\mathcal{C}$ ([**Alp26**, Corollary 5.5.12]), where the two maps are p_1 and $p_2 \circ \alpha_{ij}$, respectively. Furthermore, by the étale descent of morphisms, the local sections glue into global sections $\sigma_1, \ldots, \sigma_n\colon S \to \mathcal{C}$ such that $\sigma_k|_{S_i} = \sigma_{k,i}$. Therefore, the local families uniquely glue into a global family $(\mathcal{C} \to S, \sigma_1, \ldots, \sigma_n)$, showing that $\mathcal{M}_{g,n}^{\text{all}}$ is a stack.

To prove algebraicity, we may assume $n = 0$ without loss of generality. This reduction is justified because the forgetful morphism $\mathcal{M}_{g,n+1}^{\text{all}} \to \mathcal{M}_{g,n}^{\text{all}}$ is representable. Specifically, if a morphism $S \to \mathcal{M}_{g,n}^{\text{all}}$ is classified by a family $(\mathcal{C} \to S, \sigma_1, \ldots, \sigma_n)$, we have a natural isomorphism $S \times_{\mathcal{M}_{g,n}^{\text{all}}} \mathcal{M}_{g,n+1}^{\text{all}} \cong \mathcal{C}$, since the choice of the new section σ_{n+1} is equivalent to choosing a point in the fibers of $\mathcal{C}$ over S. Therefore, if we assume $\mathcal{M}_{g,n}^{\text{all}}$ is algebraic and admits a smooth presentation

$U \to \mathcal{M}_{g,n}^{\mathrm{all}}$ from a scheme U, then the base change $U' \coloneqq U \times_{\mathcal{M}_{g,n}^{\mathrm{all}}} \mathcal{M}_{g,n+1}^{\mathrm{all}} \to \mathcal{M}_{g,n+1}^{\mathrm{all}}$ is a surjective, smooth, and representable morphism from an algebraic space U'. Choosing a scheme presentation $V \to U'$ demonstrates that $\mathcal{M}_{g,n+1}^{\mathrm{all}}$ is also an algebraic stack.

The algebraicity of the unpointed stack $\mathcal{M}_g^{\mathrm{all}}$ can then be established by verifying Artin's axioms for algebraicity ([**Alp26**, Theorem C.7.7]). We refer to [**Sta26**, Tag 0D5A] for the details of this verification. □

We also remark that, rather than viewing $\overline{\mathcal{M}}_{g,n}$ as an open substack of $\mathcal{M}_{g,n}^{\mathrm{all}}$, we can also construct it using GIT. By Proposition 2.4.10, for any family of stable curves $(\mathcal{C} \to S, \sigma_1, \ldots, \sigma_n)$, the twisted relative dualizing sheaf $\omega_{\mathcal{C}/S}(\sum_i \sigma_i)^{\otimes 3}$ is relatively very ample, and its pushforward $E \coloneqq \pi_* \left(\omega_{\mathcal{C}/S}\left(\sum_i \sigma_i\right)^{\otimes 3}\right)$ is a vector bundle. This induces a canonical closed embedding $\mathcal{C} \hookrightarrow \mathbb{P}(E)$ over S. Therefore, one can mimic the proof of Theorem 2.3.9 and show that $\overline{\mathcal{M}}_{g,n} \cong [H'/\operatorname{PGL}_N]$, where H' is a locally closed subscheme of a suitable Hilbert scheme H parametrizing pluricanonically embedded curves. For the details of this GIT construction, see [**HM98**, Section 4.C]. Note that this construction also proves that $\overline{\mathcal{M}}_{g,n}$ is of finite type.

Having established the algebraicity (and finite type) of $\overline{\mathcal{M}}_{g,n}$, the next step is to prove that it is in fact a Deligne–Mumford stack. Our primary tool for this is the characterization of Deligne–Mumford stacks established in Theorem 2.3.30. By analyzing the deformation theory of stable curves to verify the finiteness of automorphisms, we will simultaneously establish the smoothness and the dimension of $\overline{\mathcal{M}}_{g,n}$.

Proposition 2.4.13 (Deformation Theory of Stable Curves)**.** *Let $A' \to A$ be a surjection of Artinian local rings with residue field k. Suppose that $J \coloneqq \operatorname{Ker}(A' \to A)$ satisfies $\mathfrak{m}_{A'}J = 0$. Let $(\mathcal{C} \to \operatorname{Spec} A, \sigma_1, \ldots, \sigma_n)$ be a family of prestable curves over A, which is a deformation of an n-pointed prestable curve $(C, p_1, \ldots, p_n)$ over k. Then*

(1) The group of automorphisms of a deformation of $(\mathcal{C} \to \operatorname{Spec} A, \sigma_i)$ over A' is isomorphic to
$$\operatorname{Ext}^0_{\mathcal{O}_C}\left(\Omega_C\left(\sum_{i=1}^n p_i\right), \mathcal{O}_C \otimes_k J\right).$$
(2) There are no obstructions to deforming $(\mathcal{C} \to \operatorname{Spec} A, \sigma_1, \ldots, \sigma_n)$ over A'. The set of isomorphism classes of deformations over A' is a torsor under
$$\operatorname{Ext}^1_{\mathcal{O}_C}\left(\Omega_C\left(\sum_{i=1}^n p_i\right), \mathcal{O}_C \otimes_k J\right).$$
Furthermore, the dimensions of these vector spaces satisfy
$$\dim_k \operatorname{Ext}^1_{\mathcal{O}_C}\left(\Omega_C\left(\sum_{i=1}^n p_i\right), \mathcal{O}_C\right) - \dim_k \operatorname{Ext}^0_{\mathcal{O}_C}\left(\Omega_C\left(\sum_{i=1}^n p_i\right), \mathcal{O}_C\right) = 3g-3+n.$$
(3) Moreover, if $(C, p_1, \ldots, p_n)$ is stable, then $\operatorname{Ext}^0_{\mathcal{O}_C}\left(\Omega_C\left(\sum_{i=1}^n p_i\right), \mathcal{O}_C\right) = 0$.

Corollary 2.4.14. *The stack $\overline{\mathcal{M}}_{g,n}$ is a smooth Deligne–Mumford stack of pure dimension $3g-3+n$.*

Proof. The proof proceeds identically to the proof of Theorem 2.3.35. Let $(C, p_1, \ldots, p_n)$ be an n-pointed stable curve over a field k. By Proposition 2.4.13, $\operatorname{Ext}^0_{\mathcal{O}_C}(\Omega_C(\sum p_i), \mathcal{O}_C) = 0$. This implies that the Lie algebra of the group scheme $\underline{\operatorname{Aut}}(C, p_1, \ldots, p_n)$ is trivial, so $\underline{\operatorname{Aut}}(C, p_i)$ is a finite and reduced group scheme. By Theorem 2.3.30, this finiteness guarantees that $\overline{\mathcal{M}}_{g,n}$ is a Deligne–Mumford stack.

Furthermore, because $\mathrm{Ext}^2_{\mathcal{O}_C}(\Omega_C(\sum p_i), \mathcal{O}_C) = 0$ for a one-dimensional scheme, there are no obstructions to higher-order deformations of stable curves. By the infinitesimal lifting criterion (Theorem 2.3.36), this unobstructedness proves that $\overline{\mathcal{M}}_{g,n}$ is smooth.

Finally, by Proposition 2.3.37, the dimension of the smooth stack $\overline{\mathcal{M}}_{g,n}$ is equal to the dimension of its tangent space at the point $[(C, p_1, \dots, p_n)]$. This tangent space is isomorphic to the first-order deformation space $\mathrm{Ext}^1_{\mathcal{O}_C}(\Omega_C(\sum p_i), \mathcal{O}_C)$. Since $\mathrm{Ext}^0 = 0$, the dimension formula in Proposition 2.4.13 yields $\dim \overline{\mathcal{M}}_{g,n} = \dim_k \mathrm{Ext}^1 = 3g - 3 + n$. $\square$

We are left to verify the properness and irreducibility of $\overline{\mathcal{M}}_{g,n}$. To establish properness, recall that in classical algebraic geometry, a powerful way to verify the properness of a quasi-compact morphism of Noetherian schemes is the valuative criterion, which provides a geometric strategy to verify separatedness, properness, and universal closedness. This criterion generalizes to stack theory, but with a twist. More precisely:

Theorem 2.4.15 (Valuative Criterion for Properness and Separatedness)**.** *Let $f\colon \mathcal{X} \to \mathcal{Y}$ be a quasi-compact and quasi-separated morphism of locally Noetherian algebraic stacks. Consider a 2-commutative diagram*

$$\begin{tikzcd} \operatorname{Spec} K \arrow[r] \arrow[d] & \mathcal{X} \arrow[d] \\ \operatorname{Spec} R \arrow[r] & \mathcal{Y} \end{tikzcd}$$

where R is a DVR with fraction field K. Then:

(1) The morphism f is proper if and only if f is of finite type, and for every diagram as above, there exists an extension $R \to R'$ of DVRs (with fraction fields $K \to K'$) and a lifting as below, which is unique up to unique isomorphism:

$$\begin{tikzcd} \operatorname{Spec} K' \arrow[r] \arrow[d] & \operatorname{Spec} K \arrow[r] \arrow[d] & \mathcal{X} \arrow[d, "f"] \\ \operatorname{Spec} R' \arrow[r] \arrow[urr, dashed] & \operatorname{Spec} R \arrow[r] & \mathcal{Y} \end{tikzcd}$$

(2) The morphism f is separated if and only if any two such liftings over any extension R' are uniquely isomorphic.

For the proof of this theorem, we refer to [**Sta26**, Tag 0CLY]. Note that compared to the classical case for schemes ([**Har77**, Theorem II.4.7]), we are forced to allow finite extensions of the DVR. We first use an explicit geometric example to demonstrate why this ring extension is necessary, and then conceptually explain why algebraic stacks require such base changes.

Example 2.4.16 (Base change is necessary.)**.** Consider the classifying stack $\mathsf{B}\mu_n$ over $\mathbb{C}$, where $\mu_n = \mathrm{Spec}(\mathbb{C}[t]/(t^n - 1))$ is the finite group scheme of n-th roots of unity. Let $R = \mathbb{C}[[x]]$ be the DVR with fraction field $K = \mathbb{C}((x))$. By the structure of Kummer extensions, isomorphism classes of principal μ_n-bundle over $\operatorname{Spec} K$ corresponds to elements of the group $K^\times/(K^\times)^n$. The element $x \in K^\times$ therefore defines a non-trivial principal μ_n-bundle over K, which corresponds to the field extension $L = K(x^{1/n})$ over K. This classifying map $\operatorname{Spec} K \to \mathsf{B}\mu_n$ does not extend to a map $\operatorname{Spec} R \to \mathsf{B}\mu_n$ since every principal μ_n-bundles over $\operatorname{Spec} R$ is trivial.

However, consider the extension of DVRs $R \to R'$ where $R' = \mathbb{C}[[t]]$ is given by the ramified covering $t^n = x$. The new fraction field is $K' = \mathbb{C}((t))$. When we pull back our original μ_n-bundle along $\operatorname{Spec} K' \to \operatorname{Spec} K$, we evaluate the class of x inside $(K')^\times/((K')^\times)^n$. Since $x = t^n$ in K', it becomes a trivial class, so the classifying map $\operatorname{Spec} K' \to \mathsf{B}\mu_n$ extends to the trivial classifying map $\operatorname{Spec} R' \to \mathsf{B}\mu_n$.

Now we provide a conceptual interpretation of the necessity of this base change. In classical algebraic geometry, the valuative criterion for properness has a geometric translation: let X be a nonsingular curve and $P \in X$ be a closed point. The local ring at P is a DVR, and the punctured curve $X \setminus \{P\}$ corresponds to its fraction field. To say a scheme Y is proper means that any morphism from the punctured curve $X \setminus \{P\} \to Y$ can be uniquely extended to a morphism $X \to Y$ to fill in the missing point P ([**Har77**, I. Proposition 6.8]). However, in the realm of algebraic stacks, a "point" (more precisely, a field-valued point $\mathcal{X}(k)$) is no longer merely a single point in the underlying topological space; it is a groupoid that carries the data of its non-trivial automorphisms. As a family of objects degenerates toward a central limit point possessing such automorphisms, the family can non-trivially "twist" around the puncture, generating monodromy. This inherent twisting is what causes a naive extension over the original DVR to fail. By allowing a finite base extension of the DVR, which geometrically corresponds to pulling the family back to a ramified cover of the curve, we "resolve" the monodromy created by the non-trivial automorphisms, untwisting the family and finally allowing us to complete the extension.

Returning to the proof of the properness of $\overline{\mathcal{M}}_{g,n}$, our main ingredient is the generalized valuative criterion for properness discussed above. We first establish the existence of the lifting.

Proposition 2.4.17 (Stable Reduction). *Let R be a DVR with fraction field K. Denote $\Delta = \operatorname{Spec} R$ and $\Delta^* = \operatorname{Spec} K$. If $(\mathcal{C}^* \to \Delta^*, \sigma_1^*, \dots, \sigma_n^*)$ is a family of n-pointed stable curves, then there exists an extension of DVRs $R \to R'$ and an n-pointed family $(\mathcal{C} \to \Delta' := \operatorname{Spec} R', \sigma_1, \dots, \sigma_n)$ of stable curves extending the base change of $(\mathcal{C}^* \to \Delta^*, \sigma_1^*, \dots, \sigma_n^*)$ to K', the fraction field of R'.*

Proof. The proof is highly non-trivial and is divided into six main steps. See [**HM98**, Section 3.C] for more details.

Step I: Reduce to the case where $\mathcal{C}^* \to \Delta^*$ is smooth. The key tool in this step is the consequence we proved in Proposition 2.4.2, which states that a prestable curve is stable if and only if the connected components of its normalization are stable. Suppose that the generic fiber $\mathcal{C}^*$ has δ nodes. After finite base changes if necessary, we may arrange that the j-th node is given by a K-valued point $n_j^* \in \mathcal{C}^*$ whose preimage under the normalization $\widetilde{\mathcal{C}^*} \to \mathcal{C}^*$ consists of two distinct points q_j^* and r_j^*. Let $(\widetilde{\mathcal{C}}^* \to \Delta^*, \sigma_i^*, q_j^*, r_j^*)$ be the pointed normalization, and let $(\widetilde{\mathcal{C}}_k^* \to \Delta^*, q_{kl}^*)$ be its connected components, where $\{q_{kl}^*\} = \{\sigma_i^*, q_j^*, r_j^*\} \cap \widetilde{\mathcal{C}}_k^*$. Each of these components is also stable.

Assuming that the stable reduction theorem holds for smooth families, we can construct a stable family $(\widetilde{\mathcal{C}}_k \to \Delta, q_{kl})$ extending $(\widetilde{\mathcal{C}}_k^* \to \Delta^*, q_{kl}^*)$ after further base changes. To extend the original family $\mathcal{C}^* \to \Delta^*$, we simply glue the sections q_{il} and $q_{i'l'}$ corresponding to the preimages q_j^* and r_j^*. This gluing produces an n-pointed family $(\mathcal{C} \to \Delta, \sigma_i)$ of nodal curves with δ additional sections tracking the nodes. By construction, the pointed normalization of the central fiber $\mathcal{C}_0$ is the disjoint union of the stable central fibers $(\widetilde{\mathcal{C}}_k)_0$. Therefore, the central fiber $\mathcal{C}_0$ is stable, which implies the stability of the entire family $(\mathcal{C} \to \Delta, \sigma_i)$.

Step II: Find a flat extension $(\mathcal{C} \to \Delta, \sigma_i)$ for the smooth family. By Step I, we may now assume without loss of generality that the original family $\mathcal{C}^* \to \Delta^*$ is smooth. By Proposition 2.4.10, the twisted dualizing sheaf $\left(\omega_{\mathcal{C}^*/\Delta^*}\left(\sum_i \sigma_i^*\right)\right)^{\otimes 3}$ is relatively very ample, inducing a closed immersion $\mathcal{C}^* \hookrightarrow \mathbb{P}^N_{\Delta^*}$. By the flatness criterion over Dedekind domains ([**Har77**, Proposition III.9.7]), the scheme-theoretic image $\mathcal{C}$ of $\mathcal{C}^* \hookrightarrow \mathbb{P}^N_{\Delta}$ is flat over Δ. This gives a projective family of curves $\mathcal{C} \to \Delta$

extending $\mathcal{C}^* \to \Delta^*$. Furthermore, by the properness of $\mathcal{C}$, the sections $\sigma_i^* \colon \Delta^* \to \mathcal{C}^*$ uniquely extend to sections $\sigma_i \colon \Delta \to \mathcal{C}$ via the valuative criterion for properness.

The primary difficulty at this stage is that the central fiber $\mathcal{C}_0$ may be highly singular, and the marked points $\sigma_i(0) \in \mathcal{C}_0$ may intersect each other or lie in the singular locus. Our next goal is to use resolution of singularities and base changes to tame these issues.

Step III: Replace $\mathcal{C}$ with a resolution of singularities. By successive blow-ups of the total space, there exists a projective birational morphism $\widetilde{\mathcal{C}} \to \mathcal{C}$ from a regular two-dimensional scheme $\widetilde{\mathcal{C}}$ which is an isomorphism over Δ^*, such that the reduced central fiber $(\widetilde{\mathcal{C}}_0)_{\mathrm{red}}$ is a divisor with normal crossings (a nodal curve). Since $\widetilde{\mathcal{C}}$ is regular and maps dominantly to the curve Δ, the morphism $\widetilde{\mathcal{C}} \to \Delta$ is flat. We then replace $\mathcal{C}$ with $\widetilde{\mathcal{C}}$ and the sections σ_i with their strict transforms.

Step IV: Take a base change and normalize to reduce the central fiber. After the first three steps, $\mathcal{C}$ is a regular surface and the central fiber $\mathcal{C}_0$ has no embedded points. Moreover, $(\mathcal{C}_0)_{\mathrm{red}}$ is an effective Cartier divisor. Therefore, étale locally, the map $\mathcal{C} \to \Delta$ can be explicitly described as follows:

- If $p \in (\mathcal{C}_0)_{\mathrm{red}}$ is a smooth point of the fiber, the map $\mathcal{C} \to \Delta$ near p is given by $x^a = t$, where a is the multiplicity of the irreducible component containing p.
- If $p \in (\mathcal{C}_0)_{\mathrm{red}}$ is a node, the map is locally given by $x^a y^b = t$, where the two components intersecting at p have multiplicities a and b.

Let N be the least common multiple of the multiplicities of all irreducible components of $\mathcal{C}_0$. We may assume, after a preliminary base change, that R contains a primitive N-th root of unity ρ. Let $\Delta' \to \Delta$ be a totally ramified extension of DVRs such that the uniformizer $t \in R$ pulls back to $(t')^N$ for a uniformizer $t' \in R'$. Let $\mathcal{C}' := \mathcal{C} \times_\Delta \Delta'$, and let $p' \in \mathcal{C}'$ be the unique preimage of p.

We now verify that this base change resolves the multiplicities. We start with the case where $p \in (\mathcal{C}_0)_{\mathrm{red}}$ is smooth. The space $\mathcal{C}'$ is defined étale locally by $x^a = (t')^N$ near p'. Because R is defined over a field of characteristic zero, we have the factorization:

$$x^a - (t')^N = \prod_{i=0}^{a-1} \left(x - \rho^{i\frac{N}{a}} (t')^{\frac{N}{a}} \right).$$

This factorization demonstrates that after passing to the normalization $\widetilde{\mathcal{C}}'$, the central fiber becomes reduced over this locus.

Next, we consider the case where $p \in (\mathcal{C}_0)_{\mathrm{red}}$ is a node. Here, $\mathcal{C}'$ is defined étale locally by $x^a y^b = (t')^N$. Let $d := \gcd(a, b)$. If $d > 1$, we again have a factorization:

$$x^a y^b - (t')^N = \prod_{i=0}^{d-1} \left(x^{\frac{a}{d}} y^{\frac{b}{d}} - \rho^{i\frac{N}{d}} (t')^{\frac{N}{d}} \right).$$

Taking the normalization factors the space as $\widetilde{\mathcal{C}}' \to \mathcal{C}'' \to \mathcal{C}'$, where p has d preimages in $\mathcal{C}''$, each locally defined by $x^{a/d} y^{b/d} = \mathrm{unit} \cdot (t')^{N/d}$. The central fiber is still not necessarily reduced; however, the multiplicities of the two intersecting branches are now relatively prime. We can therefore assume without loss of generality that $\gcd(a, b) = 1$.

Choose positive integers α, β such that $\alpha a - \beta b = 1$. The normalization of the domain $R'[x,y]/(x^a y^b - (t')^N)$ is then explicitly given by the ring map:

$$R'[x,y]/(x^a y^b - (t')^N) \to R'[u,v]/(uv - (t')^{N/(ab)})$$
$$x \mapsto u^b$$
$$y \mapsto v^a$$

where $u = (t')^{\alpha N}/(x^\beta y^\alpha)$ and $v = x^\beta y^\alpha/(t')^{\beta N/a}$. After this final normalization, the central fiber is reduced, possessing a single node at the unique preimage of p. This local process introduces A_{m-1}-singularities of the form $uv = (t')^m$ into the total space. Repeatedly blowing up these isolated surface singularities replaces them with chains of smooth rational curves intersecting nodally.

Step V: Arrange the marked points $\sigma_i(0) \in \mathcal{C}_0$. We repeatedly blow up any closed points in the central fiber $\mathcal{C}_0$ where the marked points $\sigma_i(0)$ hit singular nodes or collide with one another. The strict transforms of the sections are thereby separated, eventually landing as distinct points in the smooth locus of the central fiber. After replacing $\mathcal{C}$ with this blow-up and the sections with their strict transforms, we arrive at a prestable family $(\mathcal{C} \to \Delta, \sigma_1, \ldots, \sigma_n)$.

Step VI: Contract rational tails and rational bridges. To upgrade the prestable family $(\mathcal{C} \to \Delta, \sigma_i)$ into a stable family, we rely on the characterization that a prestable curve is stable if and only if it contains no rational tails or rational bridges. By Proposition 2.4.5, contracting each rational tail and bridge yields a projective morphism $\mathcal{C} \to \mathcal{C}'$ of families of nodal curves over Δ. Defining the new sections as the composition $\sigma_i' \colon \Delta \to \mathcal{C} \to \mathcal{C}'$ results in a family $(\mathcal{C}' \to \Delta, \sigma_i')$ that is stable, thus concluding the proof of Stable Reduction. □

The stable reduction theorem established above implies the existence part of the valuative criterion for the properness of $\overline{\mathcal{M}}_{g,n}$. To complete the proof that the moduli stack $\overline{\mathcal{M}}_{g,n}$ is proper, we need show that the stable limit is in fact unique, which corresponds to the separatedness of the stack.

Proposition 2.4.18 (Uniqueness of the Stable Limit). *Let R be a DVR with fraction field K. Set $\Delta = \operatorname{Spec} R$ and $\Delta^* = \operatorname{Spec} K$. If $(\mathcal{C} \to \Delta, \sigma_1, \ldots, \sigma_n)$ and $(\mathcal{D} \to \Delta, \tau_1, \ldots, \tau_n)$ are families of n-pointed stable curves, then every isomorphism $\alpha^* \colon \mathcal{C} \times_\Delta \Delta^* \xrightarrow{\cong} \mathcal{D} \times_\Delta \Delta^*$ compatible with the sections extends to a unique isomorphism $\alpha \colon \mathcal{C} \to \mathcal{D}$ over Δ such that $\tau_i = \alpha \circ \sigma_i$ for all $1 \leq i \leq n$.*

Sketch of Proof. Assume $n = 0$; the uniqueness of marked sections follows from the separatedness of the curves once an isomorphism of the underlying surfaces is established. We may reduce to the case where the generic fibers $\mathcal{C}^* \cong \mathcal{D}^*$ are smooth over Δ^* via the identical argument used in Step I of Proposition 2.4.17.

Let $\widetilde{\mathcal{C}} \to \mathcal{C}$ and $\widetilde{\mathcal{D}} \to \mathcal{D}$ be the minimal regular resolutions, ensuring the reduced central fibers $(\widetilde{\mathcal{C}}_0)_{\mathrm{red}}$ and $(\widetilde{\mathcal{D}}_0)_{\mathrm{red}}$ are normal crossing divisors. Let Γ be the closure of the graph of the isomorphism $(\mathrm{id}, \alpha^*) \colon \mathcal{C}^* \xrightarrow{\cong} \widetilde{\mathcal{C}} \times_\Delta \widetilde{\mathcal{D}}$, and let $\widetilde{\Gamma} \to \Gamma$ be a resolution of singularities. This yields a diagram of birational morphisms:

$$\begin{tikzcd} & \widetilde{\mathcal{C}} \arrow[r] & \mathcal{C} \arrow[dr] & \\ \widetilde{\Gamma} \arrow[ur] \arrow[dr] & & & \Delta \\ & \widetilde{\mathcal{D}} \arrow[r] & \mathcal{D} \arrow[ur] & \end{tikzcd}$$

The birational maps $\widetilde{\Gamma} \to \widetilde{\mathcal{C}}$ and $\widetilde{\Gamma} \to \widetilde{\mathcal{D}}$ between regular proper surfaces are isomorphisms in codimension 2 (isomorphisms outside finite sets of points). By [**Har77**, Theorem II.8.19], their relative pluricanonical sheaves share isomorphic global sections:

$$\Gamma\left(\widetilde{\mathcal{C}}, \omega_{\widetilde{\mathcal{C}}/\Delta}^{\otimes k}\right) \cong \Gamma\left(\widetilde{\mathcal{D}}, \omega_{\widetilde{\mathcal{D}}/\Delta}^{\otimes k}\right), \quad \forall k \in \mathbb{N}.$$

The morphism $\pi\colon \widetilde{\mathcal{C}} \to \mathcal{C}$ resolves the nodal singularities of $\mathcal{C}$. Because $\mathcal{C}$ is a stable family, π contracts chains of smooth rational -2-curves in the central fiber. Thus, π is exactly the stable contraction of the regular prestable family $\widetilde{\mathcal{C}}$.

Since nodal singularities are rational, the pushforward of the relative pluricanonical sheaf under the resolution matches that of the stable curve: $\pi_*\left(\omega_{\widetilde{\mathcal{C}}/\Delta}^{\otimes k}\right) \cong \omega_{\mathcal{C}/\Delta}^{\otimes k}$. By the properties of stable contractions ([**ACG11**, Chapter X, Proposition 6.7]), $\mathcal{C}$ is isomorphic to the relative canonical model of its resolution:

$$\mathcal{C} \cong \operatorname{Proj}_\Delta \bigoplus_{k\geq 0} \Gamma\left(\widetilde{\mathcal{C}}, \omega_{\widetilde{\mathcal{C}}/\Delta}^{\otimes k}\right)$$

By identical algebraic reasoning for $\mathcal{D}$, the canonical isomorphism of their relative pluricanonical rings over Δ forces $\mathcal{C} \cong \mathcal{D}$. $\square$

Before delving into the discussion of irreducibility, we introduce several fundamental operations on the moduli stack $\overline{\mathcal{M}}_{g,n}$. These include the contraction, forgetful, and gluing morphisms. We will focus specifically on proving that the forgetful morphism $\overline{\mathcal{M}}_{g,n+1} \to \overline{\mathcal{M}}_{g,n}$ serves as the universal family over $\overline{\mathcal{M}}_{g,n}$.

We have previously introduced the stable contraction of a single prestable curve (Proposition 2.4.5). It turns out that this pointwise construction extends to flat families of prestable curves (see [**Alp26**, Theorem 6.6.6]).

Proposition 2.4.19 (Stable Contraction of a Prestable Family)**.** *Let $(\mathcal{C} \to S, \sigma_i)$ be a family of n-pointed prestable curves of genus g satisfying $2g-2+n>0$. Then there exists a unique projective morphism $c\colon \mathcal{C} \to \mathcal{C}^{\mathrm{st}}$ over S such that:*

(1) The family $(\mathcal{C}^{\mathrm{st}} \to S, c\circ\sigma_i)$ is an n-pointed family of strictly stable curves of genus g.
(2) $\mathcal{O}_{\mathcal{C}^{\mathrm{st}}} = c_*\mathcal{O}_{\mathcal{C}}$ and $\mathrm{R}^1 c_*\mathcal{O}_{\mathcal{C}} = 0$.
(3) The construction of c is compatible with any base change $S' \to S$.
(4) For each $s \in S$, the restriction of c to the geometric fiber $\mathcal{C}_s$ is the stable contraction constructed in Proposition 2.4.5.

As a corollary, this induces a canonical morphism of algebraic stacks $\mathcal{M}_{g,n}^{\mathrm{pres}} \to \overline{\mathcal{M}}_{g,n}$ defined on points by $(C, p_i) \mapsto (C^{\mathrm{st}}, c(p_i))$.

The stable contraction of a prestable family allows us to define the forgetful morphism. Naively, one might attempt to simply delete the final marked point of a stable curve $(C, p_1, \ldots, p_{n+1})$ to define a map $\overline{\mathcal{M}}_{g,n+1} \to \overline{\mathcal{M}}_{g,n}$. However, deleting a marked point can cause a previously stable rational component to drop to only two special points, destroying the stability of the curve. Therefore, we must apply the stable contraction morphism to re-stabilize the curve after the point is forgotten. More precisely, the forgetful morphism is defined on objects by $(C, p_1, \ldots, p_{n+1}) \mapsto (C^{\mathrm{st}}, c(p_1), \ldots, c(p_n))$.

Proposition 2.4.20 (Universal Family)**.** *The forgetful morphism $\overline{\mathcal{M}}_{g,n+1} \to \overline{\mathcal{M}}_{g,n}$ is isomorphic to the universal family over $\overline{\mathcal{M}}_{g,n}$.*

Proof. Let $\mathcal{U}_{g,n}$ denote the universal family over the stack $\overline{\mathcal{M}}_{g,n}$. This corresponds to a universal n-pointed family of curves $(\mathcal{U}_{g,n} \to \overline{\mathcal{M}}_{g,n}, \sigma_1^{\mathrm{univ}}, \ldots, \sigma_n^{\mathrm{univ}})$. By the definition of a universal family, an object of $\mathcal{U}_{g,n}$ over a scheme S is equivalent to the data of an n-pointed family of stable curves $(\mathcal{C} \to S, \sigma_1, \ldots, \sigma_n)$ classified by a unique morphism $S \to \overline{\mathcal{M}}_{g,n}$, together with a section that fits into the following Cartesian diagram:

$$\begin{tikzcd}
\mathcal{C} \arrow[r] \arrow[d] & \mathcal{U}_{g,n} \arrow[d] \\
S \arrow[r] \arrow[u, bend left, "\sigma_i"] \arrow[ur, phantom, "\square\sigma_i^{\mathrm{univ}}"] & \overline{\mathcal{M}}_{g,n} \arrow[u, bend left]
\end{tikzcd}$$

Thus, an object of $\mathcal{U}_{g,n}$ over S is equivalent to an n-pointed stable family $(\mathcal{C} \to S, \sigma_1, \ldots, \sigma_n)$ equipped with one additional arbitrary section $\tau \colon S \to \mathcal{C}$. This data naturally defines a morphism of algebraic stacks $\Phi \colon \overline{\mathcal{M}}_{g,n+1} \to \mathcal{U}_{g,n}$ over $\overline{\mathcal{M}}_{g,n}$. The morphism Φ takes an $(n+1)$-pointed stable family $(\mathcal{D} \to S, \sigma_1, \ldots, \sigma_{n+1})$ to the data $((\mathcal{D}^{\mathrm{st}} \to S, \sigma_1^{\mathrm{st}}, \ldots, \sigma_n^{\mathrm{st}}), \sigma_{n+1}^{\mathrm{st}})$, where the first n sections are stabilized after forgetting σ_{n+1}, and the final stabilized section $\sigma_{n+1}^{\mathrm{st}}$ acts as the additional section τ.

It suffices to show that $\Phi \colon \overline{\mathcal{M}}_{g,n+1} \to \mathcal{U}_{g,n}$ is an isomorphism. By Proposition 2.4.17, both $\overline{\mathcal{M}}_{g,n+1}$ and $\mathcal{U}_{g,n}$ are proper over $\overline{\mathcal{M}}_{g,n}$, so the morphism Φ is proper and representable. Clearly Φ is quasi-finite (Definition 2.3.21), so it is finite. Furthermore, Φ is a birational morphism. It is an isomorphism between the open substack of $\overline{\mathcal{M}}_{g,n+1}$ parametrizing $(C, p_1, \ldots, p_{n+1})$ where dropping p_{n+1} does not create any rational bridges or tails, and the open substack of $\mathcal{U}_{g,n}$ parametrizing $(C, p_1, \ldots, p_n, \tau)$ where the extra point τ lies in the smooth locus of C and does not collide with the first n marked points.

Finally, we analyze the singularities of the target $\mathcal{U}_{g,n}$. Because $\overline{\mathcal{M}}_{g,n}$ is a smooth stack and the fibers of $\mathcal{U}_{g,n} \to \overline{\mathcal{M}}_{g,n}$ are nodal curves, the total space $\mathcal{U}_{g,n}$ is Cohen–Macaulay. Moreover, $\mathcal{U}_{g,n}$ is regular on the open locus where the additional section τ (corresponding to p_{n+1}) lands in the smooth locus of the curve C. The complement of this regular locus consists of points where the extra section hits a node; this complement has codimension 2 in the total space. By Serre's criterion for normality, $\mathcal{U}_{g,n}$ is normal. Because $\Phi \colon \overline{\mathcal{M}}_{g,n+1} \to \mathcal{U}_{g,n}$ is a finite birational morphism to a normal Deligne–Mumford stack, Zariski's Main Theorem guarantees that Φ is an isomorphism. $\square$

Finally, we introduce the gluing morphism along marked sections. This is a very natural geometric operation and serves as an essential tool in defining the tautological ring of the moduli space. More precisely, consider two field-valued points $[(C, p_1, \ldots, p_{n_1})] \in \overline{\mathcal{M}}_{g_1,n_1}(k)$ and $[(D, q_1, \ldots, q_{n_2})] \in \overline{\mathcal{M}}_{g_2,n_2}(k)$. Let $g = g_1 + g_2$ and $n = n_1 + n_2 - 2$. We can glue these two stable curves together by identifying the points p_{n_1} and q_{n_2} to form a single curve with a new node:

$$(C \sqcup_{p_{n_1} \sim q_{n_2}} D, p_1, \ldots, p_{n_1-1}, q_1, \ldots, q_{n_2-1})$$

To verify that this glued curve defines a moduli point in $\overline{\mathcal{M}}_{g,n}(k)$, it suffices to check that it preserves stability. This is immediate: by construction, the normalization of the glued curve is exactly the disjoint union $C \sqcup D$, and the connected components of this normalization are stable by our initial assumption.

This defines the gluing morphism for single stable curves over a field, and it also extends to flat families of curves. To construct this extension, it is sufficient to consider a general family of stable curves $(\mathcal{C} \to S, \sigma, \tau)$ over a scheme S equipped with two disjoint sections σ and τ. Because the family $\mathcal{C} \to S$ is projective, the images $\sigma(S)$ and $\tau(S)$ are closed subschemes that locally lie within a common affine open subscheme of $\mathcal{C}$. This topological separation allows us to glue the sections by taking the Ferrand pushout in the category of schemes:

$$\begin{tikzcd} S \sqcup S \arrow[r, "\sigma\sqcup\tau"] \arrow[d] & \mathcal{C} \arrow[d, "g"] \\ S \arrow[r, "\mu"] & \mathcal{C}' \end{tikzcd}$$

The resulting proper, flat family $(\mathcal{C}' \to S)$ gives the desired gluing, where the image of the section $\mu(s)$ specifies a new node in the geometric fiber $\mathcal{C}'_s$ for every $s \in S$.

Therefore we obtain the gluing morphism $\overline{\mathcal{M}}_{g_1,n_1} \times \overline{\mathcal{M}}_{g_2,n_2} \to \overline{\mathcal{M}}_{g,n}$. Because this operation only introduces a node, the gluing morphism is quasi-finite and representable. Since it is a quasi-finite morphism between proper Deligne–Mumford stacks, it is a finite morphism. Similarly there is a finite morphism $\overline{\mathcal{M}}_{g-1,n+2} \to \overline{\mathcal{M}}_{g,n}$ by gluing the final two sections.

We conclude this subsection with a brief discussion of the irreducibility of $\overline{\mathcal{M}}_{g,n}$, which completes our survey of Theorem 2.4.7. Because the stack $\overline{\mathcal{M}}_{g,n}$ is smooth, proving irreducibility is equivalent to proving connectedness. By Proposition 2.4.20, the forgetful morphism $\overline{\mathcal{M}}_{g,n+1} \to \overline{\mathcal{M}}_{g,n}$ is the universal family. Since stable curves are connected by definition, the geometric fibers of this morphism are connected, so by induction, it suffices to prove connectedness for the unpointed case $n = 0$. Finally, showing $\overline{\mathcal{M}}_g$ is connected reduces to proving two things: that the open substack $\mathcal{M}_g$ is connected, and that $\mathcal{M}_g$ is dense in $\overline{\mathcal{M}}_g$.

Proposition 2.4.21 (Density of Smooth Curves)**.** *The open substack $\mathcal{M}_g$ is dense in $\overline{\mathcal{M}}_g$. This means every stable curve of genus g deforms a smooth curve: for any stable curve C, there exists a family of stable curves $\pi\colon \mathcal{C} \to T$ over a curve T equipped with a closed point $t \in T$, such that the central fiber is $\mathcal{C}_t \cong C$ and the restriction of π over $T \setminus \{t\}$ is a family of smooth curves.*

Proof. Étale locally, a node over $\mathbb{C}$ is isomorphic to the surface singularity $xy = 0$, which is smoothed by the one-parameter local family $xy = s$. More precisely, there is a non-trivial local first-order deformation over the dual numbers

$$\operatorname{Spec} \mathbb{C}[x, y, \varepsilon]/(xy - \varepsilon, \varepsilon^2) \to \operatorname{Spec} \mathbb{C}[\varepsilon]/(\varepsilon^2).$$

By Proposition 2.4.13, the global first-order deformations of a stable curve C are parameterized by the group $\operatorname{Ext}^1(\Omega_C, \mathcal{O}_C)$. The relationship between global deformations and local deformations at the singularities is governed by the local-to-global spectral sequence for Ext groups, which yields the exact sequence:

$$\cdots \to \operatorname{Ext}^1(\Omega_C, \mathcal{O}_C) \xrightarrow{\phi} \mathrm{H}^0(C, \mathcal{E}xt^1(\Omega_C, \mathcal{O}_C)) \to \mathrm{H}^2(C, \mathcal{H}om(\Omega_C, \mathcal{O}_C)) \to \cdots.$$

The map ϕ restricts a global deformation to the local deformations at the nodes. Since C is a curve, the obstruction space vanishes: $\mathrm{H}^2(C, \mathcal{H}om(\Omega_C, \mathcal{O}_C)) = 0$. This vanishing guarantees that

ϕ is surjective, so any combination of local deformations of the nodes can be lifted to a global first-order deformation of the curve C.

Let F be the deformation functor of C. We have constructed an element in $F(\mathbb{C}[t]/(t^2))$ corresponding to this global first-order smoothing. To lift this deformation continuously to higher-order Artin rings $\mathbb{C}[t]/(t^{n+1})$, the obstructions lie in the group $\operatorname{Ext}^2(\Omega_C, \mathcal{O}_C)$. Since C is a curve, $\operatorname{Ext}^2(\Omega_C, \mathcal{O}_C) = 0$. Since all higher-order obstructions vanish, we can lift our first-order deformation indefinitely, defining a sequence of deformations that yields an element in the inverse limit $\widehat{F}(\mathbb{C}[[t]])$. By Proposition 2.1.4, this limit element corresponds to a natural transformation $h_{\mathbb{C}[[t]]} \to F$, which defines a formal family $\widehat{\mathcal{C}}$ over $\operatorname{Spec} \mathbb{C}[[t]]$.

Finally, because C is a projective curve, Grothendieck's Existence Theorem [**Ser06**, Theorem 2.5.13] guarantees that this formal deformation is effective. Applying Artin's Algebraization theorem [**Ser06**, Theorem 2.5.14] to this effective versal formal deformation yields an algebraic curve T over $\mathbb{C}$, equipped with a closed point $t \in T$ whose completed local ring is $\mathbb{C}[[t]]$, along with a family $\pi\colon \mathcal{C} \to T$ that recovers our deformation at t. This provides the geometric family where the central fiber is C and the generic fibers are smooth, concluding the proof that $\mathcal{M}_g$ is dense in $\overline{\mathcal{M}}_g$. □

After completing these technical proofs, we now turn to the connectedness of the moduli stack $\mathcal{M}_g$. Surprisingly, we can reduce this topological question to a purely combinatorial problem using the elementary Hurwitz theory established in the first chapter.

Let $\operatorname{Hur}_{g,b}$ denote the moduli space parameterizing simply branched coverings $C \to \mathbb{P}^1$ over b ordered branch points, where the domain C is a connected, smooth curve of genus g. We define a family of simply branched coverings of $\mathbb{P}^1$ of genus g over b ordered points, parameterized by a scheme S, as a commutative diagram equipped with b sections:

$$\begin{tikzcd} \mathcal{C} \arrow[rr] \arrow[dr] & & \mathcal{D} \arrow[dl] \\ & S \arrow[ur, bend right, "\sigma_i"'] & \end{tikzcd}$$

In this diagram, the map $\mathcal{C} \to S$ is a family of smooth curves of genus g (an S-point of $\mathcal{M}_g$), and the map $\mathcal{D} \to S$ with sections $\sigma_1, \ldots, \sigma_b$ is a family of smooth rational curves with b marked points (an S-point of $\mathcal{M}_{0,b}$). Furthermore, over every geometric point $s \in S(k)$, the morphism of fibers $\mathcal{C}_s \to \mathcal{D}_s$ is a covering simply branched over the points $\sigma_1(s), \ldots, \sigma_b(s)$.

This allows us to define the Hurwitz functor as:

$$\begin{aligned} \operatorname{Hur}_{g,b}\colon \mathsf{Sch}^{\mathrm{op}} &\longrightarrow \mathsf{Sets} \\ S &\longmapsto \Big\{\text{families of simply branched coverings of } \mathbb{P}^1 \text{ of genus } g \text{ over } b \text{ points}\Big\} / \sim \end{aligned}$$

Fix g and b as above, and let d be the degree of the simple covering $C \to \mathbb{P}^1$. By the Riemann–Hurwitz formula, the number of branch points is $b = 2g + 2d - 2$. Assume $d \geq 3$, which forces $b \geq 2g + 4 \geq 4$. Since $b \geq 3$, the rational curves parameterized by $\mathcal{M}_{0,b}$ have no non-trivial automorphisms. The vanishing of automorphisms guarantees that the moduli stack is representable by an algebraic space, but we can explicitly realize it as a fine moduli space. As established in Example 2.3.10, applying the unique automorphism of $\mathbb{P}^1$ fixing $0, 1$, and ∞ reduces the remaining $b-3$ points to the configuration space:

$$M_{0,b} \cong \Big(\mathbb{P}^1 \setminus \{0, 1, \infty\}\Big)^{b-3} \setminus \Delta.$$

Therefore, $M_{0,b}$ is an open subscheme of an affine space. Forgetting the domain of the cover induces a natural morphism $\mathrm{Hur}_{g,b} \to M_{0,b}$, defined on S-points by sending the covering family $\mathcal{C} \to \mathcal{D}$ to the base family of rational curves $(\mathcal{D} \to S, \sigma_1, \dots, \sigma_b) \in M_{0,b}(S)$. This morphism satisfies the following properties:

- The morphism $\mathrm{Hur}_{g,b} \to M_{0,b}$ is finite, étale, and representable.
- Since the target $M_{0,b}$ is a scheme and the morphism is representable, the functor $\mathrm{Hur}_{g,b}$ is represented by a scheme, known as the Hurwitz scheme (or Hurwitz space).
- If $g \geq 2$ and $d \geq 2g+3$ (with $b = 2g+2d-2$), the natural projection morphism forgetting the map,
$$\mathrm{Hur}_{g,b} \to \mathcal{M}_g, \qquad (C \to \mathbb{P}^1) \mapsto C$$
is surjective.

For details regarding these properties, see [**Ful69**]. We briefly sketch the proof of the final property, as it is vital for establishing the connectedness of $\mathcal{M}_g$. To prove the forgetful morphism is surjective, we need to find an ample line bundle L on C and a 2-dimensional linear system $V \subseteq \mathrm{H}^0(C, L)$ that induces a simply branched covering $C \to \mathbb{P}^1$. The degree bound $d \geq 2g+3$ ensures $\deg(\omega_C \otimes L^\vee) < 0$. By Serre duality, $\mathrm{h}^1(C, L) = \mathrm{h}^0(C, \omega_C \otimes L^\vee) = 0$. Riemann–Roch then yields $\mathrm{h}^0(C, L) = d + 1 - g$. Therefore, the Grassmannian $\mathrm{Gr}(2, \mathrm{H}^0(C, L))$ of 2-dimensional linear systems has dimension $2(d-1-g)$.

It suffices to show that the locus of valid linear systems is non-empty. A linear subspace $V \in \mathrm{Gr}(2, \mathrm{H}^0(C, L))$ fails to induce a simply branched covering $C \to \mathbb{P}^1$ if and only if one of the following geometric degenerations occurs:

(1) V has a base point.
(2) There exists a ramification point with index ≥ 3.
(3) There exist two distinct ramification points in the same fiber.

Each of these three conditions defines a closed locus of codimension at least one in the Grassmannian. Consequently, the generic element of $\mathrm{Gr}(2, \mathrm{H}^0(C, L))$ avoids all three obstructions, guaranteeing the existence of the required linear system V (see [**Ful69**, Proposition 8.1]).

Proposition 2.4.22 (Connectedness of $\mathcal{M}_g$)**.** *For $g \geq 2$ and $d \geq 2$ with $b = 2g+2d-2$, the Hurwitz scheme $\mathrm{Hur}_{g,b}$ is connected.*

Since the projection $\mathrm{Hur}_{g,b} \to \mathcal{M}_g$ is surjective when b is sufficiently large for any $g \geq 2$, it follows as a corollary that $\mathcal{M}_g$ is connected.

Sketch of Proof. By Lemma 1.1.1, we pass to the topological category to analyze the finite étale morphism $\beta\colon \mathrm{Hur}_{g,b} \to M_{0,b}$. This morphism defines an unramified covering space over the complex manifold $M_{0,b}^{\mathrm{an}}$. Fixing a base configuration of branch points $B = \{q_1, \dots, q_b\} \in M_{0,b}^{\mathrm{an}}$, the fundamental group $\pi_1(M_{0,b}^{\mathrm{an}}, B)$ acts on the discrete fiber $\mathrm{Hur}_{g,B} \coloneqq \beta^{-1}(B)$ via path lifting. The space $\mathrm{Hur}_{g,b}$ is connected if and only if this monodromy action is transitive.

By Theorem 1.2.4, the fiber $\mathrm{Hur}_{g,B}$ is in natural bijection with the set of branch data:

$$\begin{aligned}\mathrm{Hur}_{g,B} &= \{\text{simply branched coverings } C \to \mathbb{P}^1 \text{ over } B\}/\sim \\ &\cong \left\{ (\tau_1, \dots, \tau_b) \in (S_d)^b \;\middle|\; \begin{array}{c} \tau_i \text{ are transpositions}, \prod_{i=1}^b \tau_i = 1, \\ \text{and } \langle \tau_1, \dots, \tau_b \rangle \text{ acts transitively on } \{1, \dots, d\} \end{array} \right\} \Big/ \sim\end{aligned}$$

where $\sim$ denotes simultaneous conjugation by elements of S_d.

The fundamental group $\pi_1(M^{\mathrm{an}}_{0,b}, B)$ is generated by the elements $\sigma_1, \dots, \sigma_{b-1}$, where each generator σ_i corresponds to a loop that swaps q_i and q_{i+1} along a counterclockwise semicircle. The induced monodromy action of σ_i on a factorization sequence is given by the Hurwitz move:

$$\sigma_i \cdot (\tau_1, \dots, \tau_i, \tau_{i+1}, \dots, \tau_b) = (\tau_1, \dots, \tau_{i-1}, \tau_i\tau_{i+1}\tau_i^{-1}, \tau_i, \tau_{i+2}, \dots, \tau_b)$$

To establish transitivity, it suffices to show that any valid sequence can be reduced to the following target element:

$$\tau = \left(\underbrace{(1\,2), (1\,2)}_{2}, \underbrace{(1\,3), (1\,3)}_{2}, \dots, \underbrace{(1\,d-1), (1\,d-1)}_{2}, \underbrace{(1\,d), \dots, (1\,d)}_{2g+2} \right) \in \mathrm{Hur}_{g,B}$$

As shown in [**Ful69**], there always exists a finite sequence of Hurwitz moves $\sigma_{i_1} \cdots \sigma_{i_k}$ transforming an arbitrary sequence $(\tau_1, \dots, \tau_b)$ into τ. The group action is therefore transitive, showing that the Hurwitz space $\mathrm{Hur}_{g,b}$ is connected. □

This concludes the proof of Theorem 2.4.7.

2.4.3. Tautological Classes and the Tautological Ring. We conclude this section with a discussion of the tautological classes on the moduli stack $\overline{\mathcal{M}}_{g,n}$. These were introduced by Mumford in [**Mum83**], motivated by the cohomology of Grassmannians. Complex vector bundles on a topological space are classified by continuous maps to Grassmannians, and the cohomology of a Grassmannian is generated by the Chern classes of its universal subbundles. By analogy, to understand the intersection theory of the moduli space of curves, one seeks a natural set of algebraic classes that form a subring of the Chow ring. This motivates the introduction of the tautological ring. Unless specifically stated, all Chow rings are taken with rational coefficients. The formal definitions here follow Faber and Pandharipande [**FP13**].

Definition 2.4.23 (Tautological Ring)**.** Let $\pi\colon \overline{\mathcal{M}}_{g,n+1} \to \overline{\mathcal{M}}_{g,n}$ be the forgetful morphism corresponding to the universal family (Proposition 2.4.20). Consider the gluing morphisms:

$$\iota_1\colon \overline{\mathcal{M}}_{g_1,n_1} \times \overline{\mathcal{M}}_{g_2,n_2} \to \overline{\mathcal{M}}_{g_1+g_2,n_1+n_2-2}$$
$$\iota_2\colon \overline{\mathcal{M}}_{g-1,n+2} \to \overline{\mathcal{M}}_{g,n}$$

The system of tautological rings is defined as the smallest system of $\mathbb{Q}$-subalgebras

$$\mathrm{R}^*(\overline{\mathcal{M}}_{g,n}) \subset \mathrm{CH}^*(\overline{\mathcal{M}}_{g,n})$$

simultaneously satisfying the following two properties:

(1) The system is closed under pushforward via the forgetful morphisms:

$$\pi_*\colon \mathrm{R}^*(\overline{\mathcal{M}}_{g,n+1}) \to \mathrm{R}^*(\overline{\mathcal{M}}_{g,n})$$

(2) The system is closed under pushforward via all gluing morphisms:

$$\iota_{1*}\colon \mathrm{R}^*(\overline{\mathcal{M}}_{g_1,n_1}) \otimes_{\mathbb{Q}} \mathrm{R}^*(\overline{\mathcal{M}}_{g_2,n_2}) \to \mathrm{R}^*(\overline{\mathcal{M}}_{g_1+g_2,n_1+n_2-2})$$
$$\iota_{2*}\colon \mathrm{R}^*(\overline{\mathcal{M}}_{g-1,n+2}) \to \mathrm{R}^*(\overline{\mathcal{M}}_{g,n})$$

While this recursive definition is elegant, it is more geometrically illuminating to construct explicit algebraic classes that reside within this ring.

Let $\omega_\pi := \omega_{\overline{\mathcal{M}}_{g,n+1}/\overline{\mathcal{M}}_{g,n}}$ be the relative dualizing sheaf of the universal family. We define the Hodge bundle as the pushforward $\mathbb{E} := \pi_*\omega_\pi$. By Proposition 2.4.10, this pushforward is a vector bundle of rank g (Example 2.3.41). Let $\sigma_i \colon \overline{\mathcal{M}}_{g,n} \to \overline{\mathcal{M}}_{g,n+1}$ denote the i-th section of the universal family for $i = 1, \ldots, n$. We define the i-th cotangent line bundle as the pullback:

$$\mathbb{L}_i \coloneqq \sigma_i^* \omega_\pi$$

The fiber of $\mathbb{L}_i$ over a moduli point $[C, p_1, \ldots, p_n]$ is the cotangent space $T^*_{p_i}C$ of the curve at the i-th marked point.

Definition 2.4.24 (ψ-classes, κ-classes, and λ-classes)**.**

(1) The i-th ψ-class, denoted ψ_i, is the first Chern class of the i-th cotangent line bundle:

$$\psi_i \coloneqq c_1(\mathbb{L}_i) \in \mathrm{CH}^1(\overline{\mathcal{M}}_{g,n})$$

(2) The i-th κ-class, denoted κ_i, is defined via the pushforward of higher powers of the $(n+1)$-th ψ-class on the universal family:

$$\kappa_i \coloneqq \pi_* \left(\psi_{n+1}^{i+1} \right) \in \mathrm{CH}^i(\overline{\mathcal{M}}_{g,n})$$

(3) The i-th λ-class, denoted λ_i, is the i-th Chern class of the Hodge bundle:

$$\lambda_i \coloneqq c_i(\mathbb{E}) \in \mathrm{CH}^i(\overline{\mathcal{M}}_{g,n})$$

The ψ-classes belong to the tautological ring because they can be expressed as the pushforward of the self-intersection of boundary strata. Specifically:

$$-\pi_* \left(\left(\iota_{1*}([\overline{\mathcal{M}}_{g,n}] \times [\overline{\mathcal{M}}_{0,3}]) \right)^2 \right) = \psi_i.$$

Since the tautological ring is closed under pushforwards by π, the κ-classes are tautological by definition. The λ-classes are also tautological; this non-trivial fact is established by Mumford's formula [**Mum83**, (5.2)], which expresses the Chern character of the Hodge bundle $\mathbb{E}$ entirely in terms of κ-classes and boundary divisors.

The definition of the tautological ring naturally extends to open substacks and to singular cohomology. Given an open substack $U \subseteq \overline{\mathcal{M}}_{g,n}$, the tautological ring $\mathrm{R}^*(U)$ is defined as the image of the tautological ring of $\overline{\mathcal{M}}_{g,n}$ under the natural restriction map $\mathrm{CH}^*(\overline{\mathcal{M}}_{g,n}) \to \mathrm{CH}^*(U)$. Furthermore, to translate these intersection-theoretic classes into topological invariants, we utilize the cycle class map [**Ful98**, Chapter 19]:

$$\mathrm{cl} \colon \mathrm{CH}^*(\overline{\mathcal{M}}_{g,n}) \to \mathrm{H}^*(\overline{\mathcal{M}}_{g,n}; \mathbb{Q}).$$

The tautological ring in cohomology, denoted $\mathrm{RH}^*(\overline{\mathcal{M}}_{g,n})$, is defined as the image of the tautological Chow ring $\mathrm{R}^*(\overline{\mathcal{M}}_{g,n})$ under this cycle map.

The study of the tautological ring occupies a central position in Gromov–Witten theory, modern enumerative geometry, and the broader landscape of moduli problems in algebraic geometry. It remains a remarkably active and rapidly developing field of current research. Consequently, providing a comprehensive survey of this vast and beautiful subject lies well beyond the scope of this text. For a foundational introduction to the tautological rings $\mathrm{R}^*(\mathcal{M}_g)$ and $\mathrm{R}^*(\overline{\mathcal{M}}_{g,n})$, we refer the reader to [**Pan18**] and [**FP13**]. Additionally, for an excellent overview of recent breakthroughs concerning the Chow and cohomology rings of the moduli stack $\overline{\mathcal{M}}_{g,n}$, we direct the reader to the detailed survey in [**CL24**, Section 1].

CHAPTER 3

Equivariant Cohomology in Algebraic Geometry

We motivate this chapter with a brief discussion of enumerative geometry. As one of the oldest and most classical subfields of algebraic geometry, it focuses on counting geometric figures that satisfy certain given conditions. One of its most well-known and classical problems is the following:

Question 3.0.1. In the complex projective plane $\mathbb{P}^2$, given $3d-1$ general points, how many rational curves of degree d pass through these points?

Here, $3d-1$ is the "codimension" required for the intersection class to be a zero-cycle (a concept we will define rigorously later). This condition ensures that the number N_d is finite. The degree-one case, $N_1 = 1$, was well-known to ancient Greek mathematicians, as it corresponds to an axiom of Euclidean geometry (two points determine a unique line). Similarly, the degree-two case yields $N_2 = 1$ (five points determine a unique conic). However, as the degree grows, the problem becomes significantly more difficult. The answer $N_3 = 12$ was not obtained until 1848 by Steiner, and $N_4 = 620$ was computed by Zeuthen in the late 19th century. We refer the reader to [**Kle87**] for a historical review of enumerative geometry.

Thanks to the development of intersection theory by Fulton, MacPherson, and others, enumerative geometry found a solid foundation, answering Hilbert's 15th problem. In 1995, Kontsevich's revolutionary paper [**Kon95**] transformed Question 3.0.1 into a grander program framed in the language of the intersection theory on the moduli stack of stable maps. Following [**FP97**], we provide a brief survey of how this paradigm shift occurred. Unless otherwise stated, we work exclusively over the field of complex numbers $\mathbb{C}$ throughout this chapter, and all Chow groups and rings are taken with rational coefficients $\mathbb{Q}$.

Let $(C, p_1, \ldots, p_n)$ be an n-pointed quasi-stable curve (Definition 2.4.1), and let X be a non-singular projective convex variety. Consider a morphism $\mu\colon C \to X$. We call μ *stable* if, for every irreducible subcurve $E \subset C$ contracted to a point under μ, the following conditions hold:

(1) If $E \cong \mathbb{P}^1 \subset C$, then E contains at least 3 special points (Definition 2.4.1);
(2) If $E \subset C$ is of arithmetic genus 1, then E must contain at least 1 special point.

A family of stable maps of genus g with n marked points over a scheme S is given by a commutative diagram

$$\begin{tikzcd} \mathcal{C} \arrow[r, "\mu"] \arrow[d, "\pi"] & X \\ S & \end{tikzcd}$$

where $\pi\colon \mathcal{C} \to S$ is a family of n-pointed quasi-stable curves of genus g, such that for each $s \in S$, the restricted morphism $\mu|_{\mathcal{C}_s}\colon \mathcal{C}_s \to X$ is a stable map.

We now introduce the definition of the moduli stack of stable maps $\overline{\mathcal{M}}_{g,n}(X, \beta)$. Let $\beta \in \mathrm{CH}_1(X)$ be a 1-cycle class in the Chow group. The category $\overline{\mathcal{M}}_{g,n}(X, \beta)$ is the prestack parameterizing stable maps $\mu\colon (C, p_1, \ldots, p_n) \to X$ such that $\mu_*([C]) = \beta$. More precisely, the objects $\overline{\mathcal{M}}_{g,n}(X, \beta)(T)$

over a scheme T consist of families of stable maps of genus g with n marked points over T. On each fiber, the fundamental class of the domain curve is pushed forward to β. The morphisms are defined via Cartesian squares compatible with the target morphism to X. Specifically, if $T \to S$ is a morphism in $\mathsf{Sch}_{\text{ét}}$, a morphism

$$\left(\begin{tikzcd} \mathcal{C} \arrow[r, "\mu_1"] \arrow[d, "\pi"] & X \\ S & \end{tikzcd} \longrightarrow \begin{tikzcd} \mathcal{D} \arrow[r, "\mu_2"] \arrow[d, "\pi"] & X \\ T & \end{tikzcd}\right)$$

is defined as a morphism $\mathcal{D} \to \mathcal{C}$ such that the diagram

$$\begin{tikzcd} \mathcal{D} \arrow[r] \arrow[d] \arrow[rr, bend left, "\mu_2"] & \mathcal{C} \arrow[r, "\mu_1"] \arrow[d, "\pi"] & X \\ T \arrow[r] & S & \end{tikzcd}$$

commutes, with the left square being Cartesian. We collect the fundamental properties proved in [**FP97**, Theorems 1–3] for the special case where $g = 0$.

Theorem 3.0.2 (Properties of $\overline{\mathcal{M}}_{0,n}(X, \beta)$)**.**

(1) The prestack $\overline{\mathcal{M}}_{0,n}(X, \beta)$ is a smooth and proper Deligne–Mumford stack of pure dimension

$$\dim X + \int_\beta c_1(X) + n - 3,$$

admitting a projective coarse moduli space $\overline{\mathcal{M}}_{0,n}(X, \beta) \to \overline{M}_{0,n}(X, \beta)$. Furthermore, if X is a projective space, then $\overline{\mathcal{M}}_{0,n}(X, \beta)$ is irreducible and of finite type.

(2) The boundary of $\overline{\mathcal{M}}_{0,n}(X, \beta)$ is a divisor with normal crossings.

Returning to Question 3.0.1, we focus on the moduli stack $\overline{\mathcal{M}}_{0,n}(\mathbb{P}^2, d[H])$, which we abbreviate as $\overline{\mathcal{M}}_{0,n}(\mathbb{P}^2, d)$. We define flat evaluation morphisms:

$$\mathrm{ev}_i \colon \overline{\mathcal{M}}_{0,n}(\mathbb{P}^2, d) \to \mathbb{P}^2; \qquad [(C, p_1, \ldots, p_n, \mu)] \mapsto \mu(p_i), \quad i = 1, \ldots, n,$$

which naturally extend to families of stable maps. Given cohomology classes $\gamma_1, \ldots, \gamma_n \in \mathrm{CH}^*(\mathbb{P}^2)$, one defines the *Gromov–Witten invariants* as

$$\langle \gamma_1, \ldots, \gamma_n \rangle_{0,d}^{\mathbb{P}^2} \coloneqq \int_{[\overline{\mathcal{M}}_{0,n}(\mathbb{P}^2,d)]} \mathrm{ev}_1^* \gamma_1 \cup \cdots \cup \mathrm{ev}_n^* \gamma_n,$$

where $\sum_i \operatorname{codim} \gamma_i = \dim \overline{\mathcal{M}}_{0,n}(\mathbb{P}^2, d) = 3d + n - 1$. In particular, if we set $\gamma_1 = \cdots = \gamma_n = [*] \in \mathrm{CH}^*(\mathbb{P}^2)$ (the class of a point) and let $n = 3d - 1$, then $\sum_i \operatorname{codim} \gamma_i = 2(3d - 1) = \dim \overline{\mathcal{M}}_{0,n}(\mathbb{P}^2, d)$. An application of Kleiman's transversality theorem [**Har77**, Theorem 10.8] guarantees that these intersection classes meet transversely, yielding

$$N_d = \langle \gamma_1, \ldots, \gamma_{3d-1} \rangle_{0,d}^{\mathbb{P}^2}.$$

See [**FP97**, Section 7] for detailed arguments. This formulation bridges classical enumerative geometry and modern intersection theory. A careful analysis of the boundary geometry [**FP97**, Section 0.6] eventually produces the renowned recursion formula:

$$N_d = \sum_{\substack{d_1 + d_2 = d \\ d_1, d_2 > 0}} N_{d_1} N_{d_2} \left(d_1^2 d_2^2 \binom{3d - 4}{3d_1 - 2} - d_1^3 d_2 \binom{3d - 4}{3d_1 - 1} \right),$$

which completely resolves Question 3.0.1.

Within this framework, it is natural to ask for a generalization of Gromov–Witten invariants to higher-genus cases. Given cohomology classes $\gamma_1, \dots, \gamma_n \in \mathrm{CH}^*(X)$, where X is a smooth projective variety, one might naively expect the Gromov–Witten invariants to take the form

$$\int_{[\overline{\mathcal{M}}_{g,n}(X,\beta)]} \mathrm{ev}_1^* \gamma_1 \cup \dots \cup \mathrm{ev}_n^* \gamma_n.$$

Here, however, we confront a major difficulty: as the genus increases, the geometric behavior of the moduli stack $\overline{\mathcal{M}}_{g,n}(X,\beta)$ deteriorates rapidly. For $g > 0$, the moduli space generally admits singularities and is often reducible, non-reduced, and non-equidimensional. In fact, for $g \geq 1$, $\overline{\mathcal{M}}_{g,n}(X,\beta)$ satisfies Murphy's Law in algebraic geometry, in the sense of Vakil [**Vak06**]. Consequently, the classical fundamental class of $\overline{\mathcal{M}}_{g,n}(X,\beta)$ does not behave well to perform intersection theory since it is not a homogeneous fundamental class in a single degree. A substitute is required if we still wish to perform intersection theory on this moduli space. This motivates the introduction of the **virtual fundamental class**, which plays the role of the classical fundamental class. Assuming its existence, we define the higher-genus Gromov–Witten invariants as

$$\langle \gamma_1, \dots, \gamma_n \rangle_{g,\beta}^X \coloneqq \int_{[\overline{\mathcal{M}}_{g,n}(X,\beta)]^{\mathrm{vir}}} \mathrm{ev}_1^* \gamma_1 \cup \dots \cup \mathrm{ev}_n^* \gamma_n.$$

We then ask what properties should a virtual fundamental class of a general moduli space satisfy. First and most naturally, if a moduli space is sufficiently well-behaved to admit a classical fundamental class, the virtual class should recover it. Second, the virtual fundamental class must be a cycle class of the "expected dimension."

To illustrate this, consider $\overline{\mathcal{M}}_{g,0}(X,\beta)$. Let $[\mu\colon C \to X]$ be a point in the moduli space. Assume that C is smooth and μ is an immersion, so there is a well-defined normal bundle $N_\mu = N_{C/X}$. The tangent space to $\overline{\mathcal{M}}_{g,0}(X,\beta)$ at $[\mu]$ is identified with $\mathrm{H}^0(C, N_\mu)$ [**Ser06**, Lemma 3.4.7]. If we further assume that the first-order infinitesimal deformations are unobstructed—that is, $\mathrm{H}^1(C, N_\mu) = 0$—then the Hirzebruch–Riemann–Roch theorem gives

$$\dim T_{[\mu]}(\overline{\mathcal{M}}_{g,0}(X,\beta)) = \mathrm{h}^0(C, N_\mu) = (\dim X - 3)(1-g) + \int_\beta c_1(X). \tag{3.0.1}$$

We define the *expected dimension* of $\overline{\mathcal{M}}_{g,0}(X,\beta)$ as this integer. Our "expectation" stems from the hypothetical vanishing of the obstruction space for first-order infinitesimal deformations. In general, this expected dimension serves as a lower bound for the actual dimension of any irreducible component of the moduli stack.

Therefore, we impose the requirement that the virtual fundamental class of a general moduli space $\mathcal{M}$ must reside in $\mathrm{CH}_{\mathrm{expected}}(\mathcal{M})$. For $\overline{\mathcal{M}}_{g,n}(X,\beta)$, this translates to:

$$[\overline{\mathcal{M}}_{g,n}(X,\beta)]^{\mathrm{vir}} \in \mathrm{CH}_{(\dim X - 3)(1-g) + \int_\beta c_1(X) + n}\left(\overline{\mathcal{M}}_{g,n}(X,\beta)\right).$$

From this perspective, while finding virtual classes may be straightforward, the profound challenge lies in constructing a virtual class of the correct expected dimension.

A major advantage of the moduli stack of stable maps $\overline{\mathcal{M}}_{g,n}(X,\beta)$ lies in its symmetries. If X is equipped with a torus action T (for example, the torus action on a homogeneous space like $X = \mathbb{P}^n$), this action naturally induces an action on the moduli stack $\overline{\mathcal{M}}_{g,n}(X,\beta)$ via composition. More precisely, let S be a scheme and $t \in T(S)$ be an S-valued point. For any S-family of stable maps $(\pi\colon \mathcal{C} \to S, \sigma_1, \dots, \sigma_n, \mu) \in \overline{\mathcal{M}}_{g,n}(X, \beta)(S)$, the induced action is given by

$$t \cdot (\pi\colon \mathcal{C} \to S, \sigma_1, \dots, \sigma_n, \mu) = (\pi\colon \mathcal{C} \to S, \sigma_1, \dots, \sigma_n, \mu'),$$

where $\mu'\colon \mathcal{C} \to X$ is defined by the composition

$$\mathcal{C} \xrightarrow{(\pi,\mathrm{id})} S \times \mathcal{C} \xrightarrow{t\times\mu} T \times X \xrightarrow{\sigma} X,$$

with σ denoting the torus action on X.

In the genus-zero case, this equivariant structure allows us to apply the Atiyah–Bott localization formula, reducing the calculation of intersection numbers to integrals over cohomology classes supported on the fixed loci. Specifically, we have

$$[\overline{\mathcal{M}}_{0,n}(\mathbb{P}^n, d)] = \sum_{\Gamma} \iota_{\Gamma *} \frac{[\mathcal{M}_\Gamma]}{e(N_\Gamma)},$$

where $e(N_\Gamma)$ is the Euler class (the top equivariant Chern class) of the normal bundle to the fixed locus Γ in $\overline{\mathcal{M}}_{0,n}(\mathbb{P}^n, d)$. This renders the seemingly intractable calculation of the fundamental class manageable; indeed, equivariant localization has become an indispensable tool in modern enumerative geometry.

A natural subsequent question is whether we can generalize this formula to the virtual setting. More precisely, does a localization formula of the form

$$[X]^{\mathrm{vir}} = \iota_* \sum_i \frac{[X_i]^{\mathrm{vir}}}{e(N_i^{\mathrm{vir}})} \in \mathrm{CH}_*^{\mathbb{C}^*}(X) \otimes \mathbb{Q}\left[t, \frac{1}{t}\right]$$

hold for an arbitrary algebraic scheme X, or even for a Deligne–Mumford stack, equipped with a perfect obstruction theory? This turns out to be true under a mild global resolution condition, which serves as one of the core topics of this chapter.

This chapter is organized as follows. We first review equivariant intersection theory following [**EG98a, EG98b**]. This includes the foundational definitions of equivariant Chow groups and their associated properties, such as functorial operations and equivariant Chern classes. We also review the standard localization formula for a nonsingular scheme X, preparing the ground for its generalization to the virtual setting. Next, following [**BF97**], we discuss the construction of the virtual fundamental class for a Deligne–Mumford stack $\mathcal{X}$. We outline the foundational ideas behind the construction and provide basic examples, restricting our attention initially to cases where $\mathcal{X}$ is a scheme. Finally, relying on this virtual machinery, we follow Graber and Pandharipande [**GP99**] to discuss the generalization of the Atiyah–Bott localization formula to the virtual setting. This will equip us with the complete theoretical toolkit required to prove the ELSV formula.

We establish the following notational conventions for the remainder of this chapter. For simplicity, any algebraic group G is implicitly assumed to be a linear algebraic group. We use T to denote an algebraic torus, which is an algebraic group isomorphic to $(\mathbb{C}^*)^n$ for some integer $n \geq 1$. Furthermore, when discussing group actions, T will consistently refer to the acting torus. As previously stated, we work exclusively over the field of complex numbers $\mathbb{C}$, and all schemes are assumed to be algebraic schemes over $\mathbb{C}$.

3.1. Equivariant Cohomology and Equivariant Intersection Theory

The exposition in this section proceeds in two main phases. We begin with the equivariant (co)homology of an algebraic variety to motivate the concept of finite-dimensional approximation, establishing the foundational definitions and properties. We place particular emphasis on the study of torus actions, as they will be our primary focus in the subsequent sections. Following this, we transition to the study of equivariant intersection theory, the motivation for which will become clearer after our detailed discussion of its topological counterpart. Finally, we explore the localization formula for equivariant (co)homology and Chow groups on nonsingular schemes, concluding with its applications to enumerative geometry.

3.1.1. Basic Definition and Properties. Equivariant (co)homology has played an essential role in topology since the 1950s, following Borel's introduction of the so-called "Borel construction." Let X be a topological space and G be a topological group. If G acts on X freely, there is a natural definition for equivariant cohomology: one simply defines $\mathrm{H}^*_G(X) \coloneqq \mathrm{H}^*(X/G)$, the standard cohomology of the quotient space. When the action of G is not free, however, Borel introduced the balanced product $\mathsf{E}G \times^G X$ by imposing the equivalence relation $(e \cdot g, x) \sim (e, g \cdot x)$ on $\mathsf{E}G \times X$, where $\mathsf{E}G \to \mathsf{B}G \coloneqq \mathsf{E}G/G$ is the universal principal G-bundle. One then defines the equivariant cohomology as $\mathrm{H}^*_G(X) \coloneqq \mathrm{H}^*(\mathsf{E}G \times^G X)$, which circumvents the issue of the group action being non-free. Because equivariant cohomology packages the topological information of X and the representation theory $\mathrm{Rep}(G)$ into a single invariant, it has been extensively studied and applied in topology and differential geometry for decades.

However, translating this machinery to the algebraic category—where X is an algebraic variety and G is an algebraic group—presents a major hurdle: in general, the universal principal G-bundle is infinite-dimensional. Topologically, $\mathsf{E}G$ must be a weakly contractible space. To construct such a space inductively within the category of CW complexes, an infinite number of cells across increasing dimensions must be attached to sequentially "kill" the generators of the lower-degree homotopy groups, forcing the dimension to become infinite. For example, when G is the multiplicative group $\mathbb{G}_m$ (topologically $\mathbb{C}^*$), we have $\mathsf{E}G = \mathbb{C}^\infty \setminus \{0\}$ and $\mathsf{B}G = \mathbb{CP}^\infty$.

Because infinite-dimensional spaces clearly fall outside the category of algebraic schemes, new ideas are required to make equivariant tools work in algebraic geometry. This leads to Totaro's idea of *finite-dimensional approximation*. Instead of relying on an actual contractible space, we seek a finite-dimensional algebraic space E_N that is "acyclic enough"—specifically, such that $\widetilde{\mathrm{H}}^i(\mathsf{E}_N) = 0$ for $i < N$—and upon which G acts freely. More formally, we introduce the following definition:

Definition 3.1.1 (Equivariant cohomology in algebraic geometry)**.** Let G be a linear algebraic group and let X be a scheme. Then for each integer $N > 0$ there exist nonsingular finite-dimensional algebraic varieties E_N and B_N such that $\widetilde{\mathrm{H}}^i(\mathsf{E}_N) = 0$ for $i < N$, G acts freely on E_N, and $\mathsf{E}_N \to \mathsf{B}_N = \mathsf{E}_N/G$ is a principal G-bundle. One then defines the equivariant cohomology as

$$\mathrm{H}^i_G(X) \coloneqq \mathrm{H}^i(\mathsf{E}_N \times^G X), \quad \text{for } i < N.$$

Before proceeding to a general discussion of well-definedness, we first examine the fundamental case where the group is a torus.

Example 3.1.2 (Equivariant cohomology for a torus action)**.** Let $T = \mathbb{G}_m$. We can take $\mathsf{E}_N \coloneqq \mathbb{C}^N \setminus \{0\}$ and $\mathsf{B}_N \coloneqq \mathbb{P}^{N-1}$. Clearly, T acts freely on E_N via scalar multiplication: $(z_1, \ldots, z_N) \cdot s \coloneqq (sz_1, \ldots, sz_N)$. Moreover, since E_N deformation retracts onto the sphere S^{2N-1}, we have $\widetilde{\mathrm{H}}^i(\mathsf{E}_N) = \widetilde{\mathrm{H}}^i(S^{2N-1}) = 0$ for $i < 2N-1$. Therefore, the equivariant cohomology of any space X equipped with a T-action is defined as

$$\mathrm{H}^i_T(X) \coloneqq \mathrm{H}^i\left((\mathbb{C}^N \setminus \{0\}) \times^T X\right), \quad \text{for } i < 2N-1.$$

Furthermore, when X is a single point, we obtain the equivariant coefficient ring $\Lambda_T := \mathrm{H}^*_T(*) = \mathbb{Z}[t]$, where $\deg t = 2$. This follows directly from the cohomology of $\mathbb{P}^{N-1}$.

When $T = (\mathbb{G}_m)^n$, we similarly set $\mathsf{E}_N := \left(\mathbb{C}^N \setminus \{0\}\right)^n$ and $\mathsf{B}_N := \left(\mathbb{P}^{N-1}\right)^n$, with T acting on E_N component-wise. Consequently, the equivariant coefficient ring becomes $\Lambda_T = \mathbb{Z}[t_1, \ldots, t_n]$, where $\deg t_i = 2$ for all $i = 1, \ldots, n$.

For a general linear algebraic group G, we first examine the fundamental case where $G = \mathrm{GL}(n)$. The construction of E_N and B_N closely parallels the torus case seen in Example 3.1.2. We draw motivation from the classical topological setting, where the classifying space of $\mathrm{GL}(n)$ is the infinite Grassmannian $\mathrm{Gr}(n, \infty)$, and the universal principal $\mathrm{GL}(n)$-bundle is the space of linear embeddings $\mathrm{Emb}(\mathbb{C}^n, \mathbb{C}^\infty)$.

To create a finite-dimensional approximation, we naturally take $\mathsf{E}_N := \mathrm{Emb}(\mathbb{C}^n, \mathbb{C}^N)$ and $\mathsf{B}_N := \mathrm{Gr}(n, N)$. We can identify E_N with the open subvariety of $n \times N$ matrices of full rank n. There is a natural, free right-action of $\mathrm{GL}(n)$ on E_N given by matrix multiplication. Let $\Omega_{n-1} := \mathrm{Hom}(\mathbb{C}^n, \mathbb{C}^N) \setminus \mathsf{E}_N$ denote the complement, which is the closed locus of $n \times N$ matrices of rank at most $n-1$. A dimension count shows that the complex codimension of Ω_{n-1} in the affine space $\mathbb{C}^{nN}$ is $(n-(n-1))(N-(n-1)) = N-n+1$. Because E_N is the complement of a closed subvariety of real codimension $2(N-n+1)$ inside a contractible space, the long exact sequence of relative cohomology implies that $\widetilde{\mathrm{H}}^i(\mathsf{E}_N) = 0$ for $i \leq 2(N-n)$.

When G is an arbitrary linear algebraic group, we can embed G as a closed subgroup of $\mathrm{GL}(n)$ for some sufficiently large n. Since $\mathrm{GL}(n)$ acts freely on the space of embeddings, its closed subgroup G naturally inherits this free action; thus, we can retain our choice of the approximation space $\mathsf{E}_N := \mathrm{Emb}(\mathbb{C}^n, \mathbb{C}^N)$ without modification. To verify that the classifying space $\mathsf{B}_N = \mathsf{E}_N/G$ exists as a scheme, we rely on a theorem [**Bor91**, Theorem 6.8], which guarantees that the homogeneous space $\mathrm{GL}(n)/G$ is a smooth quasi-projective variety. Consequently, the quotient E_N/G can be identified with the total space of an algebraic fiber bundle over the Grassmannian $\mathrm{Gr}(n, N)$, which is the quotient $\mathsf{E}_N/\mathrm{GL}(n)$, with the fiber $\mathrm{GL}(n)/G$. Since the base is a smooth projective variety and the fiber is a smooth quasi-projective variety, the total space B_N exists as a smooth quasi-projective scheme. This constructively satisfies the existence criteria required by Definition 3.1.1.

The subsequent question is whether this definition is independent of the choice of the finite-dimensional approximation spaces E_N and B_N. More precisely, suppose that $\mathsf{E} \to \mathsf{B}$ and $\mathsf{E}' \to \mathsf{B}'$ are two principal G-bundles satisfying the conditions in Definition 3.1.1. We must show that for any space X admitting a G-action, there is a canonical isomorphism

$$\mathrm{H}^i(\mathsf{E} \times^G X) \cong \mathrm{H}^i(\mathsf{E}' \times^G X), \quad \text{for } i < N.$$

Since we assume G is linear and we work over the complex numbers $\mathbb{C}$, the proof reduces to an elegant topological maneuver known as Bogomolov's double fibration argument. We consider the product space $\mathsf{E} \times \mathsf{E}'$, equipped with the diagonal G-action. This yields a natural commutative diagram of quotient spaces:

$$\begin{tikzcd}
\mathsf{E} \times X \arrow[d] & (\mathsf{E} \times \mathsf{E}') \times X \arrow[l] \arrow[r] \arrow[d] & \mathsf{E}' \times X \arrow[d] \\
\mathsf{E} \times^G X & (\mathsf{E} \times \mathsf{E}') \times^G X \arrow[l, "p_1"'] \arrow[r, "p_2"] & \mathsf{E}' \times^G X
\end{tikzcd}$$

Notice that the left projection $p_1 \colon (\mathsf{E} \times \mathsf{E}') \times^G X \to \mathsf{E} \times^G X$ is an algebraic fiber bundle with fiber E'. Because E' is chosen to be nonsingular, this algebraic bundle remains locally trivial in the complex analytic topology. We then pass to the analytic topology via Serre's GAGA principle to compute the cohomology. By our initial assumption, the fiber E' satisfies $\widetilde{\mathrm{H}}^i(\mathsf{E}') = 0$ for $i < N$. Applying the Leray–Hirsch theorem (or analyzing the Serre spectral sequence of the fibration) in the topological category, we deduce that the pullback map p_1^* induces an isomorphism on cohomology in degrees $i < N$. By symmetry, the right projection p_2 likewise induces an isomorphism in the same range. Composing these maps provides the required canonical isomorphism between $\mathrm{H}^i(\mathsf{E} \times^G X)$ and $\mathrm{H}^i(\mathsf{E}' \times^G X)$.

Remark 3.1.3. For a general linear algebraic group G, the approximation space E and the classifying space B in Definition 3.1.1 can be chosen to be nonsingular quasi-projective varieties, as demonstrated above. However, even if X is nonsingular, the balanced product $\mathsf{E} \times^G X$ is not necessarily a scheme. Because the free action of G defines a smooth equivalence relation, the quotient is only guaranteed to exist as an algebraic space [**Alp26**, Corollary 5.5.12]. We refer the reader to [**EG98a**, Proposition 23] for sufficient conditions under which this balanced product is indeed a scheme. Most notably, the balanced product is always a scheme when X is a quasi-projective scheme equipped with a linearized G-action, or when G is $\mathrm{GL}(n)$ or any other "special group" in the sense of Chevalley (meaning all principal G-bundles are locally trivial in the Zariski topology).

On the other hand, if we consider the ordinary complex analytic topology and continue to choose a nonsingular approximation space E, the balanced product $\mathsf{E} \times^G X$ is a complex analytic space. Furthermore, if X is nonsingular, this quotient is a complex manifold. We refer to [**AF24**, Proposition 2.2.6] for details.

We next examine the basic properties of equivariant cohomology, beginning with its functoriality. Fix a linear algebraic group G. Let $\varphi\colon X \to X'$ be a G-equivariant morphism of schemes. This determines a natural morphism $\mathsf{E} \times^G X \to \mathsf{E} \times^G X'$, which in turn induces a pullback map $\varphi^*\colon \mathrm{H}^*_G(X') \to \mathrm{H}^*_G(X)$. In particular, taking X' to be a point yields a canonical map $\Lambda_G := \mathrm{H}^*_G(\mathrm{pt}) \to \mathrm{H}^*_G(X)$, endowing $\mathrm{H}^*_G(X)$ with the structure of a graded anti-commutative Λ_G-algebra.

Equivariant cohomology is also functorial with respect to the group. Let $\varphi\colon G \to G'$ be a homomorphism of linear algebraic groups, and let $f\colon X \to X'$ be a morphism that is equivariant with respect to φ. This means the following diagram commutes.

$$\begin{array}{ccc} G \times X & \xrightarrow{\varphi \times f} & G' \times X' \\ \big\downarrow{\scriptstyle \sigma_G} & & \big\downarrow{\scriptstyle \sigma_{G'}} \\ X & \xrightarrow{\quad f \quad} & X' \end{array}$$

We construct the induced pullback map $f^*\colon \mathrm{H}^*_{G'}(X') \to \mathrm{H}^*_G(X)$ as follows. Choose approximation spaces $\mathsf{E} \to \mathsf{B}$ and $\mathsf{E}' \to \mathsf{B}'$ for G and G', respectively. This yields a natural sequence of maps:

$$(\mathsf{E} \times \mathsf{E}') \times^G X \to (\mathsf{E} \times \mathsf{E}') \times^G X' \to \mathsf{E}' \times^{G'} X'.$$

Here, G acts on $\mathsf{E} \times \mathsf{E}'$ on the right via $(e, e') \cdot g := (e \cdot g, e' \cdot \varphi(g))$, and G acts on X' implicitly through φ. Note that if the reduced cohomology satisfies $\widetilde{\mathrm{H}}^i(\mathsf{E}) = \widetilde{\mathrm{H}}^i(\mathsf{E}') = 0$ for $i < N$, the Künneth formula ensures that $\widetilde{\mathrm{H}}^i(\mathsf{E} \times \mathsf{E}') = 0$ in the same range. Consequently, passing to cohomology induces the desired pullback $f^*\colon \mathrm{H}^*_{G'}(X') \to \mathrm{H}^*_G(X)$. By applying Bogomolov's double fibration argument once again, one can verify that this construction is independent of the choice of approximation spaces.

Following the philosophy discussed above, we can define equivariant geometric invariants in a parallel manner by simply translating them to the ordinary, non-equivariant setting on these finite-dimensional approximations. For example, let X be a G-space and $V \subseteq X$ be a G-invariant closed subvariety. The equivariant fundamental class $[V]^G$ can be defined as the fundamental class of the quotient $[\mathsf{E} \times^G V]$. Similarly, let $E \to X$ be an equivariant vector bundle on a non-singular X. The diagonal action of G on $\mathsf{E} \times E$ yields a quotient $\mathsf{E} \times^G E \to \mathsf{E} \times^G X$, which is an ordinary vector bundle over the balanced product. One then defines the equivariant Chern classes as $c_k^G(E) \coloneqq c_k(\mathsf{E} \times^G E)$. By applying Bogomolov's double fibration argument once again, one can verify that these definitions are independent of the choice of E. Note that, as discussed in Remark 3.1.3, $\mathsf{E} \times^G X$ is always an algebraic space and a complex manifold when X is non-singular.

Example 3.1.4 (The equivariant coefficient ring $\Lambda_{\mathrm{GL}(n)}$)**.** Consider the general linear group $G = \mathrm{GL}(n)$ and let $V \cong \mathbb{C}^n$ be its standard representation. We claim that the equivariant coefficient ring is given by $\Lambda_{\mathrm{GL}(n)} \cong \mathbb{Z}[c_1, \dots, c_n]$, where $c_i = c_i^G(V)$ are the equivariant Chern classes of V. Recall from our previous discussion that we can choose the finite-dimensional approximation space to be $\mathsf{E}_N = \mathrm{Emb}(V, \mathbb{C}^N)$, with the classifying space being the Grassmannian $\mathsf{B}_N = \mathrm{Gr}(n, N)$.

Notice that the principal G-bundle $\mathsf{E}_N \to \mathsf{B}_N$ can be naturally identified with the frame bundle of the tautological subbundle $\mathcal{S} \subset \mathrm{Gr}(n, N) \times \mathbb{C}^N$ over the Grassmannian. Consequently, the balanced product $\mathsf{E}_N \times^G V$ associated to the standard representation is canonically isomorphic to the tautological subbundle $\mathcal{S}$ itself. Since the cohomology ring of the Grassmannian $\mathrm{Gr}(n, \infty)$ is freely generated by the Chern classes of the universal tautological subbundle, it follows that $\Lambda_{\mathrm{GL}(n)}$ is a polynomial ring generated by the equivariant Chern classes of V.

Example 3.1.5 (Torus actions revisited)**.** Building upon Example 3.1.2, we further investigate the rich structural properties of torus actions. We have already seen that when $T = \mathbb{G}_m$, the equivariant coefficient ring is $\Lambda_T = \mathbb{Z}[t]$ with $\deg t = 2$. Furthermore, from the discussion in Example 3.1.4, there is a canonical identification $t = c_1^T(\mathbb{C}_1)$, where $\mathbb{C}_1 \cong \mathbb{C}$ denotes the standard one-dimensional representation of T with weight 1.

Now consider $T = (\mathbb{G}_m)^n$ acting on $V \cong \mathbb{C}^n$ via the standard diagonal representation. For each $i = 1, \dots, n$, this action isolates a one-dimensional representation $\mathbb{C}_{t_i}$ defined by $z \cdot v \coloneqq z_i \cdot v$, yielding a canonical splitting $V = \bigoplus_{i=1}^n \mathbb{C}_{t_i}$. Consequently, the equivariant Chern classes are given by $c_i^T(V) = e_i(t_1, \dots, t_n)$, where e_i is the i-th elementary symmetric polynomial and $t_i = c_1^T(\mathbb{C}_{t_i})$. This gives an identifaction of the t_i in the coefficient ring $\Lambda_T \cong \mathbb{Z}[t_1, \dots, t_n]$. If we explicitly choose the approximation spaces $\mathsf{E}_N \coloneqq (\mathbb{C}^N \setminus \{0\})^n$ and $\mathsf{B}_N \coloneqq (\mathbb{P}^{N-1})^n$ as we did in Example 3.1.2, each Chern class t_i can be also identified with the first Chern class of the tautological line bundle pulled back from the i-th factor of B_N.

Comparing this to our previous calculations, we have $\Lambda_T \cong \mathbb{Z}[t_1, \dots, t_n]$ and $\Lambda_{\mathrm{GL}(n)} \cong \mathbb{Z}[c_1, \dots, c_n]$. The natural inclusion of the torus $T \hookrightarrow \mathrm{GL}(n)$ induces a pullback ring homomorphism $\Lambda_{\mathrm{GL}(n)} \hookrightarrow \Lambda_T$ given by $c_i \mapsto e_i(t_1, \dots, t_n)$. This maps $\Lambda_{\mathrm{GL}(n)}$ isomorphically onto the subring of symmetric polynomials in Λ_T, serving as an algebraic manifestation of the splitting principle over a single point.

Finally, we can provide an intrinsic description of this ring using representation theory. Let $M \coloneqq \mathrm{Hom}(T, \mathbb{G}_m)$ be the character lattice of the torus T. Each character $\chi \in M$ determines an equivariant line bundle (a one-dimensional representation) $\mathbb{C}_\chi$ defined by $z \cdot v \coloneqq \chi(z)v$ for $z \in T$ and $v \in \mathbb{C}$. This induces a natural ring homomorphism $\mathrm{Sym}^* M \to \Lambda_T$ defined by extending the assignment $\chi \mapsto c_1^T(\mathbb{C}_\chi)$. By fixing an isomorphism $T \cong (\mathbb{G}_m)^n$, which naturally identifies $M \cong \mathbb{Z}^n$, it becomes clear that this canonical homomorphism is, in fact, an isomorphism.

We close our discussion of the topological counterpart by introducing equivariant homology, which serves as a natural bridge to equivariant Chow groups. It is well established that the most suitable topological homology framework for complex algebraic varieties is Borel–Moore homology [**Ful98**, Chapter 19]. Therefore, it is natural to construct the equivariant homology $\mathrm{H}_*^G(X)$ based on the Borel–Moore theory.

Recall that if a scheme X can be embedded as a closed subspace of an oriented real n-dimensional manifold M, we have a canonical isomorphism $\overline{\mathrm{H}}_i(X) \cong \mathrm{H}^{n-i}(M, M \setminus X)$ [**Ful98**, Chapter 19], which can be viewed as a generalized form of Poincaré duality. This suggests a natural candidate for the equivariant definition: if a G-space X admits an equivariant embedding as a closed invariant subspace of an oriented G-manifold M of real dimension n, we should define $\mathrm{H}_i^G(X) \coloneqq \mathrm{H}_G^{n-i}(M, M \setminus X)$.

By our previous construction, this relative equivariant cohomology evaluates to

$$\mathrm{H}^{n-i}\left(\mathsf{E}_N \times^G M, (\mathsf{E}_N \times^G M) \setminus (\mathsf{E}_N \times^G X)\right)$$

provided that $n - i < N$ (and thus $i > n - N$). As discussed in Remark 3.1.3, the approximation space E_N can be chosen such that the balanced product $\mathsf{E}_N \times^G M$ is at least an oriented real manifold of dimension $\dim_{\mathbb{R}} M + \dim_{\mathbb{R}} \mathsf{B} = n + \dim_{\mathbb{R}} \mathsf{E} - \dim_{\mathbb{R}} G$, and it contains $\mathsf{E}_N \times^G X$ as a closed subspace. By applying the duality theorem once again to this quotient, the relative cohomology group transforms into the Borel–Moore homology group $\overline{\mathrm{H}}_{i+\dim_{\mathbb{R}} \mathsf{B}}(\mathsf{E}_N \times^G X)$. This eliminates the auxiliary manifold M from the final expression, motivating the following definition of equivariant homology.

Definition 3.1.6 (Equivariant Homology)**.** Let G be a linear algebraic group and let X be a scheme admitting a G-action. For a sufficiently large approximation space E_N (acyclic in degrees less than N), the equivariant (Borel–Moore) homology groups of X are defined as

$$\mathrm{H}_i^G(X) \coloneqq \overline{\mathrm{H}}_{i+\dim_{\mathbb{R}} \mathsf{B}_N}(\mathsf{E}_N \times^G X) \quad \text{for } i > -N.$$

By employing Bogomolov's double fibration argument once more, or simply by passing to the cohomology case via Poincaré duality, one can easily verify that this definition is independent of the choice of the approximation space E_N.

Motivated by the construction of equivariant homology, we are now positioned to introduce the equivariant Chow groups, which follow the exact same finite-dimensional approximation philosophy. Let E be a sufficiently acyclic approximation space and $\mathsf{B} = \mathsf{E}/G$ its corresponding classifying space. The equivariant Chow group of a G-space X in degree i is defined as

$$\mathrm{CH}_i^G(X) \coloneqq \mathrm{CH}_{i+\dim \mathsf{E} - \dim G}(\mathsf{E} \times^G X) = \mathrm{CH}_{i+\dim \mathsf{B}}(\mathsf{E} \times^G X).$$

Here, we treat the balanced product $\mathsf{E} \times^G X$ as an algebraic space, as explained in Remark 3.1.3. Although we previously discussed the (naive) Chow groups of an arbitrary Deligne–Mumford stack, we emphasize that the foundational constructions from the first eight chapters of [**Ful98**] generalize to algebraic spaces with no essential modifications. Because the diagonal map $X \to X \times X$ of an algebraic space is a monomorphism representable by schemes (Theorem 2.3.31), the deformation to the normal cone still functions effectively. We refer the reader to [**EG98a**, Section 6.1] for a detailed discussion of intersection theory on algebraic spaces. Note that while we could assume X itself is merely an algebraic space admitting a G-action, we restrict our attention to algebraic schemes, as this level of generality is rarely needed in our practice.

We give the definition of the equivariant Chow groups in a more precise way, following [**EG98a**, Definition–Proposition 1]. Let X be an n-dimensional scheme equipped with an action by an algebraic group G of dimension g. One can choose an l-dimensional representation V of G containing a G-invariant open subset U upon which G acts freely, such that the codimension of the complement $V \setminus U$ is greater than $n - i$. Furthermore, U is chosen so that the quotient map $U \to U/G$ is a principal G-bundle in the category of schemes.

The i-th equivariant Chow group is then defined as:

$$\mathrm{CH}_i^G(X) \coloneqq \mathrm{CH}_{i+l-g}(X \times^G U).$$

This definition is independent of the choice of representation V and open set U, provided the complement $V \setminus U$ has sufficiently high codimension.

This algebraic formulation carries a distinct technical advantage: by utilizing Bogomolov's double fibration argument, one can apply the excision exact sequence for Chow groups alongside the canonical isomorphism between the Chow group of a vector bundle and its base scheme to verify the independence of the representation.

Moreover, the construction of the principal bundle $U \to U/G$ is identical to the algebraic approximation of the universal bundle $\mathsf{E}G \to \mathsf{B}G$ discussed in Definition 3.1.1. To explicitly construct these representation spaces, one first resolves the case for the general linear group GL_m. For an arbitrary algebraic group G, one embeds $G \hookrightarrow \mathrm{GL}_m$ and leverages the fact that the quotient GL_m / G is a smooth quasi-projective variety to restrict the free GL_m-action down to G, as detailed in [**EG98a**, Lemma 9].

We establish functoriality using the same finite-dimensional approximation philosophy discussed earlier for equivariant cohomology. However, we need to exercise additional care here: to define intersection-theoretic operations, we need to descend specific geometric properties, such as properness or flatness, to the induced morphisms between the balanced products.

Consider a property of morphisms P that is stable under base change and local on the target in the fppf topology. As we will demonstrate, equivariant Chow groups inherit the identical functorial operations of ordinary Chow groups for equivariant morphisms satisfying P. If a G-equivariant morphism $f \colon X \to Y$ satisfies P, then the induced product map $\mathrm{id} \times f \colon \mathsf{E} \times X \to \mathsf{E} \times Y$ satisfies P as well. We can relate this to the balanced products via the natural Cartesian square:

$$\begin{tikzcd} \mathsf{E} \times X \arrow[r, "\mathrm{id} \times f"] \arrow[d, "\pi_X"'] \arrow[dr, phantom, "\square"] & \mathsf{E} \times Y \arrow[d, "\pi_Y"] \\ \mathsf{E} \times^G X \arrow[r, "\overline{f}"] & \mathsf{E} \times^G Y \end{tikzcd}$$

The quotient map $\pi_Y \colon \mathsf{E} \times Y \to \mathsf{E} \times^G Y$ is a principal G-bundle, making it a smooth and surjective morphism. Therefore, by fppf descent, the induced morphism $\overline{f} \colon \mathsf{E} \times^G X \to \mathsf{E} \times^G Y$ inherits the property P.

This descent mechanism allows us to transfer operations such as proper pushforwards, flat pullbacks, and refined Gysin maps for regular embeddings into the equivariant Chow setting. To conclude the construction, one applies Bogomolov's double fibration argument once again to verify that these induced maps are independent of the choice of the approximation space E.

3.1.2. Localization Formula. In this subsection, we restrict our attention to torus actions on a scheme X.

One of the most remarkable features of equivariant cohomology and intersection theory is that global equivariant invariants can often be completely determined by the fixed points. This philosophy allows us to carry out global computations using only local data around the fixed locus. This approach is exceptionally powerful for torus actions, which, as we have seen, admit rich properties.

The foundational theorem of this philosophy is the Localization Formula. Let X be a scheme admitting a T-action, and let $X^T \subseteq X$ be the fixed locus, which is a closed subscheme of X. There is a natural pushforward map of equivariant Chow groups $\iota_* \colon \mathrm{CH}^T_*(X^T) \to \mathrm{CH}^T_*(X)$. Similarly, on the level of equivariant Borel–Moore homology, there is a pushforward map $\iota_* \colon \mathrm{H}^T_*(X^T) \to \mathrm{H}^T_*(X)$. In general, these maps are neither injective nor surjective when $X^T \neq \varnothing$. For example, consider $X = \mathbb{P}^1$ equipped with a $T = \mathbb{G}_m$ action via a non-trivial weight χ.

However, a profound result states that, up to inverting certain torsion elements, the equivariant Chow groups (and homology groups) of the fixed locus X^T are isomorphic to those of the entire space X.

Theorem 3.1.7 (Localization Formula). *Let X be a scheme admitting a torus action by T. Let S be the multiplicative set generated by all nonzero characters in $M = \mathrm{Hom}(T, \mathbb{G}_m)$. Then the pushforward induces an isomorphism*

$$S^{-1}\iota_* \colon S^{-1}\,\mathrm{CH}^T_*(X^T) \xrightarrow{\sim} S^{-1}\,\mathrm{CH}^T_*(X).$$

as Λ_T-modules. Note that $\mathrm{Sym}^* M \cong \Lambda_T$ as shown in Example 3.1.5.

An analogous isomorphism holds for equivariant Borel–Moore homology groups.

We will not reproduce the proof of the most general version of this theorem (for arbitrary schemes), and instead refer the reader to [**EG98b**, Theorem 1]. Here, we focus on the special case where X is a smooth variety. This assumption equips us with tools from the theory of algebraic groups and yields a geometric proof using the self-intersection formula to construct the inverse of $S^{-1}\iota_*$.

Lemma 3.1.8.

(1) (Smoothness of the fixed locus) Let G be a smooth linearly reductive group variety acting on a smooth variety X. Then the fixed locus subscheme X^G is smooth.

(2) (Luna's Slicing Theorem) Let G be a reductive algebraic group acting on a smooth affine scheme X. Let x be a closed point of X such that its orbit $\mathcal{O}_x$ is closed. Then there exists a G-equivariant étale neighbourhood $U \to X$ of $\mathcal{O}_x$ which is equivariantly isomorphic to an étale neighbourhood of the zero section of the normal bundle $N_{\mathcal{O}_x/X}$.

For the proofs of these foundational results, we refer to [**Mil17**, Theorem 13.1, Lemma 13.36].

Let us examine how to deduce Theorem 3.1.7 when X is a smooth variety. Let $Z \subseteq X^T$ be a connected component of the fixed locus of codimension d, and let $\iota \colon Z \hookrightarrow X$ be the natural closed immersion. By Lemma 3.1.8, Z is smooth. Consequently, the normal bundle $N_{Z/X}$ is a well-defined equivariant vector bundle of rank d. The self-intersection formula [**Ful98**, Corollary 6.3] states that the composition

$$\mathrm{CH}^T_*(Z) \xrightarrow{\iota_*} \mathrm{CH}^T_*(X) \xrightarrow{\iota^*} \mathrm{CH}^T_*(Z)$$

is given by multiplication by the equivariant top Chern class $c^T_d(N_{Z/X})$. Our immediate goal is to identify this class.

Fix $p \in Z$. Luna's Slicing Theorem (Lemma 3.1.8) provides a T-equivariant étale neighbourhood of p that is equivariantly isomorphic to a neighbourhood of the origin in its normal space T_pX. Under this local isomorphism, the fixed locus Z corresponds to the 0-weight space T_pZ. Consequently, the normal space $N_p = T_pX/T_pZ$ decomposes into non-trivial weight spaces. Since Z is connected, the multiset of weights of $N_{Z/X}$ is constant across Z.

This implies there are nonzero characters $\chi_1, \ldots, \chi_d$ such that at every $p \in Z$, the torus T acts on the fiber N_p with weights $\chi_1, \ldots, \chi_d$. If we view the equivariant Chow ring as an algebra over the coefficient ring Λ_T, the top Chern class $c_d^T(N_{Z/X})$ takes the form

$$c_d^T(N_{Z/X}) = \chi_1 \cdots \chi_d + \sum_{i=1}^{d} a_{d-i} c_i, \quad a_j \in \Lambda_T^j,\ c_i \in \mathrm{CH}^i(Z).$$

Since Z is a finite-dimensional variety, any element $c_i \in \mathrm{CH}^i(Z)$ of positive degree is nilpotent. Since $\chi_1 \cdots \chi_d$ is a product of nonzero characters, it lies in the multiplicative set S and is invertible in $S^{-1}\Lambda_T$. Therefore, $c_d^T(N_{Z/X})$ is invertible in the localized Chow group $S^{-1}\,\mathrm{CH}_*^T(Z)$, and the localized composition

$$S^{-1}\iota^* \circ S^{-1}\iota_* \colon S^{-1}\,\mathrm{CH}_*^T(Z) \to S^{-1}\,\mathrm{CH}_*^T(Z)$$

is an isomorphism. This proves that $S^{-1}\iota_*$ is injective.

To establish that $S^{-1}\iota_*$ is surjective, more technical tools are required. If the equivariant Chow ring $\mathrm{CH}^*(X^T)$ is free as a $\mathbb{Z}$-module and there exists a collection of elements in $\mathrm{CH}_T^*(X)$ whose restrictions form a basis of $\mathrm{CH}^*(X)$, surjectivity follows from the graded Nakayama's lemma [**AF24**, Chapter 5]. In the topological setting, this is proved by the Leray–Hirsch theorem. Although this technical condition is redundant for the general theorem, the proof is involved and requires results from the theory of algebraic groups; we refer to [**AF24**, Chapter 7] for details. One way to approach is to analyze the equivariant excision exact sequence

$$\mathrm{CH}_*^T(X^T) \xrightarrow{\iota_*} \mathrm{CH}_*^T(X) \to \mathrm{CH}_*^T(X \setminus X^T) \to 0$$

and show that after localization the equivariant Chow group of the complement vanishes.

This argument also leads directly to a computationally powerful corollary of Theorem 3.1.7.

Corollary 3.1.9 (Integration Formula / Bott Residue Formula)**.** *Following the notation above, let X be a smooth variety equipped with a torus action. For any class $\alpha \in S^{-1}\,\mathrm{CH}_*^T(X)$, we have*

$$\alpha = \sum_{Z \subseteq X^T} (\iota_Z)_* \left(\frac{\iota_Z^* \alpha}{c_{d_Z}^T(N_{Z/X})} \right),$$

where the sum runs over all connected components Z of X^T, $\iota_Z \colon Z \hookrightarrow X$ is the natural inclusion, and d_Z is the codimension of Z in X.

Proof. Since Theorem 3.1.7 guarantees that $S^{-1}\iota_*$ is surjective, we can write $\alpha = \sum_Z (\iota_Z)_*(\beta_Z)$ for some classes $\beta_Z \in S^{-1}\,\mathrm{CH}_*^T(Z)$. Applying the pullback $(\iota_Z)^*$ to both sides, and using the fact that the fixed components are disjoint, the self-intersection formula yields

$$\iota_Z^* \alpha = (\iota_Z)^*(\iota_Z)_*(\beta_Z) = c_{d_Z}^T(N_{Z/X}) \cdot \beta_Z.$$

Since the top Chern class of the normal bundle is invertible in the localized Chow group, we can divide by it to obtain

$$\beta_Z = \frac{\iota_Z^* \alpha}{c_{d_Z}^T(N_{Z/X})}.$$

Substituting this back into the expression for α yields the desired formula. $\square$

Example 3.1.10 (Torus action on projective spaces)**.** To illustrate the power of the localization theorems, we study the fundamental case of $T = \mathbb{G}_m$ acting on the projective space $X = \mathbb{P}^{n-1}$ with distinct weights $\chi_1, \dots, \chi_n$. The action is given by $a \cdot [z_1, \dots, z_n] = [\chi_1(a)z_1, \dots, \chi_n(a)z_n]$ for any $a \in T$.

We first compute the equivariant cohomology ring (or equivariant Chow ring) $\mathrm{H}_T^*(\mathbb{P}^{n-1})$. Let $V \cong \mathbb{C}^n$ be the standard representation. There is a commutative diagram:

$$\begin{tikzcd}
\mathsf{E} \times^T \mathbb{P}(V) \arrow[rr, "\cong"] \arrow[dr] & & \mathbb{P}(\mathsf{E} \times^T V) \arrow[dl] \\
& \mathsf{B} &
\end{tikzcd}$$

By the formula for projective bundles, the cohomology ring is

$$\mathrm{H}_T^*(\mathbb{P}^{n-1}) \cong \frac{\Lambda_T[\zeta]}{(\zeta^n + c_1\zeta^{n-1} + \cdots + c_n)},$$

where $\zeta = c_1^T(\mathcal{O}(1))$ is the hyperplane class and $c_i = c_i^T(V)$ for each $i = 1, \dots, n$. By the splitting principle (Example 3.1.5), V decomposes into weight spaces, allowing us to factor the relation as

$$\mathrm{H}_T^*(\mathbb{P}^{n-1}) \cong \frac{\Lambda_T[\zeta]}{\prod_{i=1}^n(\zeta + \chi_i)} = \frac{\mathbb{Z}[t, \zeta]}{\prod_{i=1}^n(\zeta + \chi_i)},$$

where each χ_i is identified with the first Chern class $c_1^T(\mathbb{C}_{\chi_i})$. The calculation for the equivariant Chow ring is identical.

The fixed points of the torus action on $\mathbb{P}^{n-1}$ correspond to $[1, 0, \dots, 0], \dots, [0, \dots, 0, 1]$, which we denote by $p_1, \dots, p_n$. The tangent space of $\mathbb{P}^{n-1}$ at a fixed point p_i is

$$T_{p_i}\mathbb{P}^{n-1} = \mathrm{Hom}(L_i, V/L_i) \cong \bigoplus_{j \neq i} L_i^\vee \otimes L_j,$$

where L_i is the i-th coordinate line. Because L_i is the eigenspace corresponding to χ_i, the weights of the tangent space are $\chi_j - \chi_i$. Therefore, the top Chern class $c_{n-1}^T(T_{p_i}\mathbb{P}^{n-1})$ is

$$c_{n-1}^T(T_{p_i}\mathbb{P}^{n-1}) = \prod_{j \neq i}(\chi_j - \chi_i).$$

Note that equivariant Poincaré duality identifies the homology group $\mathrm{H}_*^T(\mathbb{P}^{n-1})$ with the cohomology ring $\mathrm{H}_T^*(\mathbb{P}^{n-1})$. Under this isomorphism, Corollary 3.1.9 applies also to cohomology classes.

If we push the identity from Corollary 3.1.9 forward to a point via the projection map $\pi\colon \mathbb{P}^{n-1} \to \mathrm{pt}$, we obtain the classical Bott residue formula for the degree of a class $\alpha \in S^{-1}\mathrm{H}_T^*(\mathbb{P}^{n-1})$:

$$\int_{\mathbb{P}^{n-1}} \alpha = \sum_{i=1}^n \frac{\alpha|_{p_i}}{c_{n-1}^T(T_{p_i}\mathbb{P}^{n-1})} = \sum_{i=1}^n \prod_{j \neq i} \frac{\alpha|_{p_i}}{\chi_j - \chi_i}.$$

Specifically, setting $\alpha = \zeta = c_1^T(\mathcal{O}(1))$ to be the hyperplane class, its restriction at p_i is $-\chi_i$. Substituting α^k into the residue formula yields

$$\int_{\mathbb{P}^{n-1}} \alpha^k = \sum_{i=1}^n \prod_{j \neq i} \frac{(-\chi_i)^k}{\chi_j - \chi_i}.$$

On the other hand, we know that $\int_{\mathbb{P}^{n-1}} \alpha^k = 0$ for $k < n-1$, and 1 for $k = n-1$. This produces a non-trivial algebraic identity:

$$\delta_{k,n-1} = \sum_{i=1}^{n} \prod_{j \neq i} \frac{(-\chi_i)^k}{\chi_j - \chi_i} \quad \text{for } k \leq n-1.$$

Example 3.1.11 (Schubert's calculus for lines)**.** We use the Integration Formula to solve a classical problem in enumerative geometry: given $2n-4$ general subspaces of codimension 2 in the projective space $\mathbb{P}^{n-1}$, how many lines intersect all of them?

Lines in $\mathbb{P}^{n-1}$ correspond to 2-dimensional subspaces in $\mathbb{C}^n$. Therefore, we work over the Grassmannian $X = \mathrm{Gr}(2, \mathbb{C}^n)$, which is a smooth variety of dimension $2(n-2) = 2n-4$. Let $E_1, \dots, E_{2n-4}$ be general $(n-2)$-dimensional subspaces of $\mathbb{C}^n$. We count the number of points $E \in X$ satisfying

$$\dim(E \cap E_k) \geq 1 \quad \text{for } k = 1, \dots, 2n-4.$$

Let $\Omega_k \subseteq X$ be the Schubert variety of lines meeting E_k. Fixing a standard basis $e_1, \dots, e_n$, we can represent the first condition by setting $E_1 = \mathrm{span}(e_1, \dots, e_{n-2})$. The condition that E meets E_1 means the natural projection map from the tautological subbundle $\mathbb{S}$ to the quotient $\mathbb{C}^n/E_1$ fails to be injective. Therefore, Ω_1 is the zero locus of a section of the line bundle $\mathrm{Hom}\left(\bigwedge^2 \mathbb{S}, \bigwedge^2(\mathbb{C}^n/E_1)\right) \cong \bigwedge^2 \mathbb{S}^\vee \otimes \det(\mathbb{C}^n/E_1)$. The non-equivariant class of this Schubert variety is simply the first Chern class of the dual tautological bundle, $[\Omega_k] = c_1(\mathbb{S}^\vee)$, often denoted as the hyperplane class σ_1.

We introduce the torus action on X with weights $\chi_1, \dots, \chi_n$. The fixed points correspond to the coordinate planes $E_I = \mathrm{span}(e_a, e_b)$, indexed by 2-element subsets $I = \{a, b\} \subseteq \{1, \dots, n\}$. The tangent space at a fixed point E_I is

$$T_{[E_I]}X = \mathrm{Hom}(E_I, \mathbb{C}^n/E_I) \cong \bigoplus_{i \in I,\ j \notin I} L_i^\vee \otimes L_j.$$

Thus, the equivariant top Chern class of the tangent space is

$$c_{\mathrm{top}}^T(T_{[E_I]}X) = \prod_{i \in I,\ j \notin I} (\chi_j - \chi_i).$$

The number of intersecting lines is given by the integral

$$\int_X [\Omega_1] \cdots [\Omega_{2n-4}] = \int_X \sigma_1^{2n-4}.$$

Since this top-degree integral evaluates to a scalar, it can be computed using any equivariant lift. We choose $c_1^T(\mathbb{S}^\vee)$, whose restriction to the fixed point E_I is $(-\chi_a - \chi_b)$.

By Corollary 3.1.9, summing over the $\binom{n}{2}$ fixed points yields:

$$\int_X \sigma_1^{2n-4} = \sum_{1 \leq a < b \leq n} \frac{(-\chi_a - \chi_b)^{2n-4}}{\prod_{j \neq a,b}(\chi_j - \chi_a)(\chi_j - \chi_b)}.$$

This algebraic sum evaluates to the Catalan number:

$$C_{n-2} = \frac{1}{n-1}\binom{2n-4}{n-2}.$$

For $n = 4$, which corresponds to the number of lines meeting 4 general lines in $\mathbb{P}^3$, the formula recovers $C_2 = 2$.

3.2. Virtual Fundamental Classes and Virtual Localization

This section surveys the construction of virtual fundamental classes for moduli problems, for example, the existence of $[\overline{\mathcal{M}}_{g,n}(X,\beta)]^{\mathrm{vir}}$ posed at the beginning of this chapter. We then generalize the Integration Formula (Corollary 3.1.9) to the virtual setting, establishing the primary computational tool for these classes. The main references for this section are [**BF97, GP99**].

3.2.1. Virtual Fundamental Classes. We motivate the general construction using Fulton's intersection theory. Let X be a pure-dimensional scheme admitting a global regular embedding $i\colon X \hookrightarrow Y$ into a scheme Y. Consider the sequence of maps:

$$\mathrm{CH}_*(X) \xrightarrow{i_*} \mathrm{CH}_*(Y) \xrightarrow{\sigma} \mathrm{CH}_*(C_{X/Y}) \xrightarrow{\cong} \mathrm{CH}_*(N_{X/Y}) \xrightarrow{0^!} \mathrm{CH}_*(X).$$

Here, σ is the specialization map to the normal cone, and $0^!$ is the inverse of the flat pullback $\pi^*\colon \mathrm{CH}_*(X) \to \mathrm{CH}_*(N_{X/Y})$. This pullback is an isomorphism because $N_{X/Y}$ is a vector bundle for a regular embedding [**Ful98**, Theorem 3.3]. The operation $0^!$ represents intersection with the zero section of the normal bundle. The composition of these maps sends the fundamental class $[X]$ identically back to $[X]$.

If X is singular or not equidimensional, the classical fundamental class is ill-behaved since it lacks a homogeneous fundamental class in a single degree. However, given a global embedding into a scheme Y, the algebraic normal cone $C_{X/Y}$ remains well-defined and is equi-dimensional as long as Y is [**Ful98**, Appendix B.6.6]. If there exists a vector bundle $E_{X/Y}$ admitting a closed embedding of cones $j\colon C_{X/Y} \hookrightarrow E_{X/Y}$, we evaluate the sequence:

$$\mathrm{CH}_*(X) \xrightarrow{i_*} \mathrm{CH}_*(Y) \xrightarrow{\sigma} \mathrm{CH}_*(C_{X/Y}) \xrightarrow{j_*} \mathrm{CH}_*(E_{X/Y}) \xrightarrow{0^!_E} \mathrm{CH}_*(X).$$

This mechanism provides the first potential definition of the virtual fundamental class:

$$[X]^{\mathrm{vir}}_{E_{X/Y}} := 0^!_E[C_{X/Y}]. \tag{3.2.1}$$

The notation $[X]^{\mathrm{vir}}_{E_{X/Y}}$ emphasizes its dependence on the choice of the global embedding $X \hookrightarrow Y$ and the vector bundle $E_{X/Y}$. We refer to $E_{X/Y}$ as the obstruction bundle, a terminology we will justify shortly.

This construction closely parallels Fulton's general intersection product of X and V within an ambient scheme Y. Consider the Cartesian diagram:

$$\begin{tikzcd} W \arrow[r] \arrow[d, "g"'] \arrow[dr, phantom, "\square"] & V \arrow[d, "f"] \\ X \arrow[r, hook, "i"'] & Y \end{tikzcd}$$

When i is a regular embedding, the refined intersection product $X \cdot V$ is defined via the natural closed immersion $C_{W/V} \hookrightarrow g^* N_{X/Y}$. The product is defined by intersecting this cone with the zero section of the pullback normal bundle. The virtual fundamental class mirrors this mechanism, substituting the normal bundle with the obstruction bundle $E_{X/Y}$.

Example 3.2.1. Suppose X admits a global embedding into a pure-dimensional scheme Y. Assume its ideal sheaf $\mathcal{I}_X$ is generated by global sections $f_1, \ldots, f_n$ corresponding to effective Cartier divisors $D_1, \ldots, D_n$. This defines a canonical surjection

$$\bigoplus_{i=1}^{n} \mathcal{O}_Y(-D_i) \xrightarrow{D_i \mapsto f_i} \mathcal{I}_X.$$

Taking the relative spectrum of the symmetric algebras induces a closed immersion of the normal cone $C_{X/Y}$ into the vector bundle $E_{X/Y} \coloneqq \bigoplus_{i=1}^n \mathcal{O}_Y(D_i)|_X$. The construction of this obstruction bundle $E_{X/Y}$ explicitly depends on the choice of the generating sections $f_1, \dots, f_n$.

This setup is optimal for cases where $X \subseteq Y$ is defined by a section intersecting the zero section non-transversally. More generally, the relationship between the ambient space, the subscheme, and the vector bundle is captured by the following commutative diagram, which characterizes a broader model for the obstruction bundle:

$$\begin{tikzcd} & E \arrow[d] \\ X = Z(s) \arrow[r, hook] & Y \arrow[u, bend right, "s"'] \end{tikzcd}$$

Since X is the zero locus of the global section s, the section induces a canonical cone embedding $C_{X/Y} \hookrightarrow E|_X$. By setting the obstruction bundle to $E_{X/Y} = E|_X$, the virtual fundamental class evaluates to $[X]^{\mathrm{vir}}_E = 0^!_{E|_X}[C_{X/Y}]$.

We demonstrate this model with a fundamental example of non-transversal intersection, following [**Ful98**, Examples 4.2.2 and 6.1.4]. Let Y be a smooth surface and let A, B, and D be effective Cartier divisors on Y, where A and B intersect transversally at a point $P \notin D$. Define $A' \coloneqq A + D$ and $B' \coloneqq B + D$, and consider the scheme-theoretic intersection $X \coloneqq A' \cap B'$.

The subscheme X admits a global embedding into Y. Since A' and B' correspond to the generators of the ideal sheaf $\mathcal{I}_X$, the associated obstruction bundle is

$$E_{X/Y} = \left(\mathcal{O}_Y(A') \oplus \mathcal{O}_Y(B')\right)\big|_X .$$

To calculate the virtual fundamental class $[X]^{\mathrm{vir}}_{E_{X/Y}} = 0^!_E[C_{X/Y}]$, we apply the excess intersection formula [**Ful98**, Proposition 6.1(a)]:

$$[X]^{\mathrm{vir}}_{E_{X/Y}} = \Big\{c(E_{X/Y}) \cap s(X,Y)\Big\}_0 ,$$

where $s(X,Y)$ is the Segre class of X in Y. As computed in [**Ful98**, Example 4.2.2] via the blow-up of Y at P, the Segre class is

$$s(X,Y) = [D] + ([P] - D \cdot [D]) .$$

The first Chern class of the obstruction bundle is $c_1(E_{X/Y}) = A' + B' = A + B + 2D$. Substituting this and the Segre class into the excess intersection formula evaluates the virtual fundamental class:

$$[X]^{\mathrm{vir}}_{E_{X/Y}} = (A + B + 2D) \cdot [D] + ([P] - D \cdot [D]) = (A + B + D) \cdot [D] + [P].$$

We now address the limitations of the naive virtual fundamental class definition (3.2.1). Requiring X to admit a global embedding into a pure-dimensional or smooth scheme Y alongside a global cone immersion $C_{X/Y} \hookrightarrow E_{X/Y}$ is often overly restrictive. However, this is not a fatal drawback since we can circumvent this constraint by gluing local obstruction data to construct a global obstruction bundle.

Let $\{U_i\}$ be an open cover of X. Suppose each U_i admits a closed embedding $U_i \hookrightarrow M_i$ into a smooth scheme M_i, equipped with an obstruction bundle embedding $C_{U_i/M_i} \hookrightarrow E_{U_i/M_i}$. To glue these into a global structure, we impose compatibility conditions on both the ambient spaces and the obstruction bundles.

For a morphism of open subsets $\varphi_U \colon U_i \to U_j$, we first assume that there exists a corresponding smooth morphism $\varphi_M \colon M_i \to M_j$ such that the following diagram is Cartesian:

$$\begin{tikzcd} U_i \arrow[r, hook] \arrow[d, "\varphi_U"'] \arrow[dr, phantom, "\square"] & M_i \arrow[d, "\varphi_M"] \\ U_j \arrow[r, hook] & M_j \end{tikzcd}$$

The Cartesian and smoothness conditions guarantee that the pullback of the normal cone is isomorphic to the local normal cone (i.e., $\varphi_U^* C_{U_j/M_j} \cong C_{U_i/M_i}$), which allows us to glue the local normal cones C_{U_i/M_i} into a well-defined global one. We call the data $(U_i, M_i, \varphi_U, \varphi_M)$ a system of (Zariski) local embedding of X.

We subsequently enforce gluing conditions on the obstruction bundles. Let $E_i \coloneqq E_{U_i/M_i}$. On overlaps $U_{ij} = U_i \cap U_j$, we assume that there exist isomorphisms $\phi_{ij} \colon E_i|_{U_{ij}} \xrightarrow{\sim} E_j|_{U_{ij}}$ satisfying cocycle conditions, such that the following diagram commutes.

$$\begin{tikzcd} C_{U_i/M_i}|_{U_{ij}} \arrow[r, hook] \arrow[d, "\cong"'] & E_i|_{U_{ij}} \arrow[d, "\phi_{ij}"] \\ C_{U_j/M_j}|_{U_{ij}} \arrow[r, hook] & E_j|_{U_{ij}} \end{tikzcd}$$

These local compatibility criteria generalize the global construction and serve as the prototypical model for (perfect) obstruction theories. The mechanism developed above allows us to define the virtual fundamental class for a scheme admitting a system of local embedding.

The second, more severe drawback is the arbitrariness of the obstruction bundle in definition (3.2.1). Merely requiring a local closed immersion $C_{X/Y} \hookrightarrow E_{X/Y}$ permits choices that yield virtual fundamental classes of the incorrect expected dimension, violating the foundational principles proposed at the beginning of this chapter. The following example of twisted cubic illustrates this phenomenon.

Example 3.2.2 (Twisted Cubic in $\mathbb{P}^3$)**.** Consider the twisted cubic $X \subset \mathbb{P}^3$ defined by the 3-uple embedding $\mathbb{P}^1 \hookrightarrow \mathbb{P}^3$ via $[s,t] \mapsto [s^3, s^2t, st^2, t^3]$. Its ideal is generated by three quadrics $(xz - y^2, yw - z^2, xw - yz)$. Following the construction in Example 3.2.1, these generators correspond to global sections of $\mathcal{O}_{\mathbb{P}^3}(2)$. Since the embedding is of degree 3, the proposed obstruction bundle is

$$E_{X/\mathbb{P}^3} = \mathcal{O}_{\mathbb{P}^1}(6)^{\oplus 3}.$$

On the other hand, it is a classical result [**EVdV81**] that the normal bundle of the twisted cubic in $\mathbb{P}^3$ has rank 2 and decomposes as $N_{X/\mathbb{P}^3} = \mathcal{O}_{\mathbb{P}^1}(5) \oplus \mathcal{O}_{\mathbb{P}^1}(5)$. The embedding of the normal cone thus corresponds to an exact sequence of vector bundles:

$$0 \to N_{X/\mathbb{P}^3} \to E_{X/\mathbb{P}^3} \to Q \to 0,$$

where the quotient bundle Q is a line bundle.

By Fulton's excess intersection formula [**Ful98**, Theorem 6.3], the virtual fundamental class is computed by capping the classical fundamental class $[X]$ with the top Chern class of the quotient bundle Q:

$$[X]^{\mathrm{vir}}_{E_{X/\mathbb{P}^3}} = c_1(Q) \cap [X].$$

Therefore

$$[X]^{\mathrm{vir}}_{E_{X/\mathbb{P}^3}} = 8[\mathrm{pt}].$$

To resolve this limitation, a more intrinsic definition of the obstruction bundle $E_{X/Y}$ is required. Recall that when calculating the expected dimension of $\overline{\mathcal{M}}_{g,n}(X, \beta)$ (3.0.1), the deformation theory of stable maps fully controlled the lower bound of the dimension via the Grothendieck–Riemann–Roch theorem. This indicates a natural link between the choice of obstruction bundles and the deformation theory of the scheme X, which is entirely governed by its cotangent complex $\mathbb{L}_X$. Furthermore, the obstruction bundle construction relies on a closed immersion $C_{X/Y} \hookrightarrow E_{X/Y}$, and the normal sheaf of a scheme natively emerges from the lowest cohomological degree of the dual cotangent complex. These geometric facts motivate the use of the cotangent complex to extract the necessary intrinsic data for a "finer" obstruction theory.

Before proceeding, we collect the essential properties of cotangent complexes. We omit the technical construction of these complexes; full proofs are available in [**Ill71, Ill72**], and a modern, concise survey is provided in [**Alp26**, Appendix C].

Proposition 3.2.3 (Properties of Cotangent Complexes)**.** *Let $f\colon X \to Y$ be a morphism of schemes (resp. a morphism of finite type between Noetherian schemes). There exists a complex of flat $\mathcal{O}_X$-modules concentrated in degrees ≤ 0 with quasi-coherent (resp. coherent) cohomology. We denote its image in the upper-bounded derived category $\mathsf{D}^-(\mathrm{Qcoh}(X))$ (resp. $\mathsf{D}^-(X)$) by $\mathbb{L}_{X/Y}$. It satisfies the following properties:*

(1) The zeroth cohomology sheaf recovers the module of Kähler differentials: $\mathcal{H}^0(\mathbb{L}_{X/Y}) \cong \Omega_{X/Y}$.
(2) The morphism f is smooth if and only if f is locally of finite presentation and $\mathbb{L}_{X/Y}$ is a perfect complex concentrated in degree 0.
(3) If f is flat and locally of finite presentation, then f is a flat local complete intersection if and only if $\mathbb{L}_{X/Y}$ is a perfect complex concentrated in degrees $[-1, 0]$. Specifically, if f factors as a closed immersion $X \hookrightarrow \widetilde{X}$ into a scheme smooth over Y, then $\mathbb{L}_{X/Y}$ is isomorphic in the derived category to the complex

$$0 \to I/I^2 \to \Omega_{\widetilde{X}/Y}|_X \to 0,$$

where $\Omega_{\widetilde{X}/Y}|_X$ is placed in degree 0 and I is the ideal sheaf of X in $\widetilde{X}$.
(4) For a Cartesian square

$$\begin{tikzcd} X' \arrow[r, "g'"] \arrow[d] \arrow[dr, phantom, "\square"] & X \arrow[d, "f"] \\ Y' \arrow[r, "g"] & Y \end{tikzcd}$$

if either f or g is flat, there is a canonical isomorphism in the derived category: $Lg'^*\mathbb{L}_{X/Y} \xrightarrow{\sim} \mathbb{L}_{X'/Y'}$.
(5) The composition of scheme morphisms $X \xrightarrow{f} Y \to Z$ induces a distinguished triangle in the derived category:

$$Lf^*\mathbb{L}_{Y/Z} \to \mathbb{L}_{X/Z} \to \mathbb{L}_{X/Y} \to Lf^*\mathbb{L}_{Y/Z}[1].$$

We now generalize our previous constructions. We have discussed how to define the virtual fundamental class of a scheme admitting a system of local embedding, which allows us to glue the normal cones. Our next goal is to extend this to Deligne–Mumford stacks. By modding out the action of the tangent bundle, we generalize the gluing of scheme-theoretic normal cones to the gluing of quotient cone stacks. This mechanism incorporates intrinsic obstruction data, yielding a finer geometric construction known as the intrinsic normal cone.

Definition 3.2.4 (System of Local Embeddings)**.** Let $\mathcal{X}$ be a Deligne–Mumford stack. A system of local embeddings for $\mathcal{X}$ is the data $\{(U_i, M_i)\}_{i\in I}$ equipped with transition morphisms, where $\{U_i \to \mathcal{X}\}$ is an étale covering of $\mathcal{X}$, and each M_i is a smooth scheme admitting a closed immersion $f_i \colon U_i \to M_i$.

This data satisfies the compatibility condition that for any morphism $\varphi_U \colon U_i \to U_j$ over $\mathcal{X}$, there exists a smooth morphism $\varphi_M \colon M_i \to M_j$ such that the following diagram commutes:

$$\begin{tikzcd} U_i \arrow[r, "f_i"] \arrow[d, "\varphi_U"'] & M_i \arrow[d, "\varphi_M"] \\ U_j \arrow[r, "f_j"'] & M_j \end{tikzcd}$$

Let (U_i, M_i) be a local embedding in this system, and let $\mathcal{I}_i$ be the ideal sheaf of U_i in M_i. There is a canonical differential map $\mathcal{I}_i/\mathcal{I}_i^2 \to f_i^*\Omega_{M_i}$. Dualizing this yields a vector bundle morphism $f_i^*TM_i \to N_{U_i/M_i}$, which defines an action of the tangent bundle $f_i^*TM_i$ on the normal sheaf N_{U_i/M_i}. This naturally yields the quotient stack $[N_{U_i/M_i}/f_i^*TM_i]$.

Furthermore, the compatibility condition over a morphism $\varphi_U \colon U_i \to U_j$ yields a natural pullback map of ideal sheaves, inducing a commutative diagram:

$$\begin{tikzcd} \mathcal{I}_j/\mathcal{I}_j^2|_{U_i} \arrow[r] \arrow[d] & f_j^*\Omega_{M_j}|_{U_i} \arrow[d] \\ \mathcal{I}_i/\mathcal{I}_i^2 \arrow[r] & f_i^*\Omega_{M_i} \end{tikzcd}$$

Because we quotient by the respective tangent bundles, this diagram induces a canonical morphism of quotient stacks $\widetilde{\chi}^\vee \colon [N_{U_i/M_i}/f_i^*TM_i] \to [N_{U_j/M_j}/f_j^*TM_j]|_{U_i}$, which is an isomorphism.

The stack $[C_{U_i/M_i}/f_i^*TM_i]$ is a natural closed substack of $[N_{U_i/M_i}/f_i^*TM_i]$. It is proved in [**BF97**, Corollary 3.9] that the isomorphism $\widetilde{\chi}^\vee$ identifies the closed substack $[C_{U_i/M_i}/f_i^*TM_i]$ with the pullback $[C_{U_j/M_j}/f_j^*TM_j]|_{U_i}$. This descent condition ensures the pieces glue together, permitting the global definition of intrinsic normal cones and sheaves.

Definition 3.2.5 (Intrinsic Normal Cones and Sheaves)**.** Let $\mathcal{X}$ be a Deligne–Mumford stack equipped with a system of local embeddings $\{(U_i, M_i)\}_{i\in I}$.

(1) The intrinsic normal cone of $\mathcal{X}$ is the algebraic stack $\mathfrak{C}_\mathcal{X}$ such that for any local embedding (U_i, M_i) in the system, there is an isomorphism $\mathfrak{C}_\mathcal{X}|_{U_i} \cong [C_{U_i/M_i}/f_i^*TM_i]$.
(2) The intrinsic normal sheaf of $\mathcal{X}$ is the algebraic stack $\mathfrak{N}_\mathcal{X}$ such that for any local embedding (U_i, M_i) in the system, there is an isomorphism $\mathfrak{N}_\mathcal{X}|_{U_i} \cong [N_{U_i/M_i}/f_i^*TM_i]$.

This updated language incorporates the action of the tangent bundle on the normal cone and sheaf, inherently integrating the data of the cotangent complex. The connection of the intrinsic normal cones/sheaves and the cotangent complexes will become more explicit once we introduce the language of h^1/h^0-type stacks.

Definition 3.2.6 (h^1/h^0-type stacks)**.** Let $\mathcal{X}$ be a Deligne–Mumford stack and let $E^\bullet$ be a complex of abelian sheaves on $\mathcal{X}$ concentrated in degrees 0 and 1, with differential $d \colon E^0 \to E^1$. We define the quotient stack

$$h^1/h^0(E^\bullet) := [E^1/E^0].$$

For a complex $E^\bullet$ of arbitrary length, we define:

$$h^1/h^0(E^\bullet) := h^1/h^0(\tau_{[0,1]}(E^\bullet)),$$

where $\tau_{[0,1]}$ denotes the canonical truncation of $E^\bullet$.

A morphism $\phi\colon E^\bullet \to F^\bullet$ between two-term complexes induces a morphism of quotient stacks $h^1/h^0(\phi)\colon h^1/h^0(E^\bullet) \to h^1/h^0(F^\bullet)$. For an étale scheme $U \to \mathcal{X}$, an object $(P,f) \in h^1/h^0(E^\bullet)(U)$ consists of an E^0-torsor P and an E^0-equivariant map $f\colon P \to E^1$. The morphism $h^1/h^0(\phi)$ maps (P,f) to the pushout $(P \times^{E^0} F^0, \phi^1(f))$, where the map

$$\begin{aligned} \phi^1(f)\colon P \times^{E^0} F^0 &\longrightarrow F^1 \\ (p,\mu) &\longmapsto \phi^1(f(p)) + d(\mu) \end{aligned}$$

is well-defined via the differential $d\colon F^0 \to F^1$ of $F^\bullet$.

We collect the foundational properties of these stacks [**BF97**, Section 2].

Proposition 3.2.7 (Properties of h^1/h^0-type stacks)**.** *Let $E^\bullet$ and $F^\bullet$ be complexes of abelian sheaves on a Deligne–Mumford stack $\mathcal{X}$. A complex $L^\bullet \in \mathsf{D}(\mathcal{O}_\mathcal{X})$ satisfies condition (*) if $\mathcal{H}^i(L^\bullet) = 0$ for $i > 0$ and $\mathcal{H}^i(L^\bullet)$ is coherent for $i = 0, -1$.*

(1) A quasi-isomorphism $\phi\colon E^\bullet \to F^\bullet$ induces an isomorphism of quotient stacks $h^1/h^0(\phi)$. Consequently, the h^1/h^0-type stack is well-defined up to unique isomorphism for objects in the derived category $\mathsf{D}(\mathcal{O}_\mathcal{X})$.

(2) If $L^\bullet \in \mathsf{D}(\mathcal{O}_\mathcal{X})$ satisfies condition (*), the dual quotient stack $h^1/h^0((L^\bullet)^\vee)$ is an algebraic stack and an abelian cone stack. If $L^\bullet$ is perfect and concentrated in degrees $[-1,0]$, then $h^1/h^0((L^\bullet)^\vee)$ is a vector bundle stack. (We implicitly evaluate the dual complex in the smooth or big fppf topology).

(3) Let $\phi\colon E^\bullet \to F^\bullet$ be a morphism in $\mathsf{D}(\mathcal{O}_\mathcal{X})$ where both complexes satisfy condition (*). The induced stack morphism $h^1/h^0(\phi^\vee)\colon h^1/h^0((F^\bullet)^\vee) \to h^1/h^0((E^\bullet)^\vee)$ is a closed immersion if and only if $\mathcal{H}^0(\phi)$ is an isomorphism and $\mathcal{H}^{-1}(\phi)$ is surjective. Furthermore, $h^1/h^0(\phi^\vee)$ is an isomorphism if and only if both $\mathcal{H}^0(\phi)$ and $\mathcal{H}^{-1}(\phi)$ are isomorphisms.

We can now give an intrinsic description of Definition 3.2.5. Let (U,M) be a local embedding in a system of local embeddings for a Deligne–Mumford stack $\mathcal{X}$. Because $U \to \mathcal{X}$ is étale, we have $\mathbb{L}_\mathcal{X}|_U \cong \mathbb{L}_U$. By Proposition 3.2.3, the relative cotangent complex $\mathbb{L}_{U/M}$ is isomorphic in the derived category to the complex $[I/I^2 \to f^*\Omega_M]$ concentrated in degrees $[-1,0]$.

The natural morphism of cotangent complexes $\phi\colon \mathbb{L}_\mathcal{X}|_U \to \mathbb{L}_{U/M}$ induces an isomorphism on the cohomology sheaves $\mathcal{H}^{-1}$ and $\mathcal{H}^0$. Therefore, by Proposition 3.2.7, the dual morphism $\phi^\vee$ yields a canonical isomorphism of quotient stacks:

$$[N_{U/M}/f^*TM] \xrightarrow{\sim} h^1/h^0(\mathbb{L}_\mathcal{X}^\vee)\Big|_U .$$

This glues globally to establish the canonical isomorphism:

$$h^1/h^0(\mathbb{L}_\mathcal{X}^\vee) \cong \mathfrak{N}_\mathcal{X}.$$

Consequently, the intrinsic normal cone $\mathfrak{C}_\mathcal{X}$ is identified exactly as a closed substack of $h^1/h^0(\mathbb{L}_\mathcal{X}^\vee)$.

This proves that the intrinsic normal sheaf is entirely controlled by the cotangent complex of $\mathcal{X}$. On the other hand, we need to utilize the full information of the cotangent complex to capture all deformation obstructions and yield the correct expected dimension, which motivates the following definition:

Definition 3.2.8 (Perfect Obstruction Theory and Virtual Fundamental Class)**.** Let $\mathcal{X}$ be a Deligne–Mumford stack, and let $E^\bullet \in \mathsf{D}(\mathcal{O}_\mathcal{X})$ be a complex satisfying condition (*) from Proposition 3.2.7.

(1) A morphism $\phi\colon E^\bullet \to \mathbb{L}_\mathcal{X}$ is an *obstruction theory* if $\mathcal{H}^0(\phi)$ is an isomorphism and $\mathcal{H}^{-1}(\phi)$ is surjective. By Proposition 3.2.7, this is equivalent to the condition that the induced stack morphism $\phi^\vee\colon \mathfrak{N}_\mathcal{X} \to h^1/h^0(E^\vee)$ is a closed immersion. If $E^\bullet$ is additionally a perfect complex concentrated in degrees $[-1,0]$, then ϕ is called a *perfect obstruction theory*.

(2) Let $\phi\colon E^\bullet \to \mathbb{L}_\mathcal{X}$ be a perfect obstruction theory. The quotient stack can be represented as $h^1/h^0(E^\vee) \cong [(E^{-1})^\vee/(E^0)^\vee]$. The vector bundle $(E^{-1})^\vee$ provides an atlas for this stack, inducing the following Cartesian diagram:

$$\begin{tikzcd} C(E^\bullet) \arrow[r] \arrow[d] \arrow[dr, phantom, "\square"] & (E^{-1})^\vee \arrow[d] \\ \mathfrak{C}_\mathcal{X} \arrow[r] & h^1/h^0(E^\vee) \end{tikzcd}$$

The *virtual fundamental class* is defined as the intersection of the cone pullback $C(E^\bullet)$ with the zero section of the vector bundle $(E^{-1})^\vee$:

$$[\mathcal{X}]^{\mathrm{vir}}_{E^\bullet} \coloneqq 0^!_{(E^{-1})^\vee}[C(E^\bullet)].$$

Remark 3.2.9. In practice, it suffices to assume that $E^\bullet$ is locally isomorphic to a perfect complex in degrees $[-1,0]$ and admits a global resolution by vector bundles. The resulting virtual fundamental class is independent of the chosen global resolution [**BF97**, Proposition 5.3].

The requirement for a global resolution can be dropped, provided one extends intersection theory to general Artin stacks. The quotient stacks of h^1/h^0-type are algebraic, but not generally Deligne–Mumford. The development of this generalized intersection theory is beyond the scope of this thesis; we refer to [**Kre99**] for the construction.

Example 3.2.10 (Perfect obstruction theory on schemes)**.** We revisit the construction of the virtual fundamental class for an arbitrary scheme, addressing the core problem posed at the beginning of this section. We expand upon the zero-locus model from Example 3.2.1. Let Y be a smooth scheme, $E \to Y$ a vector bundle, and $s\colon Y \to E$ a global section. Let $X \subseteq Y$ be the subscheme defined as the zero locus of s:

$$\begin{tikzcd} & E \arrow[d] \\ X = Z(s) \arrow[r, hook] & Y \arrow[u, bend right, "s"'] \end{tikzcd}$$

We illustrate how this model induces a perfect obstruction theory on X. The section s is equivalently a bundle morphism $s^\vee\colon E^\vee \to \mathcal{O}_Y$. The image of $s^\vee$ generates the ideal sheaf $\mathcal{I}$ defining X. Restricting to X yields a canonical surjection $E^\vee|_X \to \mathcal{I}/\mathcal{I}^2$.

Composing this with the differential map $\mathcal{I}/\mathcal{I}^2 \to \Omega_Y|_X$ produces the following commutative diagram:

$$\begin{tikzcd} E^\vee|_X \arrow[r] \arrow[d, "s^\vee|_X"'] & \Omega_Y|_X \arrow[d, "\mathrm{id}"] \\ \mathcal{I}/\mathcal{I}^2 \arrow[r, "d"'] & \Omega_Y|_X \end{tikzcd} \tag{3.2.2}$$

We define the two-term complex $E^\bullet$ concentrated in degrees $[-1,0]$ as:

$$E^\bullet \coloneqq [E^\vee|_X \to \Omega_Y|_X].$$

The commutative diagram above induces a morphism in the derived category $\phi\colon E^\bullet \to \mathbb{L}_X$. Because $\mathcal{H}^0(\phi)$ is an isomorphism and $\mathcal{H}^{-1}(\phi)$ is surjective, ϕ constitutes a perfect obstruction theory on X.

This construction yields the desired virtual fundamental class $[X]^{\mathrm{vir}}_{E^\bullet}$. According to Definition 3.2.8, this class is obtained by the refined intersection product of the graph of the section s with the zero section of E. In the literature, this localized top Chern class is referred to as the *refined Euler class* of E, as detailed in [**GP99**, Section 2]. See also [**Ful98**, Chapter 14.1] for detailed discussion of localized Chern classes.

Furthermore, it is a foundational result that *any* perfect obstruction theory on a scheme is locally of the form (3.2.2), as demonstrated in [**Tod21**, pp. 6–7].

Overall, the preceding discussion reveals the general principle underlying virtual fundamental classes: defining the virtual fundamental class of a moduli space—such as the moduli stack of stable maps $\overline{\mathcal{M}}_{g,n}(X,\beta)$ introduced at the beginning of this chapter—ultimately reduces to finding a well-behaved two-term obstruction theory. For a detailed discussion of the canonical obstruction theory on $\overline{\mathcal{M}}_{g,n}(X,\beta)$, we refer to [**GP99**, Appendix B].

3.2.2. Virtual Localization. We now generalize the Integration Formula (Corollary 3.1.9) to the virtual setting. Recall that for a smooth variety X equipped with a $\mathbb{C}^*$-action, the classical localization formula states

$$[X] = \sum_{Z\subseteq X^{\mathbb{C}^*}} (\iota_Z)_* \left(\frac{[Z]}{c^{\mathbb{C}^*}_{\mathrm{top}}(N_{Z/X})}\right),$$

where $[X]$ and $[Z]$ are the equivariant fundamental classes of X and a component of the fixed locus $X^{\mathbb{C}^*}$, respectively. We follow the standard convention of denoting the top equivariant Chern class as the equivariant Euler class $e(N)$.

Throughout this subsection, all classes (including virtual fundamental classes) are $\mathbb{C}^*$-equivariant, though we suppress the torus action subscript $\mathbb{C}^*$ to simplify the notation.

Let X be an algebraic scheme over $\mathbb{C}$ equipped with a torus action and a $\mathbb{C}^*$-equivariant perfect obstruction theory. Although the classical localization formula fails for singular X, the virtual fundamental class machinery extends it to

$$[X]^{\mathrm{vir}} = \sum_{X_i\subseteq X^{\mathbb{C}^*}} \iota_* \left(\frac{[X_i]^{\mathrm{vir}}}{e(N_i^{\mathrm{vir}})}\right).$$

Here, X_i are the components of the fixed locus $X^{\mathbb{C}^*}$ equipped with an induced perfect obstruction theory, and $e(N_i^{\mathrm{vir}})$ is the Euler class of the virtual normal bundle. To make this formula rigorous, we need to resolve two questions:

(1) Given a $\mathbb{C}^*$-equivariant perfect obstruction theory $E^\bullet$ on X, how does it induce a perfect obstruction theory on X_i to define $[X_i]^{\mathrm{vir}}$?
(2) What is the definition of the virtual normal bundle N_i^{vir}, and how is its Euler class defined?

To address the first question, we assume X admits a $\mathbb{C}^*$-equivariant closed immersion into a smooth scheme (or Deligne–Mumford stack) Y. Let $\mathcal{I}$ be the ideal sheaf of X in Y. The canonical truncation of the cotangent complex of X is quasi-isomorphic to the two-term complex:

$$\tau_{\geq -1}\mathbb{L}_X \cong [\mathcal{I}/\mathcal{I}^2|_X \to \Omega_Y|_X].$$

A perfect obstruction theory is given by a two-term complex of vector bundles $E^\bullet = [E^{-1} \to E^0]$ on X and a morphism $E^\bullet \to \mathbb{L}_X$ that induces an isomorphism on $\mathcal{H}^0$ and a surjection on $\mathcal{H}^{-1}$. If this morphism is $\mathbb{C}^*$-equivariant, $E^\bullet$ constitutes a $\mathbb{C}^*$-equivariant perfect obstruction theory.

By Lemma 3.1.8, the components Y_i of the fixed locus $Y^{\mathbb{C}^*}$ are smooth. For any $\mathbb{C}^*$-invariant closed subscheme $X \subseteq Y$, we have $X^T = X \cap Y^T$ scheme-theoretically. Let $X_i = X \cap Y_i$ denote the components of the fixed locus $X^{\mathbb{C}^*}$, which may be singular or reducible.

For a $\mathbb{C}^*$-equivariant coherent sheaf $\mathcal{F}$ on a fixed component X_i, the representation theory of the torus dictates that $\mathcal{F}$ decomposes into a direct sum of weight spaces

$$\mathcal{F} = \bigoplus_{\chi \in \mathrm{Hom}(\mathbb{C}^*, \mathbb{G}_m)} \mathcal{F}^\chi \cong \bigoplus_{k \in \mathbb{Z}} \mathcal{F}^k.$$

Following the notation of [**GP99**, p. 6], we define the fixed part of $\mathcal{F}$ as the weight-zero component $\mathcal{F}^f := \mathcal{F}^0$, and the moving part as $\mathcal{F}^m := \bigoplus_{\chi \neq 0} \mathcal{F}^\chi$. This yields a canonical splitting $\mathcal{F} = \mathcal{F}^f \oplus \mathcal{F}^m$.

Restricting the complex $E^\bullet$ to X_i, we obtain a complex $E_i^\bullet$ with a fixed part $E_i^{\bullet,f}$. The original obstruction morphism yields a pullback map $\phi_i \colon E_i^\bullet \to \mathbb{L}_X|_{X_i}$. Composing this with the natural restriction $\mathbb{L}_X|_{X_i} \to \mathbb{L}_{X_i}$ yields a canonical map $E_i^\bullet \to \mathbb{L}_{X_i}$.

Because X_i is fixed pointwise by $\mathbb{C}^*$, the torus acts trivially on $\mathbb{L}_{X_i}$, meaning $\mathbb{L}_{X_i}^f = \mathbb{L}_{X_i}$. Furthermore, from the scheme-theoretic intersection $X_i = X \cap Y_i$, we have $\Omega_Y|_{Y_i}^f \cong \Omega_{Y_i}$, which restricts to $\Omega_X|_{X_i}^f \cong \Omega_{X_i}$. Taking the fixed part of the composite morphism naturally isolates a map:

$$\psi_i \colon E_i^{\bullet,f} \to \mathbb{L}_{X_i}.$$

By [**GP99**, Proposition 1], the morphism ψ_i inherently constitutes a perfect obstruction theory on X_i. This answers the first question: the canonical fixed-part perfect obstruction theory directly induces the virtual fundamental class $[X_i]^{\mathrm{vir}}$.

We next define the virtual normal bundle N_i^{vir} by dualizing the obstruction theory. Let $E_\bullet := (E^\bullet)^\vee$ denote the virtual tangent complex, which is concentrated in degrees $[0,1]$. Restricting this to the fixed component X_i yields the complex $E_{\bullet,i}$. The virtual normal bundle N_i^{vir} is defined as the moving part of this restricted complex, denoted by $E_{\bullet,i}^m$. In some literature, the virtual normal bundle is also defined as $[B_0] - [B_1]$ K-theoretically, where $[B_0 \to B_1]$ represents the two-term complex $E_{\bullet,i}^m$.

When X is smooth, this recovers the classical normal bundle. In the smooth case, the canonical perfect obstruction theory is simply the cotangent bundle $E^\bullet = \Omega_X$ placed in degree 0 (Proposition 3.2.3). Its dual complex is the tangent bundle $E_\bullet = T_X$. Restricting to the fixed locus X_i, the torus action induces a direct sum decomposition $T_X|_{X_i} = (T_X|_{X_i})^f \oplus (T_X|_{X_i})^m$. The fixed part is exactly the tangent bundle of the fixed locus T_{X_i}, and the moving part is precisely the classical normal bundle $N_{X_i/X}$. Therefore, $N_i^{\mathrm{vir}} = (T_X|_{X_i})^m = N_{X_i/X}$, matching classical normal bundle.

The equivariant Euler class (or top equivariant Chern class) of the virtual normal bundle is defined as follows. Let the moving complex $E_{\bullet,i}^m$ be represented by a two-term complex of vector bundles $[B_0 \to B_1]$, where B_k sits in degree k. The virtual Euler class is the ratio of their classical equivariant Euler classes:

$$e(N_i^{\mathrm{vir}}) = e([B_0 \to B_1]) := \frac{e(B_0)}{e(B_1)}.$$

Since B_0 and B_1 consist entirely of moving weight spaces, their equivariant Euler classes are products of non-zero characters. Therefore, these classes are invertible in the localized equivariant Chow ring in Theorem 3.1.7.

The coefficient ring $\Lambda_{\mathbb{C}^*}$ in the $\mathbb{Q}$-coefficient is $\Lambda_{\mathbb{C}^*} \cong \mathbb{Q}[t]$. Therefore, the localized ring used in the Localization Formula is explicitly:

$$\mathrm{CH}_*^{\mathbb{C}^*}(X) \otimes_{\mathbb{Q}[t]} \mathbb{Q}[t, t^{-1}].$$

In this localized ring, the division in the virtual Euler class formula is well-defined.

We are now prepared to state and prove the main theorem of virtual localization [**GP99**, Sections 2–3].

Theorem 3.2.11. *Let X be an algebraic scheme admitting a $\mathbb{C}^*$-action and a $\mathbb{C}^*$-equivariant perfect obstruction theory, with a $\mathbb{C}^*$-equivariant closed embedding into a nonsingular scheme (or a smooth Deligne–Mumford stack) Y. Then in the localized Chow ring $\mathrm{CH}_*^{\mathbb{C}^*}(X) \otimes_{\mathbb{Q}[t]} \mathbb{Q}[t, t^{-1}]$, we have:*

$$[X]^{\mathrm{vir}} = \sum_{X_i \subseteq X^{\mathbb{C}^*}} \iota_{i*} \left(\frac{[X_i]^{\mathrm{vir}}}{e(N_i^{\mathrm{vir}})} \right),$$

where $\iota_i \colon X_i \hookrightarrow X$ are the inclusion maps of the fixed components.

Proof. We focus on the global zero-locus case from Example 3.2.1. Suppose there is a $\mathbb{C}^*$-equivariant vector bundle $V \to Y$ with a $\mathbb{C}^*$-invariant section v, such that X is the zero locus of v. As shown previously, the two-term complex of bundles on X

$$E^\bullet := [V^\vee|_X \to \Omega_Y|_X]$$

constitutes a perfect obstruction theory on X, which is clearly $\mathbb{C}^*$-equivariant. The virtual fundamental class $[X]^{\mathrm{vir}}$ is precisely the refined Euler class $e_{\mathrm{ref}}(V)$ (Example 3.2.1).

Let Y_i be the components of the fixed locus $Y^{\mathbb{C}^*}$, so $X_i = X \cap Y_i$. Because v is $\mathbb{C}^*$-invariant, its restriction to Y_i lies entirely within the weight-zero subbundle V_i^f, where $V_i := V|_{Y_i}$. Therefore, the perfect obstruction theory induced on X_i is exactly the one obtained from the bundle V_i^f and the restricted section $v|_{Y_i}$. Consequently, the virtual fundamental class of each fixed component X_i is the refined Euler class of the fixed part: $[X_i]^{\mathrm{vir}} = e_{\mathrm{ref}}(V_i^f)$.

Next, we identify the virtual normal bundle, defined as the moving part of the complex dual to $E^\bullet$, which is $[TY|_X \to V|_X]$. Restricting to X_i, the moving part of the tangent bundle TY is simply the honest normal bundle $N_{Y_i/Y}|_{X_i}$. Thus, the virtual normal bundle is the complex

$$N_i^{\mathrm{vir}} = [N_{Y_i/Y}|_{X_i} \to V_i^m|_{X_i}].$$

By definition, its Euler class in the localized ring is

$$e(N_i^{\mathrm{vir}}) = \frac{e(N_{Y_i/Y})}{e(V_i^m)}.$$

Therefore, substituting these into the localization formula, we are left to show that the following holds in the localized Chow ring of X:

$$e_{\mathrm{ref}}(V) = \sum_i \iota_{i*} \frac{e_{\mathrm{ref}}(V_i^f) \cap e(V_i^m)}{e(N_{Y_i/Y})}. \tag{3.2.3}$$

Because Y is smooth, the classical Integration Formula (Corollary 3.1.9) applies to Y:

$$[Y] = \sum_i j_{i*} \frac{[Y_i]}{e(N_{Y_i/Y})},$$

where $j_i\colon Y_i \hookrightarrow Y$. Capping both sides with the refined Euler class $e_{\mathrm{ref}}(V)$ yields:

$$e_{\mathrm{ref}}(V) = \sum_i \iota_{i*} \frac{e_{\mathrm{ref}}(V_i) \cap [Y_i]}{e(N_{Y_i/Y})}.$$

Finally, observe that over Y_i, the bundle splits equivariantly as $V_i = V_i^f \oplus V_i^m$. Because the section $v|_{Y_i}$ is $\mathbb{C}^*$-invariant, its projection onto the moving part V_i^m is identically zero. Therefore, by [**Ful98**, Example 14.1.3] we have

$$e_{\mathrm{ref}}(V_i) = e_{\mathrm{ref}}(V_i^f) \cap e(V_i^m).$$

Substituting this factorization into the sum concludes (3.2.3) and completes the proof for the global zero-locus case.

The general case relies on reducing arbitrary perfect obstruction theories to this localized zero-locus model, which is technically heavier. We refer to [**GP99**, Section 3] for those full details. □

Remark 3.2.12. In applications, the virtual localization formula is usually applied to moduli spaces, which are typically Deligne–Mumford stacks. The foundational work of Graber and Pandharipande [**GP99**, Appendix C] also generalizes the virtual localization formula to the case where X is a Deligne–Mumford stack, under the condition that X admits a global $\mathbb{C}^*$-equivariant embedding into a smooth Deligne–Mumford stack. This global embedding requirement has since been proven redundant in subsequent literature [**CKL17**].

In practice, applying the virtual localization formula requires two main geometric identifications. First, one needs to identify the virtual fundamental class of the fixed locus, which naturally arises from the $\mathbb{C}^*$-fixed part of the moduli space's perfect obstruction theory. Second, one needs to identify the virtual normal bundle, whose Euler class is typically evaluated via the cohomology ring of the fixed locus.

The first example of this machinery is the moduli stack of stable maps, with detailed localization calculations provided in [**GP99**, Section 4]. Our derivation of the ELSV formula in the following chapter will rely on exactly this geometric philosophy, demonstrating these virtual techniques in an explicit manner.

CHAPTER 4

Proof of the ELSV Formula

In Chapter 2, we studied the moduli stack of curves, $\mathcal{M}_g$, which is a smooth Deligne–Mumford stack. We noted that, by Yoneda's lemma, understanding the classification problem of algebraic curves naturally reduces to understanding the cohomology ring (or the Chow ring) of the stack $\mathcal{M}_g$. Within this mysterious cohomology ring, the tautological classes (Definition 2.4.24) are the "easiest classes one can imagine" (Ravi Vakil). They serve as an algebraic analogue to the Chern classes of the universal subbundle $\mathcal{S}$ on the Grassmannian, freely generating the stable cohomology.

Historically, there was little reason to expect this tautological ring to be well-behaved until Faber proposed his remarkable conjecture [**Fab99b**] about its full structure. Recall that all Chow rings and cohomology rings discussed here are taken with rational coefficients.

Conjecture 4.0.1 (Faber's Conjecture)**.** For any $g \geq 2$, the tautological ring (Definition 2.4.23) $\mathrm{R}^*(\mathcal{M}_g)$ has the following structure:

(1) $\mathrm{R}^d(\mathcal{M}_g) = 0$ for $d \geq g-1$;
(2) $\mathrm{R}^{g-2}(\mathcal{M}_g) \cong \mathbb{Q}$;
(3) The multiplication pairing
$$\mathrm{R}^d(\mathcal{M}_g) \times \mathrm{R}^{g-2-d}(\mathcal{M}_g) \to \mathrm{R}^{g-2}(\mathcal{M}_g)$$
is perfect.

The first two parts (the vanishing and socle properties) of Conjecture 4.0.1 were proven by Looijenga [**Loo95**] and Faber [**Fab97**], respectively. The third part—predicting that the tautological ring is Gorenstein—was verified by Faber for $g \leq 23$ by constructing sufficiently many tautological relations, known as the Faber–Zagier relations. However, the conjecture remains open for $g \geq 24$, where the Faber–Zagier method fails to produce enough relations to guarantee a perfect pairing [**Pan18**].

One can also extend the perfect pairing conjecture to the Deligne–Mumford compactification $\overline{\mathcal{M}}_{g,n}$ for $2g-2+n > 0$. Unfortunately, this extended property is generally false. The tautological ring $\mathrm{R}^*(\overline{\mathcal{M}}_{2,n})$ is not Gorenstein for $n \geq 20$ ([**PT14, Pet16**]), and this was recently generalized by Canning [**Can25**], who showed that $\mathrm{R}^*(\overline{\mathcal{M}}_{g,n})$ is not Gorenstein whenever $g \geq 2$ and $2g+n \geq 24$.

Although the Gorenstein conjecture oversimplified the mysterious structure of tautological rings, it sparked a major wave of research into the intersection theory of the moduli stacks $\overline{\mathcal{M}}_{g,n}$ and $\mathcal{M}_{g,n}$. If the perfect pairing conjecture were true, the entire structure of the tautological ring would be recovered completely by the computation of top-level intersection numbers on $\overline{\mathcal{M}}_{g,n}$ (thanks to the vanishing and socle properties). This motivated mathematicians to study the following fundamental problem:

Question 4.0.2 (Hodge Integrals)**.** Compute the **Hodge integrals** (top intersection numbers):
$$\langle \tau_{a_1} \ldots \tau_{a_n} \lambda_{b_1} \ldots \lambda_{b_g} \rangle_g := \int_{[\overline{\mathcal{M}}_{g,n}]} \psi_1^{a_1} \cdots \psi_n^{a_n} \cdot \lambda_1^{b_1} \cdots \lambda_g^{b_g} \in \mathbb{Q},$$

where $\sum_{i=1}^{n} a_i + \sum_{k=1}^{g} k \cdot b_k = 3g - 3 + n$.

Remarkably, Faber [**Fab99a**] proposed a powerful algorithm that reduces the evaluation of these general Hodge integrals to the computation of purely ψ-class intersection numbers (i.e., the case where $b_1 = \ldots = b_g = 0$).

The evaluation of the integrals in Question 4.0.2 remained an open problem for a long time until Kontsevich's celebrated proof of Witten's conjecture, which provided a systematic method to compute all purely ψ-class intersection numbers.

Theorem 4.0.3 (Witten's Conjecture / Kontsevich's Theorem)**.** *Consider the generating function (also known as the free energy in physics) of the pure ψ-class intersection numbers $\langle \tau_{a_1} \cdots \tau_{a_n} \rangle_g$, defined as:*

$$F(\mathbf{t}) \coloneqq \sum_{g=0}^{\infty} \sum_{n=1}^{\infty} \frac{1}{n!} \sum_{a_1,\ldots,a_n \geq 0} \langle \tau_{a_1} \cdots \tau_{a_n} \rangle_g t_{a_1} \cdots t_{a_n}.$$

Let $Z(\mathbf{t}) \coloneqq \exp(F(\mathbf{t}))$ be the formal exponential of F, which is called the ***partition function*** *of two-dimensional topological quantum gravity. Then $Z(\mathbf{t})$ is a τ-function for the KdV hierarchy. Equivalently, the function $U \coloneqq \frac{\partial^2 F}{\partial t_0^2}$ satisfies the classical KdV equation*

$$\frac{\partial U}{\partial t_1} = U \frac{\partial U}{\partial t_0} + \frac{1}{12} \frac{\partial^3 U}{\partial t_0^3}.$$

This equation allows us to inductively solve for the coefficients of F, and hence compute all of these intersection numbers.

Meanwhile, a wave of systematic study regarding the general structure of Hurwitz numbers (Definition 1.1.4) arose in combinatorial theory. In this chapter, we fix the following notation for these special Hurwitz numbers. Let $\infty \in \mathbb{P}^1$ and let $\mu = (\mu_1, \ldots, \mu_n)$ be a partition of a positive integer d. We define $H_{g,\mu}$ to be the Hurwitz number counting (weighted by the automorphism factor $1/|\operatorname{Aut}(\mu)|$) all branched coverings $f \colon C \to \mathbb{P}^1$ from a connected smooth curve C of genus $g(C) = g$, such that f has ramification profile μ over ∞, and simple ramification over r other fixed branch points. Note that by the Riemann–Hurwitz formula, the parameters g, r, d, and n are constrained by

$$2g - 2 = -2d + (r + d - n) = r - d - n.$$

Our goal is to understand the general structure of $H_{g,\mu}$. For genus zero, Hurwitz himself [**Hur01**] proved that

$$H_{0,\mu} = \frac{r!}{|\operatorname{Aut}(\mu)|} d^{n-3} \prod_{i=1}^{n} \left(\frac{\mu_i^{\mu_i}}{\mu_i!} \right). \tag{4.0.1}$$

Goulden and Jackson [**GJ99**] later studied the higher-genus cases and conjectured a structural formula for $H_{g,\mu}$ in general. Their conjecture can be phrased as follows:

Conjecture 4.0.4. For each g and n, there is a symmetric polynomial $P_{g,n}$ in n variables, whose terms have specific homogeneous degrees, such that

$$H_{g,\mu} = \frac{r!}{|\operatorname{Aut}(\mu)|} \prod_{i=1}^{n} \left(\frac{\mu_i^{\mu_i}}{\mu_i!} \right) P_{g,n}(\mu_1, \ldots, \mu_n).$$

This brings us to a culmination point. The evaluation of Hodge integrals (Question 4.0.2) and the search for the general structure of $H_{g,\mu}$ (Conjecture 4.0.4) beautifully converge. The following theorem, known as the ELSV formula [**ELSV01**], establishes a remarkable bridge between combinatorics, Gromov–Witten theory, and theoretical physics.

Theorem 4.0.5 (Ekedahl–Lando–Shapiro–Vainshtein, 2001)**.** *Following the notation above, we have*

$$H_{g,\mu} = \frac{r!}{|\operatorname{Aut}(\mu)|}\prod_{i=1}^{n}\frac{\mu_i^{\mu_i}}{\mu_i!}\int_{[\overline{\mathcal{M}}_{g,n}]}\frac{\Lambda^\vee(1)}{\prod_{j=1}^{n}(1-\mu_j\psi_j)},\ \textit{when } 2g-2+n>0.$$

Here, $\Lambda^\vee(1) = \sum_{i=0}^{g}(-1)^i\lambda_i$ *is the alternating sum of the* λ*-classes, and the fraction in the integrand is formally expanded as*

$$\frac{1}{1-\mu_i\psi_i} = \sum_{t\geq 0}(\mu_i\psi_i)^t.$$

Before moving on, we first take a closer look at how powerful Theorem 4.0.5 truly is.

First, combinatorially trivial or easily derived relations among Hurwitz numbers can yield new identities and relations among Hodge integrals. For example, in the base case of $g=0$, we can apply the classical Hurwitz formula (4.0.1) to see that over $\overline{\mathcal{M}}_{0,n}$, the pure ψ-integrals satisfy

$$\int_{\overline{\mathcal{M}}_{0,n}}\psi_1^{a_1}\cdots\psi_n^{a_n} = \binom{n-3}{a_1,\ldots,a_n},$$

which recovers a classical theorem in intersection theory.

Conversely, well-known results in algebraic geometry can compute Hurwitz numbers that are incredibly difficult to derive manually via combinatorics or representation theory. To see this explicitly, consider the case $g=1$ and $n=1$, where the ramification profile over ∞ is a single partition $\mu=(d)$. By the Riemann–Hurwitz formula, the number of simple branch points is $r=d+1$. The ELSV formula yields

$$H_{1,(d)} = \frac{(d+1)!}{1}\frac{d^d}{d!}\int_{\overline{\mathcal{M}}_{1,1}}\frac{1-\lambda_1}{1-d\psi_1}.$$

Since the dimension of $\overline{\mathcal{M}}_{1,1}$ is 1, we only need to extract the degree-1 terms from the formal expansion of the integrand, which gives $d\psi_1-\lambda_1$. It is a classical result that $\int_{\overline{\mathcal{M}}_{1,1}}\psi_1 = \frac{1}{24}$ and $\int_{\overline{\mathcal{M}}_{1,1}}\lambda_1 = \frac{1}{24}$. Substituting these values into our integral yields $\frac{d-1}{24}$. Consequently, the ELSV formula immediately gives us a closed-form expression for the Hurwitz number:

$$H_{1,(d)} = \frac{(d-1)(d+1)d^d}{24}.$$

Finding this exact polynomial form purely through the character theory of the symmetric group is highly non-trivial, yet it falls out of the ELSV formula as a direct consequence of basic intersection theory.

More importantly, Theorem 4.0.5 provides a systematic way to compute all the Hodge integrals posed in Question 4.0.2. Let us define the integral expression as a function of the partition parts:

$$P_{g,n}(\mu_1,\ldots,\mu_n) \coloneqq \int_{[\overline{\mathcal{M}}_{g,n}]}\frac{\Lambda^\vee(1)}{\prod_{i=1}^{n}(1-\mu_i\psi_i)}.$$

By expanding the denominator, we see that this is a symmetric polynomial in the variables μ_i whose terms are of degree $3g-3+n-k$, where $k=0,\ldots,g$, and the coefficients of $\mu_1^{d_1}\cdots\mu_n^{d_n}$ are exactly the Hodge integrals of the form

$$\int_{[\overline{\mathcal{M}}_{g,n}]}\psi_1^{d_1}\cdots\psi_n^{d_n}\lambda_k.$$

Because the Hurwitz numbers $H_{g,\mu}$ can be computed purely combinatorially (e.g., via the cut-and-join recursion [**CM16**, Chapter 10]), one can evaluate this polynomial at a sufficiently large finite number of integer points μ_i. By polynomial interpolation, one can then recover all the coefficients of $P_{g,n}$. Once these coefficients are known, Faber's algorithm [**Fab99a**] can be applied to deduce all other Hodge numbers. See [**Kaz09**] for details.

Perhaps the most surprising consequence is that, in addition to offering the second method for calculating all purely ψ-class intersection numbers, the ELSV formula can be used to deduce Witten's conjecture (Theorem 4.0.3). The details of this spectacular proof were carried out by Okounkov and Pandharipande [**OP09**]. This application reveals that the ELSV formula is not just a computational tool, but also a foundational bridge connecting enumerative combinatorics, Gromov–Witten theory, and theoretical physics.

The primary goal of this chapter is to present a rigorous proof of the ELSV formula (Theorem 4.0.5) utilizing the machinery developed in the preceding chapters. In Section 4.1, we briefly survey the algebraic proof via virtual localization techniques following [**GV03**]. We analyze the technical difficulties encountered when applying localization in the moduli stack of stable curves $\overline{\mathcal{M}}_g(\mathbb{P}^1, d)$, which naturally motivates the need for and construction of the moduli space of relative stable maps. Section 4.2 defines relative stable maps and collects several fundamental properties of their associated moduli spaces. We demonstrate the theoretical advantages of working in the relative setting. This includes the straightforward identification of torus-fixed loci and the property that the moduli space of relative stable maps admits a canonical perfect obstruction theory. Finally, in Section 4.3, we carry out the explicit computations following the procedure in [**GP99**, Section 4], thereby completing the proof of the ELSV formula.

4.1. Graber and Vakil's Proof Revisit

Building on the ideas from the original proof in [**ELSV01**], Graber and Vakil provided the first rigorous proof using virtual localization. In this subsection, we briefly survey their approach and the technical difficulties they resolved. The main reference for this subsection is [**GV03**].

To begin, they needed a parameter space that naturally parameterizes branched coverings $f\colon C \to \mathbb{P}^1$ with the required ramification data. They considered $\overline{\mathcal{M}}_{g,0}(\mathbb{P}^1, d)$, the moduli space parameterizing stable maps of degree d from a nodal curve C to $\mathbb{P}^1$. This space naturally accommodates branched coverings with prescribed ramification profiles.

To connect this moduli stack with virtual localization techniques, Graber and Vakil adapted the treatment of the Lyashko–Looijenga map from the original ELSV proof [**ELSV01**]. Specifically, they utilized the extension of the branch morphism on the moduli stack $\overline{\mathcal{M}}_g(\mathbb{P}^1, d)$ introduced by Fantechi and Pandharipande in [**FP02**].

More precisely, let $f\colon C \to \mathbb{P}^1$ be a morphism of degree d, where C is a nonsingular algebraic curve of genus g. The ramification divisor R of f on C is defined by requiring the sequence

$$0 \to f^*\omega_{\mathbb{P}^1} \to \omega_C \to \omega_C|_R \to 0$$

to be exact. The branch divisor $\mathrm{br}(f)$ on $\mathbb{P}^1$ is then defined as the pushforward $f_*(R)$. Equivalently, one can define the ramification divisor on C via the length of the relative differentials:

$$R := \sum_{P \in C} \mathrm{length}(\Omega_{C/\mathbb{P}^1})_P \cdot P,$$

and then push this divisor R down to $\mathbb{P}^1$. By the Riemann–Hurwitz formula, the branch divisor $\mathrm{br}(f)$ is a divisor on $\mathbb{P}^1$ of degree

$$b = 2g - 2 + 2d.$$

This yields a set-theoretic map

$$\mathrm{br}\colon \mathcal{M}_g(\mathbb{P}^1, d) \to \mathrm{Sym}^b\,\mathbb{P}^1,$$

where $\mathcal{M}_g(\mathbb{P}^1, d)$ is the moduli stack parameterizing morphisms $f\colon C \to \mathbb{P}^1$ of degree d from a smooth curve C of genus g. Surprisingly, although the compactified space $\overline{\mathcal{M}}_g(\mathbb{P}^1, d)$ can be highly singular and ill-behaved, the open substack $\mathcal{M}_g(\mathbb{P}^1, d) \subseteq \overline{\mathcal{M}}_g(\mathbb{P}^1, d)$ is a smooth Deligne–Mumford stack of expected dimension b. Furthermore, the restriction of the virtual fundamental class $[\overline{\mathcal{M}}_g(\mathbb{P}^1, d)]^{\mathrm{vir}}$ coincides with the ordinary fundamental class of $\mathcal{M}_g(\mathbb{P}^1, d)$.

The proof of this fact relies on the construction of the canonical perfect obstruction theory on $\overline{\mathcal{M}}_g(\mathbb{P}^1, d)$, which induces a tangent-obstruction long exact sequence. More generally, let $\overline{\mathcal{M}} \coloneqq \overline{\mathcal{M}}_{g,n}(X, \beta)$. The canonical perfect obstruction theory is constructed from the deformation theory of maps. Consider the natural forgetful morphism

$$\tau\colon \overline{\mathcal{M}} \to \mathfrak{M},$$

where $\mathfrak{M}$ is the algebraic stack of quasi-stable curves. There is a relative cotangent complex $\mathbb{L}_\tau$ determined by a distinguished triangle in the derived category (Proposition 3.2.3):

$$\tau^*\mathbb{L}_{\mathfrak{M}} \to \mathbb{L}_{\overline{\mathcal{M}}} \to \mathbb{L}_\tau \to \tau^*\mathbb{L}_{\mathfrak{M}}[1].$$

On the other hand, the deformation theory of nodal curves is well-understood (Proposition 2.4.13). The deformation theory of maps $f\colon C \to X$ from a fixed domain curve C is likewise classical [**Ser06**, Section 3.4], with the tangent and obstruction spaces given by $\mathrm{H}^0(C, f^*TX)$ and $\mathrm{H}^1(C, f^*TX)$, respectively. One uses this deformation data to define a canonical relative perfect obstruction theory [**BF97**, Section 6] via

$$\widetilde{E}^\bullet \coloneqq [\mathrm{R}^\bullet\pi_*(f^*TX)]^\vee \to \mathbb{L}_\tau.$$

Here, $\pi\colon \mathcal{C} \to \overline{\mathcal{M}}$ is the universal curve over the moduli stack $\overline{\mathcal{M}}$. One then constructs the absolute perfect obstruction theory $E^\bullet$ of $\overline{\mathcal{M}}$ by completing the commutative diagram of distinguished triangles:

$$\begin{tikzcd}
\tau^*\mathbb{L}_{\mathfrak{M}} \arrow[r] \arrow[d, "\mathrm{id}"] & E^\bullet \arrow[r] \arrow[d, dashed, "\phi"] & \widetilde{E}^\bullet \arrow[r, "\xi"] \arrow[d] & \tau^*\mathbb{L}_{\mathfrak{M}}[1] \arrow[d, "\mathrm{id}"] \\
\tau^*\mathbb{L}_{\mathfrak{M}} \arrow[r] & \mathbb{L}_{\overline{\mathcal{M}}} \arrow[r] & \mathbb{L}_\tau \arrow[r] & \tau^*\mathbb{L}_{\mathfrak{M}}[1]
\end{tikzcd} \tag{4.1.1}$$

The top row is defined to be the distinguished triangle induced by the connecting morphism ξ. The bottom row is simply the standard distinguished triangle of cotangent complexes. Here ξ is defined in the way such that the right square of the above diagram commutes. We refer to [**Beh97**] for the technical details of this construction and a comprehensive treatment of the canonical perfect obstruction theory of $\overline{\mathcal{M}}_{g,n}(X, \beta)$.

Let $f\colon (C, p_1, \ldots, p_n) \to X$ represent a moduli point in $\overline{\mathcal{M}}$. Taking the dual of the top row of (4.1.1) yields a long exact sequence called the tangent-obstruction sequence:

$$\begin{aligned}
0 \to \mathrm{Ext}^0(\Omega_C(P), \mathcal{O}_C) &\to \mathrm{H}^0(C, f^*TX) \to \mathrm{Tan}(f) \\
\to \mathrm{Ext}^1(\Omega_C(P), \mathcal{O}_C) &\to \mathrm{H}^1(C, f^*TX) \to \mathrm{Obs}(f) \to 0.
\end{aligned}$$

This sequence is one of the most powerful computational tools afforded by the canonical perfect obstruction theory.

One can then apply the infinitesimal lifting criterion for smoothness (Theorem 2.3.36) to deduce that $\mathcal{M}_g(\mathbb{P}^1, d)$ is smooth. This amounts to showing that for every moduli point $[f : C \to \mathbb{P}^1]$, the obstruction space vanishes, i.e., $\mathrm{Obs}(f) = 0$. Since C is smooth, we have $\mathrm{Ext}^1(\Omega_C, \mathcal{O}_C) \cong \mathrm{H}^1(C, T_C)$. Because the map $\mathrm{H}^1(C, T_C) \to \mathrm{H}^1(C, f^*T\mathbb{P}^1)$ is induced by the sheaf map $T_C \to f^*T\mathbb{P}^1$ (whose cokernel is a torsion sheaf supported entirely at the ramification points, meaning its H^1 is zero), the map on cohomology is surjective. Therefore, $\mathrm{Obs}(f) = 0$, which immediately implies smoothness.

To summarize, we have established the following result using the tangent-obstruction theory:

Proposition 4.1.1 (Smoothness of $\mathcal{M}_g(\mathbb{P}^1, d)$)**.** *Let $d \geq 1$ be an integer. The open substack $\mathcal{M}_g(\mathbb{P}^1, d) \subseteq \overline{\mathcal{M}}_g(\mathbb{P}^1, d)$ is a smooth Deligne–Mumford stack of expected dimension $b = 2g-2+2d$, and the restriction of the virtual fundamental class $[\overline{\mathcal{M}}_g(\mathbb{P}^1, d)]^{\mathrm{vir}}$ is precisely the ordinary fundamental class of $\mathcal{M}_g(\mathbb{P}^1, d)$.*

We continue our discussion of the branch morphism $\mathrm{br} \colon \mathcal{M}_g(\mathbb{P}^1, d) \to \mathrm{Sym}^b \mathbb{P}^1$. Suppose that $[f : C \to \mathbb{P}^1]$ is a stable map. We define the extended branch divisor $\mathrm{br}(f)$ as follows. Let $N \subset C$ be the cycle of nodes of C, and let $\nu \colon \widetilde{C} \to C$ be the normalization of C. Let $A_1, \ldots, A_a$ be the irreducible components of $\widetilde{C}$ that dominate $\mathbb{P}^1$, and let $a_i \colon A_i \to \mathbb{P}^1$ denote the natural restricted maps. For these dominating components, the classical branch divisor $\mathrm{br}(a_i)$ is well-defined.

Let $B_1, \ldots, B_b$ be the components of $\widetilde{C}$ that are contracted by f over $\mathbb{P}^1$, with $f(B_j) = q_j \in \mathbb{P}^1$. We then define the extended branch divisor for π as

$$\mathrm{br}(f) \coloneqq \sum_{i=1}^{a} \mathrm{br}(a_i) + \sum_{j=1}^{b} (2g(B_j) - 2)\,[q_j] + 2f_*(N). \tag{4.1.2}$$

One can verify set-theoretically that this formula associates to f an effective divisor of degree $b = 2g - 2 + 2d$ on $\mathbb{P}^1$. This gives a set-theoretic extension $\mathrm{br} \colon \overline{\mathcal{M}}_g(\mathbb{P}^1, d) \to \mathrm{Sym}^b \mathbb{P}^1$. The main consequence established by Fantechi and Pandharipande [**FP02**] is that this extended map br is, in fact, a morphism of Deligne–Mumford stacks.

With this preparatory work complete, we can finally establish the setup for the proof. We first fix the notation. Let C be a smooth curve of genus g, and let $f : C \to \mathbb{P}^1$ be a map with ramification profile $\mu = (\mu_1, \ldots, \mu_n)$ over ∞, where μ is a partition of the degree d. By the Riemann–Hurwitz formula, the number of simple ramification points of f away from ∞ is $r = 2g - 2 + d + n$. Let $k = \sum_{i=1}^{n} (\mu_i - 1) = d - n$ be the total ramification index over ∞. The branch morphism is then given by

$$\mathrm{br} \colon \overline{\mathcal{M}}_g(\mathbb{P}^1, d) \to \mathrm{Sym}^b(\mathbb{P}^1),$$

where $b = r + k = 2g - 2 + 2d$ is the total degree of the branch divisor.

Consider the linear subspace $L_\infty \subset \mathrm{Sym}^b \mathbb{P}^1$ consisting of divisors of the form $k[\infty] + D$, where D is an effective divisor of degree $\deg D = b - k = r$. Since f has ramification profile μ over ∞, its branch divisor $\mathrm{br}(f)$ naturally belongs to L_∞. The pullback of the linear space L_∞ along the branch morphism yields a substack M of expected dimension r. More precisely, M is defined by the following Cartesian square:

$$\begin{tikzcd} M \arrow[r, hook] \arrow[d] \arrow[dr, phantom, "\square"] & \overline{\mathcal{M}}_g(\mathbb{P}^1, d) \arrow[d, "\mathrm{br}"] \\ L_\infty \arrow[r, hook] & \mathrm{Sym}^b(\mathbb{P}^1) \end{tikzcd}$$

Within M, we define M^μ (equipped with the reduced substack structure) to be the locus parameterizing maps from irreducible curves that have exactly the ramification profile μ over ∞. Through this construction, one can show that M^μ is irreducible with expected dimension r, and that it appears as a component in M with multiplicity $m_\mu = k! \prod_{i=1}^n \frac{\mu_i^{\mu_i - 1}}{\mu_i!}$. We refer the reader to [**OP09**, Section 7.3.3] for a detailed proof of this multiplicity.

One might ask why we isolate such a specific substack M. The reason is that, within this specifically chosen algebraic "container", the Hurwitz number can be recovered directly as

$$H_{g,\mu} = \frac{1}{|\mathrm{Aut}\,\mu|} \int_{[M^\mu]} \mathrm{br}^*[p], \tag{4.1.3}$$

where p is the Poincaré dual of the point class of $L_\infty \subset \mathrm{Sym}^b \mathbb{P}^1$.

To see why this holds, observe that the fiber $\mathrm{br}^{-1}\left(\sum_{i=1}^r [z_i] + k[\infty]\right)$ over a general point must lie within the smooth locus $M^\mu \cap \mathcal{M}_g(\mathbb{P}^1, d)$; otherwise, there would be a double point in the branch divisor (4.1.2) away from ∞. On the other hand, recall from Proposition 4.1.1 that the open stack $\mathcal{M}_g(\mathbb{P}^1, d)$ is smooth. Therefore, by Bertini's theorem [**Har77**, Theorem 8.18], the preimage of a general divisor $\sum_{i=1}^r [z_i] + k[\infty]$ intersects the substack M^μ transversally, resulting in a finite number of smooth, reduced points in $M^\mu \cap \mathcal{M}_g(\mathbb{P}^1, d)$. By definition, the number of these intersection points (weighted by their automorphism groups) is exactly the Hurwitz number $H_{g,\mu}$.

However, working within this algebraic container $M^\mu \subset M$ makes the localization calculations exceedingly involved. One utilizes the standard $\mathbb{C}^*$-action on $\mathbb{P}^1$, which naturally induces an action on $\overline{\mathcal{M}}_g(\mathbb{P}^1, d)$ and renders the branch morphism $\mathbb{C}^*$-equivariant. We denote by F_l the fixed components of the moduli stack $\overline{\mathcal{M}}_g(\mathbb{P}^1, d)$, which were systematically studied by Graber and Pandharipande in [**GP99**, Section 4]. Our next step is to localize the integral

$$\int_{[M^\mu]} \mathrm{br}^*[p]$$

to each fixed locus F_l.

To calculate the contribution of this integral over each fixed component, a more refined series of classes on M is required. In [**GV03**], Graber and Vakil define

$$[M]^{\mathrm{vir}} = \sum_n \iota_* \Gamma_n, \quad \Gamma_n \in \mathrm{CH}_*^{\mathbb{C}^*}(M_n),$$

where the M_n are the irreducible components of M. In particular, $\Gamma_\mu = m_\mu[M^\mu]$ as shown previously. Next, they localize each Γ_n to F_l, defining classes $\eta_{l,n} \in \mathrm{CH}_*^{\mathbb{C}^*}(F_l) \otimes_{\mathbb{Q}[t]} \mathbb{Q}[t, t^{-1}]$ such that

$$\sum_l \iota_* \eta_{l,n} = \Gamma_n.$$

These classes are uniquely determined by the virtual localization formula (Theorem 3.1.7). Alternatively, one can first localize M at the fixed loci F_l by introducing classes ξ_l such that $[M]^{\mathrm{vir}} = \sum_l \iota_* \xi_l$; then we obtain $\xi_l = \sum_n \eta_{l,n}$.

The final consequence is that we ultimately only need to consider the contribution from a single distinguished component, F_0. The general point of F_0 parameterizes a stable map featuring a single irreducible genus g component contracted over $0 \in \mathbb{P}^1$, along with n rational tails of degrees μ_i $(1 \le i \le n)$ mapping to $\mathbb{P}^1$, each totally ramified over 0 and ∞. Graber and Vakil proved [**GV03**, Proposition 4.7] that

$$m_\mu \int_{[M^\mu]} \mathrm{br}^*[p] = \int_{F_0} \mathrm{br}^*[p] \cap \xi_0. \tag{4.1.4}$$

As one might expect, carrying out this calculation is a combinatorial nightmare, detailed across [**GV03**, Lemmas 4.2–4.5]. Once isolated to F_0, however, one can finally apply the explicit calculations from [**GP99**, Section 4] to reach the final answer.

The technical nightmare here arises primarily from the container $M^\mu \subset M$. We have very little global information about the substack M and its irreducible components, making the calculation difficult because we need to idenitfy the virtual fundamental classes and normal bundles for the virtual localization formula. This explains why Graber and Vakil went to such extreme lengths to establish (4.1.4), allowing them to reduce the problem to known results. From this perspective, the space $\overline{\mathcal{M}}_g(\mathbb{P}^1, d)$ is far from the optimal choice to contain all the branched maps $f\colon C \to \mathbb{P}^1$; it is simply too "big" and allows for too much degenerate geometric "noise" that ultimately fails to contribute to the integral.

Motivated by the analytical difficulties of handling these fixed loci, Graber and Vakil hypothesized the existence of an ideal moduli space $\mathcal{M}$ of "relative stable maps" that would bypass these complications entirely. They proposed that this optimal proper Deligne–Mumford stack should satisfy three key properties (see [**GV03**, §5]):

(1) $\mathcal{M}$ contains an open substack U parameterizing maps where the preimage of ∞ consists of exactly n smooth points with the prescribed ramification profile $\mu = (\mu_1, \ldots, \mu_n)$.
(2) It admits a Fantechi–Pandharipande branch morphism to $\operatorname{Sym}^b \mathbb{P}^1$. Under this map, the fiber over the specific branch divisor $k[\infty] + r[0]$ is precisely the single contributing fixed locus F_0.
(3) $\mathcal{M}$ carries a $\mathbb{C}^*$-equivariant perfect obstruction theory. On the open substack U, this naturally restricts to $R^\bullet\pi_*\big(f^*T\mathbb{P}^1 \otimes \mathcal{O}(-\sum \mu_i x_i)\big)^\vee$ (relatively over $\mathfrak{M}$), where the x_i are the marked preimages of ∞.

If such conditions are satisfied, the need for applying combinatorial analyses of the fixed loci in (4.1.4) is eliminated. One could directly apply virtual localization to conclude the proof by repeating the calculation in [**GP99**, Section 4]. The natural question was whether such a compact Deligne–Mumford stack actually exists. While a space with some of these properties was first introduced in the symplectic category [**LR01**], the rigorous construction in the algebraic category was achieved by J. Li in [**Li01, Li02**], forming the core subject of the next subsection.

4.2. Moduli Spaces of Relative Stable Maps

Our next goal is to introduce the moduli space of relative stable maps and collect some of its most important properties. The main references for this section are [**Li01, Li02, GV05**].

We first fix the notation. Suppose that g is a non-negative integer and $\mu = (\mu_1, \ldots, \mu_n)$ is a partition of a positive integer d. Let $\mathcal{M}_g(\mathbb{P}^1, \mu)$ be the moduli space of ramified covers $f\colon (C, x_1, \cdots, x_n) \to (\mathbb{P}^1, \infty)$ of degree d from a smooth curve C of genus g to $\mathbb{P}^1$, such that the ramification profile over the distinguished point $\infty \in \mathbb{P}^1$ is exactly specified by the partition μ (i.e., $f^{-1}(\infty) = \sum_{k=1}^n \mu_k x_k$ as Cartier divisors), and all other branch points are simple. Our goal is to find a better compactification of $\mathcal{M}_g(\mathbb{P}^1, \mu)$ than the moduli stack of stable curves $\overline{\mathcal{M}}_g(\mathbb{P}^1, d)$ discussed in Section 4.1. The following notion of a relative stable map serves this exact purpose.

Before defining relative stable maps, we clarify the notion of an expanded target. Let m be a non-negative integer, and let

$$\mathbb{P}^1(m) \coloneqq \mathbb{P}^1_1 \cup \cdots \cup \mathbb{P}^1_m$$

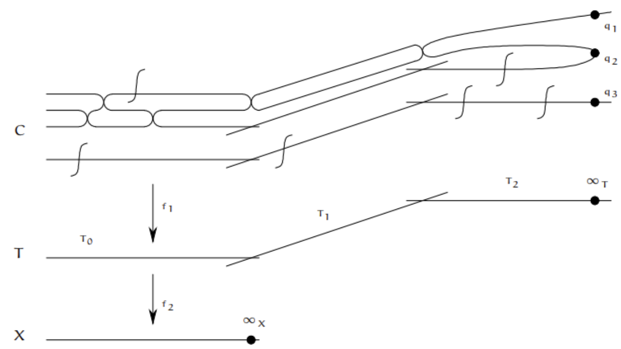

FIGURE 1. Relative Stable Maps (From [Vak08])

be a chain of m copies of $\mathbb{P}^1$ joined at nodes $q_l^1 = \mathbb{P}_l^1 \cap \mathbb{P}_{l+1}^1$ for $l = 1, \ldots, m-1$. We fix smooth points $q_0^1 \in \mathbb{P}_1^1$ and $\infty_m = q_m^1 \in \mathbb{P}_m^1$. We then define the expansion

$$\mathbb{P}^1[m] \coloneqq \mathbb{P}_0^1 \cup \mathbb{P}^1(m)$$

by identifying $\infty_0 \in \mathbb{P}_0^1$ with $q_0^1 \in \mathbb{P}_1^1$. In particular, $\mathbb{P}^1[0] \coloneqq \mathbb{P}_0^1 = \mathbb{P}^1$. We call the component $\mathbb{P}_0^1$ the *root component* of $\mathbb{P}^1[m]$, and the components $\mathbb{P}_1^1, \ldots, \mathbb{P}_m^1$ the *bubble components*.

Because there are technically multiple "infinity" points in the construction of the expansion $\mathbb{P}^1[m]$, we will be precise: ∞_0 refers to the original infinity point on the root component $\mathbb{P}_0^1$ (which becomes a node for $m > 0$), and ∞_m refers to the smooth infinity point on the final bubble component $\mathbb{P}_m^1$.

Definition 4.2.1 (Relative Maps)**.** Let $(C, x_1, \ldots, x_n)$ be a connected nodal curve of arithmetic genus g with distinct smooth marked points $x_1, \ldots, x_n$. A relative map to $\mathbb{P}^1$ modulo ∞ consists of the following data:

(1) A morphism $f_1 \colon (C, x_1, \ldots, x_n) \to (\mathbb{P}^1[m], \infty_m)$, such that $f_1^*(\infty_m) = \sum_{k=1}^n \mu_k x_k$ as a Cartier divisor. In particular, the full set-theoretic preimage $f_1^{-1}(\infty_m)$ consists entirely of the smooth marked points of C.
(2) A projection $f_2 \colon \mathbb{P}^1[m] \to \mathbb{P}^1$ that contracts all the bubble components $\mathbb{P}_i^1$ $(i \geq 1)$ to $\infty \in \mathbb{P}^1$, and restricts to an isomorphism from the root component $\mathbb{P}_0^1$ to $\mathbb{P}^1$, taking ∞_0 to $\infty \in \mathbb{P}^1$.
(3) *Predeformation Condition (Kissing Condition)*: The preimage of each node q_i^1 of $\mathbb{P}^1[m]$ is a union of nodes of C. Moreover, at any such node of C, the two branches map to the two distinct branches of q_i^1, and their local degrees of branching are identical.

By a slight abuse of notation, we usually refer to f_1 simply as the relative map. Due to the predeformation condition, if $f_1 \colon C \to \mathbb{P}^1[m]$ is a relative map of total degree d, the restriction of f_1 to the preimage of each component $f_1^{-1}(\mathbb{P}_i^1)$ also has degree d.

An isomorphism between two relative maps is defined by the following commutative diagram, where all horizontal maps are isomorphisms:

$$\begin{tikzcd}
(C, \mathbf{x}) \arrow[r, "\cong"] \arrow[d, "f_1"] & (C', \mathbf{x}') \arrow[d, "f_1'"] \\
(\mathbb{P}^1[m], \infty_m) \arrow[r, "\cong"] \arrow[d, "f_2"] & (\mathbb{P}^1[m], \infty_m) \arrow[d, "f_2'"] \\
(\mathbb{P}^1, \infty) \arrow[r, "\mathrm{id}"] & (\mathbb{P}^1, \infty)
\end{tikzcd}$$

Because the bottom map is required to be the strict identity, the middle isomorphism must preserve the identification of $\mathbb{P}^1_0$ with $\mathbb{P}^1$ and is therefore the identity on the root component $\mathbb{P}^1_0$. On the bubble components, however, the isomorphism is allowed to be a non-trivial scaling ($\mathbb{C}^*$-action) that preserves the nodes.

As is standard, a relative map $f\colon (C, \mathbf{x}) \to (\mathbb{P}^1[m], \infty_m)$ is called *stable* if its automorphism group $\mathrm{Aut}(f)$ is finite. We have the following explicit criterion for stability, the proof of which is similar to that of Proposition 2.4.2.

Proposition 4.2.2 (Characterization of Relative Stable Maps)**.** *Let $f\colon (C, \mathbf{x}) \to (\mathbb{P}^1[m], \infty_m)$ be a relative map. Then f is stable if and only if all of the following hold:*

(1) Any irreducible component of C of geometric genus zero that is contracted by f_1 contains at least three special points (nodes or markings);

(2) Any irreducible component of C of geometric genus one that is contracted by f_1 contains at least one special point;

(3) For every bubble component $\mathbb{P}^1_i$ $(0 < i < m)$, it is not the case that every domain component mapping to $\mathbb{P}^1_i$ takes the totally ramified form $[x, y] \mapsto [x^q, y^q]$ relative to the nodes $[0, 1] = q^1_{i-1}$ and $[1, 0] = q^1_i$. The same condition holds for the final bubble $\mathbb{P}^1_m$ relative to the node and ∞_m.

To define the moduli stack of relative stable maps, we define the notion of a family of relative stable maps over a base scheme S. Crucially, the number of bubble components m is allowed to jump across the fibers of a family. Therefore, a family is defined using a flat family of expanded targets rather than a fixed $\mathbb{P}^1[m]$.

Definition 4.2.3 (Family of Relative Stable Maps)**.** Let S be a scheme. A family of relative stable maps over S consists of the data $(\pi_{\mathcal{C}}\colon \mathcal{C} \to S, \{x_i\}_{i=1}^n, \pi_{\mathcal{P}}\colon \mathcal{P} \to S, \mathcal{D}, f, \rho)$, where:

(1) $\pi_{\mathcal{C}}\colon \mathcal{C} \to S$ is a flat family of nodal curves with n disjoint sections $x_i\colon S \to \mathcal{C}$;

(2) $\pi_{\mathcal{P}}\colon \mathcal{P} \to S$ is a flat family of targets, equipped with a section $\mathcal{D} \subset \mathcal{P}$, such that every geometric fiber $(\mathcal{P}_s, \mathcal{D}_s)$ is isomorphic to an expanded target $(\mathbb{P}^1[m], \infty_m)$ for some integer $m \geq 0$;

(3) $f\colon \mathcal{C} \to \mathcal{P}$ is an S-morphism such that $f^*\mathcal{D} = \sum_{k=1}^n \mu_k x_k$;

(4) $\rho\colon \mathcal{P} \to \mathbb{P}^1 \times S$ is a stabilizing contraction over S;

(5) For every geometric point $s \in S$, the restriction $f_s\colon \mathcal{C}_s \to \mathcal{P}_s$ satisfies the predeformation condition and is a relative stable map.

$$\begin{tikzcd}
\mathcal{C} \arrow[r, "f"] \arrow[d, "\pi_{\mathcal{C}}"'] & \mathcal{P} \arrow[r, "\rho"] \arrow[dl, "\pi_{\mathcal{P}}"] & \mathbb{P}^1 \times S \arrow[dll] \\
S & &
\end{tikzcd}$$

We can now define the prestack of relative stable maps, denoted by $\overline{\mathcal{M}}_g(\mathbb{P}^1, \mu)$. We remark that the theory of relative stable maps can be vastly generalized to maps into any smooth complex projective variety X relative to a smooth divisor $D \subset X$. Note that when $D = \varnothing$, it naturally recovers the standard moduli space of (absolute) stable maps; this is where the term "relative" stems from.

We collect some fundamental properties of this moduli stack. We refer the reader to [**Li01**] and [**Li02**] for the foundational proofs.

Theorem 4.2.4. *The moduli prestack of relative stable maps $\overline{\mathcal{M}}_g(\mathbb{P}^1, \mu)$ is a proper Deligne–Mumford stack equipped with a canonical perfect obstruction theory of expected dimension $r = 2g - 2 + d + n$, where μ is a partition of d of length n. Consequently, it admits a well-defined virtual fundamental class:*

$$[\overline{\mathcal{M}}_g(\mathbb{P}^1, \mu)]^{\mathrm{vir}} \in \mathrm{CH}_r\left(\overline{\mathcal{M}}_g(\mathbb{P}^1, \mu)\right).$$

Finally, we discuss how to induce the standard $\mathbb{C}^*$-action on the moduli stack of relative stable maps. We first consider the simple case of a single moduli point represented by a map $[f_1 \colon (C, \mathbf{x}) \to (\mathbb{P}^1[m], \infty_m)]$. The action of $t \in \mathbb{C}^*$ is given by the following commutative diagram:

$$\begin{tikzcd}
(C, \mathbf{x}) \arrow[r, "f_1"] \arrow[d, "="] & (\mathbb{P}^1[m], \infty_m) \arrow[r, "f_2"] \arrow[d, "\widetilde{\sigma}_t"] & \mathbb{P}^1 \arrow[d, "\sigma_t"] \\
(C, \mathbf{x}) \arrow[r, "f_1"] & (\mathbb{P}^1[m], \infty_m) \arrow[r, "f_2"] & \mathbb{P}^1
\end{tikzcd}$$

Here, σ_t is the standard $\mathbb{C}^*$-action on $\mathbb{P}^1$ fixing 0 and ∞, and $\widetilde{\sigma}_t$ is its unique lift to the expanded target that preserves all nodes and the marked point ∞_m. Specifically, on the root component $\mathbb{P}^1_0$, $\widetilde{\sigma}_t$ is simply the standard action. On each bubble component $\mathbb{P}^1_i$ $(1 \le i \le m)$, we can assign projective coordinates $[x_i : y_i]$ such that the lower node q^1_{i-1} is $[0 : 1]$ and the upper node q^1_i (or ∞_m) is $[1 : 0]$. The lifted action on this bubble is then defined by the standard scaling $t \cdot [x_i : y_i] = [tx_i : y_i]$.

To extend this $\mathbb{C}^*$-action to any family of relative stable maps over a base scheme S, one applies the action fiberwise to the target. The standard $\mathbb{C}^*$-action lifts to a canonical automorphism on the flat family of expanded targets $\pi_{\mathcal{P}} \colon \mathcal{P} \to S$. The action on the moduli stack $\overline{\mathcal{M}}_g(\mathbb{P}^1, \mu)$ is then defined by post-composing the relative map $f \colon \mathcal{C} \to \mathcal{P}$ with this target automorphism.

Thanks to the existence of the $\mathbb{C}^*$-equivariant perfect obstruction theory, we can apply the virtual localization formula (Theorem 3.2.11) directly to the moduli stack $\overline{\mathcal{M}}_g(\mathbb{P}^1, \mu)$ to conclude the proof of Theorem 4.0.5. It is worthwhile to remark that the technical prerequisite for applying virtual localization in Theorem 3.2.11—namely, constructing a global $\mathbb{C}^*$-equivariant embedding into a smooth Deligne–Mumford stack—is non-trivial. However, this embedding is carefully constructed in [**GV05**]. This remarkable paper by Graber and Vakil also includes a refined analysis of virtual localization on $\overline{\mathcal{M}}_{g,n}(X, D)$ for general relative stable maps, providing important applications to the study of the tautological ring. We also refer to [**FP05**] for a detailed analysis of the moduli stack of relative stable maps and its profound connection to tautological classes on the moduli space of curves.

4.3. Final Calculations

As promised, we conclude this section by applying the virtual localization formula (Theorem 3.2.11) to the moduli stack of relative stable maps $\overline{\mathcal{M}}_g(\mathbb{P}^1, \mu)$ to obtain the ELSV formula (Theorem 4.0.5). The calculation in this section closely follows the procedure in [**GP99**, Section 4], adapting the absolute setting to the relative one. We divide the calculation into the following steps: First, we identify the Hurwitz number $H_{g,\mu}$ with the degree of the (relative) branch morphism. Next, we lift this integral to the equivariant setting and localize it to the fixed components. We then identify the contributing fixed locus, analyze its virtual normal bundle, and compute its contribution to the integral. These steps assemble to yield the final ELSV formula.

4.3.1. Localizations. In the previous section, we discussed the absolute branch morphism

$$\mathrm{br}_{\mathrm{abs}}\colon \overline{\mathcal{M}}_g(\mathbb{P}^1, d) \to \mathrm{Sym}^b\,\mathbb{P}^1$$

via the extension (4.1.2), establishing that it is a morphism of Deligne–Mumford stacks. One can analogously define the relative branch morphism on the moduli stack of relative stable maps:

$$\mathrm{br}\colon \overline{\mathcal{M}}_g(\mathbb{P}^1, \mu) \to \mathrm{Sym}^r\,\mathbb{P}^1.$$

To do this for a moduli point $[(C, \mathbf{x}) \to (\mathbb{P}^1[m], \infty_m) \to \mathbb{P}^1]$, we first evaluate the absolute branch divisor of the composed map $f_2 \circ f_1\colon C \to \mathbb{P}^1$. Because the relative map carries a fixed ramification profile μ over ∞_m, this absolute branch divisor is guaranteed to take the form $D_{\mathrm{free}} + k[\infty]$, where $k = d - n$. We define the relative branch morphism to be exactly this residual free divisor $D_{\mathrm{free}} \in \mathrm{Sym}^r\,\mathbb{P}^1$. This relative branch morphism remains a well-defined morphism of Deligne–Mumford stacks, and it is $\mathbb{C}^*$-equivariant.

Using the standard identification $\mathbb{P}^r \xrightarrow{\cong} \mathrm{Sym}^r\,\mathbb{P}^1$ given by

$$[a_0\colon \cdots : a_r] \mapsto \mathrm{div}\left(\sum_{i=0}^{r} a_i x^i y^{r-i}\right),$$

we can regard br as a morphism from $\overline{\mathcal{M}}_g(\mathbb{P}^1, \mu)$ to $\mathbb{P}^r$.

We claim that

$$\begin{aligned} H_{g,\mu} &= \frac{1}{|\operatorname{Aut}(\mu)|}\deg(\mathrm{br}) \\ &= \frac{1}{|\operatorname{Aut}(\mu)|}\int_{[\overline{\mathcal{M}}_g(\mathbb{P}^1,\mu)]^{\mathrm{vir}}} \mathrm{br}^* e^r. \end{aligned} \tag{4.3.1}$$

Here, $e \in \mathrm{CH}^*(\mathbb{P}^r)$ is the generator of the Chow ring $\mathrm{CH}^*(\mathbb{P}^r) = \mathbb{Q}[e]/(e^{r+1})$, making e^r the Poincaré dual of the point class in $\mathbb{P}^r$. Recall again that all Chow and cohomology rings are taken with $\mathbb{Q}$-coefficients. The class $[\overline{\mathcal{M}}_g(\mathbb{P}^1, \mu)]^{\mathrm{vir}} \in \mathrm{CH}_r\left(\overline{\mathcal{M}}_g(\mathbb{P}^1, \mu)\right)$ is the virtual fundamental class introduced in Theorem 4.2.4.

The justification for equality (4.3.1) mirrors the argument for (4.1.3). First, similar to the proof of Proposition 4.1.1, one uses the canonical perfect obstruction theory (constructed in [**Li02**]) to show that the open substack $\mathcal{M}_g(\mathbb{P}^1, \mu) \subset \overline{\mathcal{M}}_g(\mathbb{P}^1, \mu)$ is a smooth Deligne–Mumford stack of expected dimension r. This is done by verifying that for every moduli point $[f\colon (C, \mathbf{x}) \to \mathbb{P}^1]$, the obstruction space vanishes ($\mathrm{Obs}(f) = 0$), allowing the use of the infinitesimal lifting criterion for smoothness (Theorem 2.3.36). Consequently, the ordinary fundamental class of the open locus is simply the restriction of the virtual fundamental class $[\overline{\mathcal{M}}_g(\mathbb{P}^1, \mu)]^{\mathrm{vir}}$.

After establishing smoothness, we apply Bertini's theorem. It is clear that the preimage $\mathrm{br}^{-1}(\sum_{i=1}^r [z_i])$ of r distinct points must fall within the smooth locus $\mathcal{M}_g(\mathbb{P}^1, \mu)$; otherwise, the extension formula (4.1.2) would force a multiple point in the branch divisor. By Bertini's theorem [**Har77**, Theorem 8.18], the preimage of a general divisor intersects $\overline{\mathcal{M}}_g(\mathbb{P}^1, \mu)$ transversally at a finite number of smooth points. By the definition of the Hurwitz number, this intersection calculation yields exactly (4.3.1).

Next, we consider the standard $\mathbb{C}^*$-action on $\mathbb{P}^1$ given by

$$t \cdot [z_0 \colon z_1] \coloneqq [tz_0 \colon z_1].$$

For any $r \in \mathbb{Z}_{\geq 1}$, this induces a $\mathbb{C}^*$-action on $\mathbb{P}^r$. Under our identification $\mathbb{P}^r \cong \operatorname{Sym}^r \mathbb{P}^1$, this action is given by $t \cdot [a_0 \colon \cdots \colon a_r] \coloneqq [a_0 \colon t^{-1}a_1 \colon \cdots \colon t^{-r}a_r]$. This action has $r+1$ isolated fixed points $[1 \colon 0 \colon \cdots \colon 0], \ldots, [0 \colon \cdots \colon 0 \colon 1]$, which we denote by $p_0, \ldots, p_r$, respectively. By construction, the branch morphism $\mathrm{br} \colon \overline{\mathcal{M}}_g(\mathbb{P}^1, \mu) \to \mathbb{P}^r$ is $\mathbb{C}^*$-equivariant with respect to these actions.

We can now lift equation (4.3.1) to equivariant intersection theory to apply virtual localization. Recall that the $\mathbb{C}^*$-equivariant Chow ring of $\mathbb{P}^r$ is given by:

$$\mathrm{CH}^*_{\mathbb{C}^*}(\mathbb{P}^r) = \frac{\mathbb{Q}[e,u]}{\langle e(e-u)\cdots(e-ru)\rangle},$$

where $e = c_1^{\mathbb{C}^*}(\mathcal{O}_{\mathbb{P}^r}(1))$ and $u = c_1^{\mathbb{C}^*}(\mathbb{C}_1)$ is the equivariant parameter. Here, $\mathbb{C}_1$ denotes the one-dimensional representation of weight 1 (see Example 3.1.10).

We replace the point class e^r with an equivariant lift and localize the integral. Let F index the connected components of the fixed locus in $\overline{\mathcal{M}}_g(\mathbb{P}^1, \mu)$. We have

$$\begin{aligned} H_{g,\mu} &= \frac{1}{|\operatorname{Aut}(\mu)|} \int_{[\overline{\mathcal{M}}_g(\mathbb{P}^1,\mu)]^{\mathrm{vir}}_{\mathbb{C}^*}} \mathrm{br}^* \left(\prod_{i=0}^{r-1} (e - iu) \right) \\ &= \frac{1}{|\operatorname{Aut}(\mu)|} \sum_F \int_{[F]^{\mathrm{vir}}} \frac{\mathrm{br}^* \left(\prod_{i=0}^{r-1}(e-iu)\right)|_F}{e_{\mathbb{C}^*}(N_F^{\mathrm{vir}})}. \end{aligned} \tag{4.3.2}$$

Because br is $\mathbb{C}^*$-equivariant, the image of any fixed component F must be one of the fixed points $\{p_0, \ldots, p_r\}$ in $\mathbb{P}^r$. We can therefore group the fixed components by their image, defining

$$F_i = \overline{\mathcal{M}}_g(\mathbb{P}^1,\mu)^{\mathbb{C}^*} \cap \mathrm{br}^{-1}(p_i) \quad \text{for } i = 0, \ldots, r.$$

Equation (4.3.2) then becomes:

$$H_{g,\mu} = \frac{1}{|\operatorname{Aut}(\mu)|} \sum_{i=0}^{r} \int_{[F_i]^{\mathrm{vir}}} \frac{\mathrm{br}^* \left(\prod_{j=0}^{r-1}(e-ju)\right)|_{F_i}}{e_{\mathbb{C}^*}(N_{F_i}^{\mathrm{vir}})}. \tag{4.3.3}$$

To evaluate the numerator, we restrict the class to the fixed points. Since the weight of $\mathcal{O}_{\mathbb{P}^r}(1)$ at p_i is $i \cdot u$, the restriction of e to p_i is precisely $i \cdot u$. This yields:

$$\begin{aligned} \mathrm{br}^* \left(\prod_{j=0}^{r-1} (e - ju) \right) \Bigg|_{F_i} &= \prod_{j=0}^{r-1} (e-ju) \Bigg|_{p_i} \\ &= \prod_{j=0}^{r-1} (i-j)u \\ &= \begin{cases} 0, & 0 \leq i \leq r-1; \\ r!\ u^r, & i = r. \end{cases} \end{aligned}$$

Consequently, all terms in the sum for $i < r$ vanish, and (4.3.3) simplifies to:

$$H_{g,\mu} = \frac{1}{|\operatorname{Aut}(\mu)|} \int_{[F_r]^{\mathrm{vir}}} \frac{r!\ u^r}{e_{\mathbb{C}^*}(N_{F_r}^{\mathrm{vir}})}. \tag{4.3.4}$$

Our final tasks are to identify the structure of the contributing substack F_r and to compute the Euler class of its virtual normal bundle $N^{\text{vir}}_{F_r}$. This will be the primary focus of the following subsections.

4.3.2. Identification of the substack F_r. To identify the fixed locus F_r, we first need to understand the geometric meaning of the condition $\operatorname{br}(f) = p_r$. Under our identification $\mathbb{P}^r \cong \operatorname{Sym}^r \mathbb{P}^1$, the fixed point $p_r = [0 : \cdots : 0 : 1]$ corresponds to the divisor $r[0]$. This implies that for any relative stable map in F_r, all r free simple branch points have collided at the single point $0 \in \mathbb{P}^1$.

Because all the extra branching occurs over 0 and no extra ramification points collide with ∞, the target space does not need to degenerate to absorb them. Consequently, the target remains unexpanded (i.e., $m = 0$). All the topology (i.e., the genus and the internal nodes) of the domain curve is forced into the fiber over 0, leaving only simple rational tails connecting 0 to ∞.

We formalize this geometric picture into the following algebraic characterization.

Proposition 4.3.1 (Characterization of F_r)**.** *Let $[f\colon (C, \mathbf{x}) \to (\mathbb{P}^1[m], \infty_m)]$ be a moduli point in $\overline{\mathcal{M}}_g(\mathbb{P}^1, \mu)$. It belongs to the fixed locus F_r if and only if f satisfies all of the following conditions:*

(1) We have $m = 0$, so the target is simply $\mathbb{P}^1$, and the marked points $x_1, \dots, x_n$ map to $\infty \in \mathbb{P}^1$.
(2) The domain curve C is the union of a connected, arithmetic genus g subcurve C_0 and n smooth rational components $D_1, \dots, D_n$. Each D_i intersects C_0 at exactly one nodal point y_i.
(3) The subcurve C_0 is contracted by f to the point $0 \in \mathbb{P}^1$.
(4) Each rational component D_i contains the marked point x_i. The restriction $f|_{D_i}\colon D_i \to \mathbb{P}^1$ is a totally ramified cover of degree μ_i, branching exclusively over 0 (at the node y_i) and over ∞ (at the marked point x_i). Specifically, in local coordinates, $f|_{D_i}$ is of the form $z \mapsto z^{\mu_i}$.
(5) The contracted core $(C_0, y_1, \dots, y_n)$, viewed as a genus g curve with n marked points (the attachment nodes), is stable. In other words, $[(C_0, y_1, \dots, y_n)] \in \overline{\mathcal{M}}_{g,n}$.

Proof. The first condition is clear: because there is no extra ramification colliding with ∞_m besides the fixed profile μ, the target requires no bubble components, forcing $m = 0$. Thus, $[f]$ actually lies in the absolute space $\overline{\mathcal{M}}_g(\mathbb{P}^1, d)$ with a fixed profile over ∞.

By $\mathbb{C}^*$-equivariance, the irreducible components of C fall into two distinct categories:

- Components that are contracted to $0 \in \mathbb{P}^1$ under f. (Note that no component can be contracted to ∞, otherwise the preimage $f^{-1}(\infty)$ would be infinite, violating the relative profile condition).
- Components D_i that dominate $\mathbb{P}^1$. By $\mathbb{C}^*$-equivariance, these must be fully ramified over 0 and ∞, taking the local form $z \mapsto z^d$.

Let us first analyze the dominating components D_i. By the Riemann–Hurwitz formula, a curve covering $\mathbb{P}^1$ branched at only two points must be rational, so $D_i \cong \mathbb{P}^1$. Moreover, since the relative map requires $f^{-1}(\infty) = \sum_{k=1}^{n} \mu_k x_k$, there must be exactly n such components, which we denote by $D_1, \dots, D_n$. Each D_i contains exactly one marked point x_i mapping to ∞. To satisfy the ramification profile, we must have $\deg(f|_{D_i}) = \mu_i$, so $f|_{D_i}$ is locally of the form $z \mapsto z^{\mu_i}$.

Next, we analyze the components contracted to $0 \in \mathbb{P}^1$. Let C_0 be the union of all such contracted components. We claim that C_0 must be connected. Suppose for the sake of contradiction that C_0 consists of disjoint pieces, say $A \sqcup B$. Since the total curve C is connected, there must exist a dominating component D_i serving as a rational bridge between A and B. This would require D_i to intersect A and B at two distinct nodes, both of which must map to 0. However, the map $z \mapsto z^{\mu_i}$ has only a single point mapping to 0. This contradiction proves that the contracted locus C_0 is connected.

Since C is connected, the core C_0 must attach to each tail D_i at exactly one node, which we call y_i. Furthermore, because $D_i \cong \mathbb{P}^1$ contributes zero to the arithmetic genus of the total curve C, the arithmetic genus of C_0 itself must be exactly g by the genus formula.

Finally, because the map f is already defined on the rational tails (up to finite automorphisms), the stability of the entire relative map $f\colon (C, \mathbf{x}) \to \mathbb{P}^1$ reduces exactly to the stability of the contracted core. Therefore, the nodal curve $(C_0, y_1, \dots, y_n)$ must be stable, so it represents a moduli point in $\overline{\mathcal{M}}_{g,n}$. □

As a corollary of this geometric description, any automorphism of a map $[f] \in F_r$ consists of an automorphism of the core curve C_0 and independent phase rotations on each of the totally ramified tails. By the predeformation condition, the automorphisms of the tails relative to the core are given by:

$$\prod_{i=1}^{n} \mathrm{Aut}\left(f|_{D_i}\colon (D_i, y_i, x_i) \to (\mathbb{P}^1, 0, \infty)\right) = \prod_{i=1}^{n} \mathbb{Z}/\mu_i.$$

Moreover, the above characterization shows that an element $[f] \in F_r$ is completely and uniquely determined by the choice of the core data $(C_0, y_1, \cdots, y_n)$, which is simply a moduli point in $\overline{\mathcal{M}}_{g,n}$. Therefore, the fixed locus F_r has the structure of a gerbe over $\overline{\mathcal{M}}_{g,n}$, naturally yielding the stacky isomorphism

$$F_r \cong \overline{\mathcal{M}}_{g,n} \times \prod_{i=1}^{n} \mathsf{B}\left(\mathbb{Z}/\mu_i\right).$$

4.3.3. Identification of the Virtual Normal Bundle. Having identified the structure of the fixed substack F_r, our next goal is to compute the contribution of its virtual normal bundle $N_{F_r}^{\mathrm{vir}}$ to the localization formula (4.3.4). Our calculation follows the procedure detailed by Graber and Pandharipande in [**GP99**, Section 4], adapted to the relative setting. The essential tool for this computation is the tangent-obstruction long exact sequence.

Define sheaves T^1 and T^2 on F_r via the cohomology of the restricted canonical (dual) perfect obstruction theory $E_\bullet$ of $\overline{\mathcal{M}}_g(\mathbb{P}^1, \mu)$:

$$0 \to T^1 \to E_0|_{F_r} \to E_1|_{F_r} \to T^2 \to 0. \tag{4.3.5}$$

Following the arguments of the proof of Proposition 4.1.1, the canonical $\mathbb{C}^*$-equivariant perfect obstruction theory yields a long exact sequence of sheaves over the substack F_r:

$$\begin{aligned} 0 \to \mathrm{Ext}^0\left(\Omega_C\left(\sum_{i=1}^{n} x_i\right), \mathcal{O}_C\right) &\to \mathrm{H}^0\left(C, f^*T\mathbb{P}^1\left(-\sum_{i=1}^{n} \mu_i x_i\right)\right) \to T^1 \\ \to \mathrm{Ext}^1\left(\Omega_C\left(\sum_{i=1}^{n} x_i\right), \mathcal{O}_C\right) &\to \mathrm{H}^1\left(C, f^*T\mathbb{P}^1\left(-\sum_{i=1}^{n} \mu_i x_i\right)\right) \to T^2 \to 0. \end{aligned} \tag{4.3.6}$$

Note that by our analysis in Proposition 4.3.1, every moduli point in F_r takes the form $[f\colon (C, \mathbf{x}) \to (\mathbb{P}^1, \infty)]$, where the target is unexpanded ($m = 0$). Therefore, the deformation space $\mathrm{Def}(f)$ and obstruction space $\mathrm{Obs}(f)$ are precisely the T^1 and T^2 terms in the sequence above.

This exact sequence has six potentially non-trivial terms, which we denote by $B_1, \ldots, B_6$ respectively. For notational simplicity, we abbreviate a moduli point in F_r simply as $(C, \mathbf{x}, f)$. Recall also from Proposition 4.3.1 that the domain curve decomposes as $C = C_0 \cup \bigcup_{i=1}^n D_i$, where the n-pointed contracted core $(C_0, y_1, \ldots, y_n)$ is stable. We will denote this core moduli point by $(C_0, \mathbf{y}) \in \overline{\mathcal{M}}_{g,n}$.

For $i = 1, 2$, let $T^{i,f}$ and $T^{i,m}$ denote the fixed and moving parts of T^i under the $\mathbb{C}^*$-action, respectively. By the definition of virtual normal bundles and the sequence (4.3.5), the virtual tangent and normal bundles of F_r are given by K-theoretic differences:

$$T^{\mathrm{vir}}_{F_r} = T^{1,f} - T^{2,f}; \tag{4.3.7}$$

$$N^{\mathrm{vir}}_{F_r} = T^{1,m} - T^{2,m}. \tag{4.3.8}$$

We have in K-theory,

$$[T^1] - [T^2] = ([B_4] - [B_5]) - ([B_1] - [B_2]).$$

Taking the moving parts of both sides yields

$$[N^{\mathrm{vir}}_{F_r}] = ([B_4^m] - [B_5^m]) - ([B_1^m] - [B_2^m]).$$

Taking the equivariant Euler class gives:

$$\frac{1}{e_{\mathbb{C}^*}(N^{\mathrm{vir}}_{F_r})} = \frac{e_{\mathbb{C}^*}(B_1^m) e_{\mathbb{C}^*}(B_5^m)}{e_{\mathbb{C}^*}(B_4^m) e_{\mathbb{C}^*}(B_2^m)}. \tag{4.3.9}$$

To explicitly evaluate the right-hand side of (4.3.9), a finer geometric analysis of the terms B_1, B_2, B_4, and B_5 is required. We divide this task into two subsections.

4.3.4. Contribution from B_1 and B_4. We begin by analyzing the deformation theory of the domain curve, captured by B_1 and B_4.

- **Analysis of B_1:** The term B_1 represents the infinitesimal automorphisms of the marked domain curve:
$$B_1 = \mathrm{Aut}(C, \mathbf{x}) \cong \mathrm{Ext}^0\left(\Omega_C\left(\sum_{i=1}^n x_i\right), \mathcal{O}_C\right).$$
Decomposing C into its irreducible components according to Proposition 4.3.1, we find that
$$B_1 = \mathrm{Aut}(C_0, \mathbf{y}) \oplus \bigoplus_{i=1}^n \mathrm{Aut}(D_i, y_i, x_i).$$
Since the core curve $(C_0, \mathbf{y})$ is stable, it has no infinitesimal automorphisms, so $\mathrm{Aut}(C_0, \mathbf{y}) = 0$. The tails D_i are rational curves fixed at two points, so $\mathrm{Aut}(D_i, y_i, x_i) \cong \mathrm{H}^0(\mathbb{P}^1, T\mathbb{P}^1(-0-\infty)) \cong \mathbb{C}$. Since scaling the target does not induce any translation of these vector fields, this space carries weight zero as a $\mathbb{C}^*$-representation. Therefore, $B_1 = \mathbb{C}_0^{\oplus n}$, and consequently $B_1^m = 0$.
- **Analysis of B_4:** The term B_4 represents the infinitesimal deformations of the marked curve:
$$B_4 = \mathrm{Def}(C, \mathbf{x}) \cong \mathrm{Ext}^1\left(\Omega_C\left(\sum_{i=1}^n x_i\right), \mathcal{O}_C\right).$$

The local-to-global Ext sequence for smoothing the nodes y_i of C yields the exact sequence of $\mathbb{C}^*$-representations:

$$0 \to T_{(C_0,\mathbf{y})}\overline{\mathcal{M}}_{g,n} \to \mathrm{Ext}^1\left(\Omega_C\left(\sum_{i=1}^n x_i\right), \mathcal{O}_C\right) \to \bigoplus_{i=1}^n T_{y_i}C_0 \otimes T_{y_i}D_i \to 0.$$

Here, the first term represents deformations that preserve all the nodes, which corresponds to the Zariski tangent space of the moduli space of stable curves at the core. The final term represents the independent smoothing of each node y_i. Since the action on the target is trivial on C_0, $T_{y_i}C_0$ has weight zero. On the tail D_i, the map $f|_{D_i}\colon z \mapsto z^{\mu_i}$ must absorb the target rotation of weight 1, forcing the local coordinate z at y_i to have weight $1/\mu_i$. Therefore, $T_{y_i}D_i$ carries weight $1/\mu_i$.

Summarizing the analysis above, we have isolated the fixed and moving parts:

$$\begin{aligned} B_1^f &= \mathbb{C}_0^{\oplus n}; \quad B_1^m = 0; \\ B_4^f &= T_{(C_0,\mathbf{y})}\overline{\mathcal{M}}_{g,n}; \\ B_4^m &= \bigoplus_{i=1}^n T_{y_i}C_0 \otimes T_{y_i}D_i \\ &\cong \bigoplus_{i=1}^n (\mathbb{L}_i^\vee)_{(C_0,\mathbf{y})} \otimes \mathbb{C}_{1/\mu_i}. \end{aligned}$$

Recalling the definition of the ψ-classes on $\overline{\mathcal{M}}_{g,n}$ (Definition 2.4.24), the equivariant Euler class of B_4^m is computed via its first Chern class. Pulling back to $\overline{\mathcal{M}}_{g,n}$ via the canonical morphism $\rho\colon \overline{\mathcal{M}}_{g,n} \to F_r$, we obtain:

$$\begin{aligned} \rho^*\left(\frac{e_{\mathbb{C}^*}(B_1^m)}{e_{\mathbb{C}^*}(B_4^m)}\right) &= \frac{1}{\prod_{i=1}^n(\frac{u}{\mu_i} - \psi_i)} \\ &= \prod_{i=1}^n \frac{\mu_i}{u - \mu_i\psi_i}, \end{aligned} \tag{4.3.10}$$

where $\psi_i = c_1(\mathbb{L}_i) \in \mathrm{CH}^1(\overline{\mathcal{M}}_{g,n})$.

4.3.5. Contribution of B_2 and B_5. To calculate the contribution from B_2 and B_5, our analysis utilizes a technique from Graber and Pandharipande. By using the normalization exact sequence to resolve all the nodes of the domain curve C, we can "embed" the terms $\mathrm{H}^i(C, f^*T\mathbb{P}^1(-\sum \mu_i x_i))$ into a computable long exact sequence.

Recall that for the normalization map $\nu\colon \tilde{C} \to C$ of a nodal curve, and for any locally free sheaf $\mathcal{F}$ on C, there is a short exact sequence:

$$0 \to \mathcal{F} \to \nu_*\nu^*\mathcal{F} \to \bigoplus_{i=1}^n \mathcal{F}_{y_i} \to 0,$$

where $\{y_1, \dots, y_n\}$ is the set of nodes of C, and $\mathcal{F}_{y_i}$ is the fiber of $\mathcal{F}$ at the node y_i. This short exact sequence induces the following long exact sequence in cohomology:

$$0 \to \mathrm{H}^0(C, \mathcal{F}) \to \mathrm{H}^0(\tilde{C}, \nu^*\mathcal{F}) \to \bigoplus_{i=1}^n \mathcal{F}_{y_i} \to \mathrm{H}^1(C, \mathcal{F}) \to \mathrm{H}^1(\tilde{C}, \nu^*\mathcal{F}) \to 0. \tag{4.3.11}$$

Let $[f\colon (C,\mathbf{x}) \to \mathbb{P}^1]$ be a moduli point in F_r. By Proposition 4.3.1, the (partial) normalization of C naturally decomposes as

$$\nu\colon \tilde{C} = C_0 \sqcup \bigsqcup_{i=1}^{n} D_i \to C.$$

We apply the sequence (4.3.11) to this specific normalization and the locally free sheaf $\mathcal{F} \coloneqq f^*T\mathbb{P}^1(-\sum_{i=1}^{n} \mu_i x_i)$. Recognizing that $B_2 = \mathrm{H}^0(C,\mathcal{F})$ and $B_5 = \mathrm{H}^1(C,\mathcal{F})$, we obtain

$$0 \to B_2 \to \mathrm{H}^0(C_0,\mathcal{F}) \oplus \bigoplus_{i=1}^{n} \mathrm{H}^0(D_i,\mathcal{F}) \to \bigoplus_{i=1}^{n} \mathcal{F}_{y_i} \to B_5 \to \mathrm{H}^1(C_0,\mathcal{F}) \oplus \bigoplus_{i=1}^{n} \mathrm{H}^1(D_i,\mathcal{F}) \to 0.$$

We analyze these terms piece by piece. First, on the contracted core C_0, the map f is constant with image $0 \in \mathbb{P}^1$, so the pullback sheaf simplifies to $\mathcal{F}|_{C_0} \cong \mathcal{O}_{C_0} \otimes T_0\mathbb{P}^1$. Therefore,

$$\mathrm{H}^0(C_0,\mathcal{F}) \cong \mathrm{H}^0(C_0,\mathcal{O}_{C_0}) \otimes T_0\mathbb{P}^1 \cong \mathbb{C}_1; \tag{4.3.12}$$

$$\mathrm{H}^1(C_0,\mathcal{F}) \cong \mathrm{H}^1(C_0,\mathcal{O}_{C_0}) \otimes T_0\mathbb{P}^1 \cong (\mathbb{E}^\vee)_{(C_0,\mathbf{y})} \otimes \mathbb{C}_1, \tag{4.3.13}$$

where the second isomorphism follows from Serre duality and the definition of the Hodge bundle $\mathbb{E}$ (Example 2.3.41). Furthermore, at each node y_i, the map sends y_i to 0, so the fiber $\mathcal{F}_{y_i}$ is $T_0\mathbb{P}^1 \cong \mathbb{C}_1$.

Next, we calculate the cohomology on the rational tails D_i. As established in Proposition 4.3.1, $D_i \cong \mathbb{P}^1$ and the restriction $f|_{D_i}$ is given by $z \mapsto z^{\mu_i}$. The tangent bundle $T\mathbb{P}^1 \cong \mathcal{O}(2)$, so the pullback $f^*T\mathbb{P}^1$ has degree $2\mu_i$. Twisting by $-\mu_i x_i$ yields $\mathcal{F}|_{D_i} \cong \mathcal{O}_{\mathbb{P}^1}(\mu_i)$. The global sections of this bundle correspond to polynomials in z of degree up to μ_i. Since the target $\mathbb{C}^*$-action forces the local domain coordinate z to scale with weight $1/\mu_i$, the basis sections $\{1, z, z^2, \ldots, z^{\mu_i}\}$ carry equivariant weights $0, 1/\mu_i, 2/\mu_i, \ldots, 1$ respectively. Since the degree $\mu_i \geq 1 > -2$, the first cohomology group vanishes. Therefore,

$$\mathrm{H}^k(D_i,\mathcal{F}) \cong \begin{cases} \mathbb{C}_0 \oplus \bigoplus_{j=1}^{\mu_i} \mathbb{C}_{j/\mu_i}, & k = 0; \\ 0, & k = 1. \end{cases} \tag{4.3.14}$$

The long exact sequence allows us to compute the moving parts $[B_2^m] - [B_5^m]$ by taking the alternating sum of the moving components of the sequence:

$$\begin{aligned} [B_2^m] - [B_5^m] &= [\mathrm{H}^0(C_0,\mathcal{F})^m] + \sum_{i=1}^{n} [\mathrm{H}^0(D_i,\mathcal{F})^m] - \sum_{i=1}^{n} [\mathcal{F}_{y_i}^m] - [\mathrm{H}^1(C_0,\mathcal{F})^m] \\ &= \mathbb{C}_1 \oplus \bigoplus_{i=1}^{n} \bigoplus_{j=1}^{\mu_i} \mathbb{C}_{j/\mu_i} \ominus \mathbb{C}_1^{\oplus n} \ominus \left((\mathbb{E}^\vee)_{(C_0,\mathbf{y})} \otimes \mathbb{C}_1 \right). \end{aligned}$$

By taking the equivariant Euler class and recalling the definition of the λ-classes (Definition 2.4.24), we pull back to $\overline{\mathcal{M}}_{g,n}$ via ρ to obtain the final required ratio:

$$\rho^* \left(\frac{e_{\mathbb{C}^*}(B_5^m)}{e_{\mathbb{C}^*}(B_2^m)} \right) = \left(\prod_{i=1}^{n} \frac{\mu_i^{\mu_i}}{\mu_i!} \right) u^{n-d-1} \cdot \Lambda_g^\vee(u), \tag{4.3.15}$$

where $\Lambda_g^\vee(u) = \sum_{i=0}^{g} (-1)^i \lambda_i u^{g-i}$.

4.3.6. Putting them together. By (4.3.9), (4.3.10), and (4.3.15), we have the following relation in the localized Chow ring of $\overline{\mathcal{M}}_{g,n}$:

$$\frac{1}{\rho^*(e_{\mathbb{C}^*}(N_{F_r}^{\mathrm{vir}}))} = \left(\prod_{i=1}^n \frac{\mu_i^{\mu_i}}{\mu_i!}\right) \frac{\mu_1 \cdots \mu_n \cdot \Lambda_g^\vee(u) u^{n-d-1}}{\prod_{i=1}^n (u - \mu_i \psi_i)}. \tag{4.3.16}$$

On the other hand, we consider the pushforward $\rho_*\colon \mathrm{CH}_*(\overline{\mathcal{M}}_{g,n}) \to \mathrm{CH}_*(F_r)$. Recall from Proposition 4.3.1 that the fixed locus has the stacky structure $F_r \cong \overline{\mathcal{M}}_{g,n} \times \prod_{i=1}^n \mathsf{B}(\mathbb{Z}/\mu_i)$. The automorphism group of the rational tails yields a stabilizer of order $\prod_{i=1}^n \mu_i$. Consequently, the pushforward of the fundamental class picks up this fractional degree:

$$[F_r] = \frac{1}{\mu_1 \cdots \mu_n} \rho_* [\overline{\mathcal{M}}_{g,n}]. \tag{4.3.17}$$

We now combine (4.3.4), (4.3.16), and (4.3.17) to evaluate the integral. By the projection formula, we obtain

$$\begin{aligned} H_{g,\mu} &= \frac{1}{|\mathrm{Aut}(\mu)|} \int_{[F_r]} \frac{r! u^r}{e_{\mathbb{C}^*}(N_{F_r}^{\mathrm{vir}})} \\ &= \frac{1}{|\mathrm{Aut}(\mu)|} \frac{r!}{\mu_1 \cdots \mu_n} \int_{[\overline{\mathcal{M}}_{g,n}]} \frac{u^r}{\rho^*(e_{\mathbb{C}^*}(N_{F_r}^{\mathrm{vir}}))} \\ &= \frac{r!}{|\mathrm{Aut}(\mu)|} \prod_{i=1}^n \frac{\mu_i^{\mu_i}}{\mu_i!} \int_{[\overline{\mathcal{M}}_{g,n}]} \frac{\Lambda_g^\vee(u) \cdot u^{r+n-d-1}}{\prod_{i=1}^n (u - \mu_i \psi_i)}. \end{aligned}$$

To reach the final formula, we factor u^n out of the denominator to normalize the ψ-classes, rewriting $\prod_{i=1}^n (u - \mu_i \psi_i) = u^n \prod_{i=1}^n (1 - \mu_i \psi_i / u)$. This changes the overall power of u outside the fraction to u^{r-d-1}.

At this point, we rely on a dimension argument. The integration over the fundamental class of $\overline{\mathcal{M}}_{g,n}$ will annihilate any term in the formal expansion that does not have a codimension exactly equal to the dimension of the space, which is $3g - 3 + n$.

Recall that the degree of λ_i is i, and the degree of ψ_j is 1. The term $\Lambda_g^\vee(u) = \sum_{i=0}^g (-1)^i \lambda_i u^{g-i}$ carries u^{g-i} alongside λ_i, and the expansion of the denominator carries u^{-k_j} alongside $\psi_j^{k_j}$. The total power of u of the term $\lambda_i \psi_1^{k_1} \cdots \psi_n^{k_n}$ is therefore

$$(r - d - 1) + (g - i) - \sum_{j=1}^n k_j.$$

Because we only integrate terms where $i + \sum_{j=1}^n k_j = 3g - 3 + n$, and because the Riemann–Hurwitz formula dictates $r = 2g - 2 + d + n$, this total power of u becomes

$$(2g - 2 + d + n - d - 1) + g - (3g - 3 + n) = (2g - 3 + n) + g - 3g + 3 - n = 0.$$

Since the equivariant parameter u cancels out from every surviving term of top codimension, we can safely set $u = 1$ to simplify the integrand. This yields the final evaluation:

$$H_{g,\mu} = \frac{r!}{|\mathrm{Aut}(\mu)|} \prod_{i=1}^n \frac{\mu_i^{\mu_i}}{\mu_i!} \int_{[\overline{\mathcal{M}}_{g,n}]} \frac{\Lambda_g^\vee(1)}{\prod_{i=1}^n (1 - \mu_i \psi_i)}.$$

This finishes the proof of the ELSV formula (Theorem 4.0.5).

Bibliography

[ACG11] Enrico Arbarello, Maurizio Cornalba, and Phillip A. Griffiths. *Geometry of algebraic curves. Volume II*, volume 268 of *Grundlehren der mathematischen Wissenschaften*. Springer, Heidelberg, 2011. 59, 61, 62, 65, 72

[AF24] David Anderson and William Fulton. *Equivariant cohomology in algebraic geometry*, volume 210 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2024. 6, 85, 90

[AKL+25] Dhyan Aranha, Adeel A. Khan, Alexei Latyntsev, Hyeonjun Park, and Charanya Ravi. Virtual localization revisited. *Adv. Math.*, 479:Paper No. 110434, 53, 2025. 7

[Alp25] Jarod Alper. Evolution of stacks and moduli. *Notices Amer. Math. Soc.*, 72(11):1248–1261, 2025. 11

[Alp26] Jarod Alper. Stacks and moduli. Draft available online, January 2026. Version 2026-01-05, `https://sites.math.washington.edu/~jarod/moduli.pdf`. 2, 25, 27, 33, 34, 35, 37, 45, 59, 66, 67, 72, 85, 96

[Beh97] K. Behrend. Gromov-Witten invariants in algebraic geometry. *Invent. Math.*, 127(3):601–617, 1997. 108

[Beh04] K. Behrend. Cohomology of stacks. In *Intersection theory and moduli*, volume XIX of *ICTP Lect. Notes*, pages 249–294. Abdus Salam Int. Cent. Theoret. Phys., Trieste, 2004. 54, 55

[BF97] K. Behrend and B. Fantechi. The intrinsic normal cone. *Invent. Math.*, 128(1):45–88, 1997. 6, 82, 93, 97, 98, 99, 108

[Bor91] Armand Borel. *Linear algebraic groups*, volume 126 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1991. 84

[BX11] Kai Behrend and Ping Xu. Differentiable stacks and gerbes. *J. Symplectic Geom.*, 9(3):285–341, 2011. 30, 32

[Can25] Samir Canning. The tautological ring of $\overline{\mathcal{M}}_{g,n}$ is rarely Gorenstein. *Geom. Topol.*, 29(7):3905–3919, 2025. 104

[CKL17] Huai-Liang Chang, Young-Hoon Kiem, and Jun Li. Torus localization and wall crossing for cosection localized virtual cycles. *Adv. Math.*, 308:964–986, 2017. 7, 103

[CL06] Bohui Chen and An-Min Li. Symplectic virtual localization of Gromov–Witten invariants. `https://arxiv.org/abs/math/0610370`, 2006. 6

[CL24] Samir Canning and Hannah Larson. On the Chow and cohomology rings of moduli spaces of stable curves. *J. Eur. Math. Soc.*, 2024. 78

[CM16] Renzo Cavalieri and Eric Miles. *Riemann surfaces and algebraic curves*, volume 87 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 2016. 7, 9, 107

[DM69] P. Deligne and D. Mumford. The irreducibility of the space of curves of given genus. *Inst. Hautes Études Sci. Publ. Math.*, pages 75–109, 1969. 60

[EG98a] Dan Edidin and William Graham. Equivariant intersection theory. *Invent. Math.*, 131(3):595–634, 1998. 6, 82, 85, 87, 88

[EG98b] Dan Edidin and William Graham. Localization in equivariant intersection theory and the Bott residue formula. *Amer. J. Math.*, 120(3):619–636, 1998. 6, 82, 89

[EKS84] Allan L. Edmonds, Ravi S. Kulkarni, and Robert E. Stong. Realizability of branched coverings of surfaces. *Trans. Amer. Math. Soc.*, 282(2):773–790, 1984. 9

[ELSV01] Torsten Ekedahl, Sergei Lando, Michael Shapiro, and Alek Vainshtein. Hurwitz numbers and intersections on moduli spaces of curves. *Invent. Math.*, 146(2):297–327, 2001. 5, 106, 107

[EVdV81] D. Eisenbud and A. Van de Ven. On the normal bundles of smooth rational space curves. *Math. Ann.*, 256(4):453–463, 1981. 95

[Fab97] Carel F. Faber. A non-vanishing result for the tautological ring of $\mathcal{M}_g$. `https://arxiv.org/abs/math/9711219`, 1997. 104

[Fab99a] Carel Faber. Algorithms for computing intersection numbers on moduli spaces of curves, with an application to the class of the locus of Jacobians. In *New trends in algebraic geometry (Warwick, 1996)*, volume 264 of *London Math. Soc. Lecture Note Ser.*, pages 93–109. Cambridge Univ. Press, Cambridge, 1999. 105, 107

[Fab99b] Carel Faber. A conjectural description of the tautological ring of the moduli space of curves. In *Moduli of curves and abelian varieties*, volume E33 of *Aspects Math.*, pages 109–129. Friedr. Vieweg, Braunschweig, 1999. 104

[FGI+05] Barbara Fantechi, Lothar Göttsche, Luc Illusie, Steven L. Kleiman, Nitin Nitsure, and Angelo Vistoli. *Fundamental Algebraic Geometry: Grothendieck's FGA Explained*, volume 123 of *Mathematical Surveys and Monographs*. American Mathematical Society, 2005. 20, 22, 23, 24

[FM98] Barbara Fantechi and Marco Manetti. Obstruction calculus for functors of Artin rings. I. *J. Algebra*, 202(2):541–576, 1998. 14, 15

[For81] Otto Forster. *Lectures on Riemann surfaces*, volume 81 of *Graduate Texts in Mathematics*. Springer-Verlag, New York-Berlin, 1981. 5, 7

[FP97] W. Fulton and R. Pandharipande. Notes on stable maps and quantum cohomology. In *Algebraic geometry—Santa Cruz 1995*, volume 62, Part 2 of *Proc. Sympos. Pure Math.*, pages 45–96. Amer. Math. Soc., Providence, RI, 1997. 79, 80

[FP00] C. Faber and R. Pandharipande. Hodge integrals and Gromov-Witten theory. *Invent. Math.*, 139(1):173–199, 2000. 5

[FP02] B. Fantechi and R. Pandharipande. Stable maps and branch divisors. *Compositio Math.*, 130(3):345–364, 2002. 5, 107, 109

[FP05] C. Faber and R. Pandharipande. Relative maps and tautological classes. *J. Eur. Math. Soc. (JEMS)*, 7(1):13–49, 2005. 114

[FP13] C. Faber and R. Pandharipande. Tautological and non-tautological cohomology of the moduli space of curves. In *Handbook of moduli. Vol. I*, volume 24 of *Adv. Lect. Math. (ALM)*, pages 293–330. Int. Press, Somerville, MA, 2013. 77, 78

[Ful69] William Fulton. Hurwitz schemes and irreducibility of moduli of algebraic curves. *Ann. of Math. (2)*, 90:542–575, 1969. 76, 77

[Ful98] William Fulton. *Intersection theory*, volume 2 of *Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge. A Series of Modern Surveys in Mathematics*. Springer-Verlag, Berlin, second edition, 1998. 12, 56, 57, 58, 78, 87, 89, 93, 94, 95, 100, 103

[GJ99] I. P. Goulden and D. M. Jackson. The number of ramified coverings of the sphere by the double torus, and a general form for higher genera. *J. Combin. Theory Ser. A*, 88(2):259–275, 1999. 105

[GP99] T. Graber and R. Pandharipande. Localization of virtual classes. *Invent. Math.*, 135(2):487–518, 1999. 4, 5, 7, 82, 93, 100, 101, 102, 103, 107, 110, 111, 114, 118

[GR84] Hans Grauert and Reinhold Remmert. *Coherent analytic sheaves*, volume 265 of *Grundlehren der mathematischen Wissenschaften*. Springer-Verlag, Berlin, 1984. 2

[GV03] Tom Graber and Ravi Vakil. Hodge integrals and Hurwitz numbers via virtual localization. *Compositio Math.*, 135(1):25–36, 2003. 5, 107, 110, 111

[GV05] Tom Graber and Ravi Vakil. Relative virtual localization and vanishing of tautological classes on moduli spaces of curves. *Duke Math. J.*, 130(1):1–37, 2005. 7, 111, 114

[Har77] Robin Hartshorne. *Algebraic geometry*, volume No. 52 of *Graduate Texts in Mathematics*. Springer-Verlag, New York-Heidelberg, 1977. 2, 3, 4, 17, 25, 37, 53, 54, 61, 62, 64, 68, 69, 72, 80, 110, 115

[HM98] Joe Harris and Ian Morrison. *Moduli of curves*, volume 187 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1998. 59, 67, 69

[Hur01] A. Hurwitz. Ueber die Anzahl der Riemann'schen Flächen mit gegebenen Verzweigungspunkten. *Math. Ann.*, 55(1):53–66, 1901. 1, 105

[Ill71] Luc Illusie. *Complexe cotangent et déformations. I*, volume Vol. 239 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin-New York, 1971. 96

[Ill72] Luc Illusie. *Complexe cotangent et déformations. II*, volume Vol. 283 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin-New York, 1972. 96

[IT92] Y. Imayoshi and M. Taniguchi. *An introduction to Teichmüller spaces*. Springer-Verlag, Tokyo, 1992. 11

[Kaz09] M. Kazarian. KP hierarchy for Hodge integrals. *Adv. Math.*, 221(1):1–21, 2009. 107

[Kle87] Steven L. Kleiman. Intersection theory and enumerative geometry: a decade in review. In *Algebraic geometry, Bowdoin, 1985 (Brunswick, Maine, 1985)*, volume 46, Part 2 of *Proc. Sympos. Pure Math.*, pages 321–370. Amer. Math. Soc., Providence, RI, 1987. 79

[Kon95] Maxim Kontsevich. Enumeration of rational curves via torus actions. In *The moduli space of curves (Texel Island, 1994)*, volume 129 of *Progr. Math.*, pages 335–368. Birkhäuser Boston, Boston, MA, 1995. 79

[Kre99] Andrew Kresch. Cycle groups for Artin stacks. *Invent. Math.*, 138(3):495–536, 1999. 59, 99

[Li01] Jun Li. Stable morphisms to singular schemes and relative stable morphisms. *J. Differential Geom.*, 57(3):509–578, 2001. 6, 111, 114

[Li02] Jun Li. A degeneration formula of GW-invariants. *J. Differential Geom.*, 60(2):199–293, 2002. 6, 111, 114, 115

[Loo95] Eduard Looijenga. On the tautological ring of $\mathcal{M}_g$. *Invent. Math.*, 121(2):411–419, 1995. 104

[LR01] An-Min Li and Yongbin Ruan. Symplectic surgery and Gromov–Witten invariants of Calabi–Yau 3-folds. *Invent. Math.*, 145(1):151–218, 2001. 5, 111

[Mil17] J. S. Milne. *Algebraic groups*, volume 170 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2017. The theory of group schemes of finite type over a field. 89

[Mir95] Rick Miranda. *Algebraic curves and Riemann surfaces*, volume 5 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 1995. 7

[Mum83] David Mumford. Towards an enumerative geometry of the moduli space of curves. In *Arithmetic and geometry, Vol. II*, volume 36 of *Progr. Math.*, pages 271–328. Birkhäuser Boston, Boston, MA, 1983. 77, 78

[Noo12] Behrang Noohi. Homotopy types of topological stacks. *Adv. Math.*, 230(4-6):2014–2047, 2012. 30, 54, 55

[Ols16] Martin Olsson. *Algebraic spaces and stacks*, volume 62 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 2016. 36

[OP06a] A. Okounkov and R. Pandharipande. The equivariant Gromov-Witten theory of $\mathbf{P}^1$. *Ann. of Math. (2)*, 163(2):561–605, 2006. 4

[OP06b] A. Okounkov and R. Pandharipande. Gromov-Witten theory, Hurwitz theory, and completed cycles. *Ann. of Math. (2)*, 163(2):517–560, 2006. 4

[OP06c] A. Okounkov and R. Pandharipande. Virasoro constraints for target curves. *Invent. Math.*, 163(1):47–108, 2006. 4

[OP09] A. Okounkov and R. Pandharipande. Gromov–Witten theory, Hurwitz numbers, and matrix models. In *Algebraic geometry—Seattle 2005. Part 1*, volume 80, Part 1 of *Proc. Sympos. Pure Math.*, pages 325–414. Amer. Math. Soc., Providence, RI, 2009. 5, 107, 110

[Pan18] Rahul Pandharipande. A calculus for the moduli space of curves. In *Algebraic geometry: Salt Lake City 2015*, volume 97.1 of *Proc. Sympos. Pure Math.*, pages 459–487. Amer. Math. Soc., Providence, RI, 2018. 78, 104

[Pet16] Dan Petersen. Tautological rings of spaces of pointed genus two curves of compact type. *Compos. Math.*, 152(7):1398–1420, 2016. 104

[PT14] Dan Petersen and Orsola Tommasi. The Gorenstein conjecture fails for the tautological ring of $\overline{\mathcal{M}}_{2,n}$. *Invent. Math.*, 196(1):139–161, 2014. 4, 104

[Ray70] Michel Raynaud. *Faisceaux amples sur les schémas en groupes et les espaces homogènes*, volume Vol. 119 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin-New York, 1970. 36

[Rie57] B. Riemann. Theorie der Abel'schen Functionen. *J. Reine Angew. Math.*, 54:115–155, 1857. 3, 11

[Rie13] Bernhard Riemann. *Bernhard Riemann "Über die Hypothesen, welche der Geometrie zu Grunde liegen"*. Klassische Texte der Wissenschaft. Springer Spektrum, 2013. 3

[Ser06] Edoardo Sernesi. *Deformations of algebraic schemes*, volume 334 of *Grundlehren der mathematischen Wissenschaften*. Springer-Verlag, Berlin, 2006. 14, 16, 37, 52, 75, 81, 108

[SGA03] *Revêtements étales et groupe fondamental (SGA 1)*, volume 3 of *Documents Mathématiques (Paris)*. Société Mathématique de France, Paris, 2003. 3

[Sta26] The Stacks project authors. The stacks project. `https://stacks.math.columbia.edu`, 2026. 26, 31, 37, 38, 39, 40, 43, 44, 47, 48, 50, 57, 66, 67, 68

[Tod21] Yukinobu Toda. *Recent progress on the Donaldson-Thomas theory—wall-crossing and refined invariants*, volume 43 of *SpringerBriefs in Mathematical Physics*. Springer, Singapore, [2021] ©2021. 100

[Vak06] Ravi Vakil. Murphy's law in algebraic geometry: badly-behaved deformation spaces. *Invent. Math.*, 164(3):569–590, 2006. 81

[Vak08] R. Vakil. The moduli space of curves and Gromov-Witten theory. In *Enumerative invariants in algebraic geometry and string theory*, volume 1947 of *Lecture Notes in Math.*, pages 143–198. Springer, Berlin, 2008. 112

[Vis89] Angelo Vistoli. Intersection theory on algebraic stacks and on their moduli spaces. *Invent. Math.*, 97(3):613–670, 1989. 12, 55, 57, 58

[WLZ+25] Yiru Wang, Bingqian Li, Yi Zhou, Zhiqiang Wei, Yu Ye, Yiqian Shi, and Bin Xu. The Hurwitz existence problem and the prime-degree conjecture: A computational perspective. https://arxiv.org/abs/2512.06545, 2025. 9