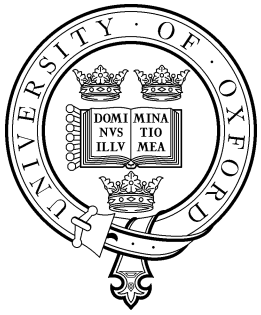
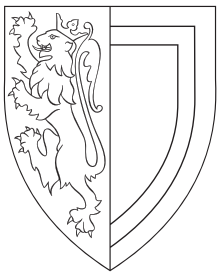
Moduli Spaces of Unstable Curves and Sheaves via Non-Reductive Geometric Invariant Theory





Joshua James Jackson

Balliol College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity Term 2018

Abstract

Many moduli problems in algebraic geometry can be posed using Geometric Invariant Theory (GIT). However, as with all such tools, if we are to have any hope of obtaining a well-behaved moduli space, certain objects, which are in whatever sense degenerate, must be ruled 'unstable' and left out of the classification. The goal of this thesis is to show how new results in non-reductive GIT may be used to construct moduli spaces that parameterise these unstable objects. Once a moduli problem is posed using GIT, this is done by studying the associated instability stratification and using non-reductive GIT to take quotients of the unstable strata, effectively splitting the moduli problem up into manageable pieces. We apply this method in two examples.

First, we consider moduli spaces of coherent sheaves over a projective scheme. The instability stratification in this case is closely related to the stratification by Harder-Narasimhan type. However, in order to apply the relevant theorems in non-reductive GIT, in fact one must perform a blow-up process analogous to partial desingularisation in the reductive case. This corresponds to refining the stratification by Harder-Narasimhan type. We give a sheaf-theoretic interpretation of this refinement, and use it to construct moduli spaces of such sheaves, in the case of Harder-Narasimhan length 2. We also include some results towards the length 3 case, and some remarks on generalisation to arbitrary lengh.

Our second application is moduli of projective algebraic curves. Here it is not clear a priori what the instability stratification should mean intrinsically, which leads us to define the notion of Rosenlicht-Serre type for curves. We show that in the case of a curve with only unibranch singularities this data essentially coincides with the semigroup of the singularities, and give a construction of moduli spaces of unibranch curves of fixed Rosenlicht-Serre type.

Declaration

The results in this thesis on Variation of Non-Reductive Geometric Invariant Theory are from a collaboration with Gergely Bérczi and Frances Kirwan. This material, presented in Section 3.4 is taken from [BJK], which will appear in a future edition of Surveys in Differential Geometry.

The results that appear in Chapter 5 on Moduli Spaces of Sheaves of Fixed Harder-Narasimhan Type (i.e. unstable sheaves) are from a collaboration with Gergely Bérczi, Victoria Hoskins and Frances Kirwan. They are taken from a preprint in preparation, [BHJK], entitled 'Quotients of unstable subvarieties and moduli spaces of sheaves of fixed Harder-Narasimhan type, II', though the analysis appearing in Section 5.3.2 onwards was written by me for this thesis. Parts of this paper have also been used in Chapter 4.

To the best of my knowledge, all other results which are not original work by the author have been duly cited.

Joshua James Jackson

June 2018

Acknowledgements

First, I offer my profound and humble thanks to Frances Kirwan. If this thesis is a testament to anything, it is your extraordinary skill and dedication as a supervisor, and, I can only assume, your indefatigable patience. I am also greatly indebted to Gergely Bérczi, for his guidance and support throughout my time as a graduate student; to Victoria Hoskins, whose thesis inspired me to pursue this project; to Eloise Hamilton, who provided the appendix, and whose conversation and work on curves was very useful in the preparation of this thesis; and to Vicky Neale, who taught me what I know of how to teach, and offered much helpful and sympathetic advice throughout this whole process. I would also like to acknowledge the work of the many educators of every sort, who have encouraged me throughout my life to seek a better understanding of the world.

Thanks are due also to Alex Torzewski and David Watson, from whom I have learned a great deal over the years. Our time at the blackboard in Warwick, and later in Cambridge, probably did more to put me on the path to a DPhil in Mathematics than anything else.

I would like to thank my friends from St Catharine's College, Cambridge; especially Joe Harper, for taking in a poor mathmo who didn't quite understand how to be a person, and is still working on it. That meant more to me than you will probably ever understand, and is proof, if proof were needed, that it is possible to be simultaneously sarcastic and kind.

This thesis bears the crest of Balliol College - a place that would by now be nothing more to me than a set of seldom-visited buildings if it were not for those I met there, who have helped make Oxford my home. To the best of my hazy recollection, in order of appearance: Laure Bouwyn, Sarah Reeve, Duncan Fraser, Gillian Bolsover, Jess Hutchens, Ben Pope, Dave Miller, Mirela Ivanova, and, our latest addition, Ed Coghill. You are all, in your gloriously disparate and idiosyncratic ways, frighteningly brilliant people, and, more than that, the most wonderful friends. It is to you that this thesis is dedicated.[1]

To my parents: thank you for always, without hesitation, doing everything you could. Nothing could have replaced your unconditional love and unwavering faith in me. Without that I do not know who I would be, and I regret that I have nothing better with which to repay it.

[1]While I am at it, it would be remiss of me not to acknowledge many happy hours spent in the company of Hanash, Aenigma, Wolf, Storm, Orla, Eascane, Cawl, Bryn, and Ranger. It has been a privilege and a joy to go on this journey with you. May you all succeed in your quests.

Finally, and above all, Elspeth. In a life filled with golden moments of absurd good luck, meeting you has been by far the greatest.

To my friends, the degenerate and the singular

Contents

1 Introduction **1**

2 Reductive GIT and its Variation **6**

2.1 Mumford's GIT 7

2.2 Variation of Reductive GIT 9

2.3 Instability Stratification 10

3 Non-Reductive GIT and its Variation **16**

3.1 First steps in Non-Reductive GIT 16

3.1.1 Motivations 17

3.1.2 Attempts to glue affines 19

3.1.3 Transferral to a reductive action 22

3.2 Example: ordered points on $\mathbb{P}^1$ up to translation 24

3.2.1 The Linearisation 25

3.2.2 The 1-parameter subgroups 25

3.2.3 The Embedding and the weights 26

3.2.4 Calculation of the semistable locus 27

3.3 The $\hat{U}$ Theorems 30

3.3.1 Statement of the theorems 30

3.3.2 The $\hat{U}$ blow-up process 34

3.4 Variation of Non-Reductive GIT 39

3.4.1 Change of distinguished $\mathbb{G}_m$ . . . 39
3.4.2 Change of linearisation . . . 43
3.4.3 When semistability does not coincide with stability . . . 44
3.4.4 Summary of NRVGIT results . . . 45
3.4.5 Points and lines in $\mathbb{P}^2$ . . . 47

4 Moduli Spaces of Unstable Objects 49

4.1 Reductive quotients of the unstable strata . . . 50
4.1.1 Toy Example: unordered points on $\mathbb{P}^1$ . . . 53
4.2 Quotients using Non-Reductive GIT . . . 56
4.2.1 When we have condition $(\star)$ . . . 58
4.2.2 Toy Example revisited . . . 59
4.2.3 Comparison of the two applications . . . 61

5 Moduli Spaces of Sheaves of Fixed Harder-Narasimhan type 63

5.1 Moduli Spaces of Semistable Sheaves . . . 64
5.1.1 Simpson's construction . . . 64
5.2 Stratifications for Sheaves . . . 70
5.2.1 Harder-Narasimhan Stratification . . . 70
5.2.2 Instability stratification for sheaves . . . 71
5.3 Moduli of Unstable Sheaves . . . 74
5.3.1 Stabilisers in U_τ . . . 75
5.3.2 The case $\ell = 2$ . . . 78
5.3.3 Towards $\ell = 3$ . . . 84
5.3.4 Remarks on $\ell > 3$ . . . 88

6 Moduli Spaces of Singular Curves 89

6.1 Constructions for Stable Curves . . . 89
6.1.1 GIT construction of $\overline{M}_g$ . . . 89

6.1.2 Baldwin-Swinarski construction of $\overline{M}_{g,n}$ 92
6.2 Preliminary Study of Singular Curves 92
6.2.1 Rosenlicht-Serre Construction 92
6.2.2 Local study of curve singularities 95
6.2.3 Unibranch singularities 99
6.3 Moduli of Curves of fixed Rosenlicht-Serre type 102
6.3.1 From local to global 103
6.3.2 The GIT problem 109
6.3.3 The associated graded curve 116
6.3.4 Construction of $M_{g,\sigma}$ 118
6.3.5 Remarks on generalisations 123

Appendix Diagrams and Algorithm for Projective Completions 125

1 Diagrams for GIT, reductive and non-reductive 125
1.1 Classical GIT for reductive groups 125
1.2 GIT for graded extensions of unipotent groups 127
1.3 GIT for internally graded unipotent radical 130
2 Algorithm for internally graded quotients 132
2.1 Replacement algorithm 133
2.2 Blow-up algorithm 136

Bibliography I

Chapter 1

Introduction

This thesis has two main aims. Firstly, to present a way of dealing with moduli problems in algebraic geometry, which we call 'Moduli Spaces of Unstable Objects'. Secondly, to explore the application of this method to the two most important examples. The central idea is far from being a new one,[1] and may be stated very simply: we treat instability itself as just another discrete invariant. By doing this, one can expand the traditional classifications of algebro-geometric objects, to include *all* objects of a particular type, rather than just the so-called 'stable' ones. Making this notion precise, however, and actually carrying it out in practice, requires some technical tools that have only recently become available, with the advent of powerful new methods in Geometric Invariant Theory (GIT): specifically, a viable and widely applicable theory for non-reductive groups, in the form of the $\hat{U}$ Theorems.

We begin in Chapter 2 by reviewing the key features of Mumford's classical GIT for reductive groups, first presented in [MFK94], with a focus on the case in which a reductive linear algebraic group G acts linearly on a projective variety X. As well as being used extensively throughout this thesis, these are the properties that one wishes to carry over to the non-reductive setting. We make particular mention of

[1] In its essence it goes back at least as far as [Kir84].

so-called Variation of GIT ([DH98] [Tha96]), which studies the dependence of the GIT quotient on the choice of linearisation, and of the instability stratification of the unstable locus. This stratification, which, morally, sorts the points of X according to how unstable they are for the action of G, is a key ingredient in what follows. In effect, the first stage in constructing moduli spaces of unstable objects is to find an intrinsic moduli interpretation of this stratification.

In Chapter 3 we give an overview of the current state of non-reductive GIT. Beginning by explaining some of the difficulties involved, and the differences from the reductive case, we then go on to introduce the $\hat{U}$ Theorems themselves, which allow one, under certain reasonable conditions, to recover essentially all of the things that make Mumford's GIT so useful: one has a semistable locus, readily calculable by a Hilbert-Mumford criterion, with a projective categorical quotient, restricting to a geometric quotient on a certain subset. In addition, we include in Section 3.4 some original results showing that there is also a good analogue of Variation of GIT in this setting.

We bring all of this together in Chapter 4, in which we outline the general point of view of moduli of unstable objects. A quick way for a reader who is already familiar with reductive GIT to get a rough sense of what this thesis is about is to read Section 2.3, followed by this chapter. We show how the $\hat{U}$ Theorems may be used to quotient unstable strata in the instability stratification of a reductive GIT problem, and demonstrate how this improves on a previous method that uses reductive GIT. We illustrate this with a simple example: that of n-tuples of unordered points on the projective line.

The remainder of the thesis treats two applications. Chapter 5 considers the moduli

of coherent sheaves on a projective base scheme, taking as our GIT set-up Simpson's construction of moduli spaces of semistable sheaves, given in [Sim94]. In this case, the resultant instability stratification of the Quot scheme is well understood, and known to be closely related to the stratification by Harder-Narasimhan type; this is as one would expect, since both stratifications in effect measure the failure of a sheaf to be semistable. By work of Hoskins and Kirwan ([HK12]), one may associate to each Harder-Narasimhan type (that is, to each tuple of polynomials giving the Hilbert polynomials of the successive quotients in the Harder-Narasimhan filtration of the sheaf) a certain index in the instability stratification, such that each Harder-Narasimhan stratum is a closed subscheme of its associated instability stratum. The challenge is to use the $\hat{U}$ Theorems to perform quotients of these strata, in order to obtain moduli spaces of sheaves of fixed Harder-Narasimhan type. This turns out to be non-trivial: due to the presence of certain unipotent automorphisms, the most straightforward of the $\hat{U}$ Theorems do not apply, and one must perform a certain blow up according to the dimensions of stabilisers in the derived series of the relevant unipotent radical. This corresponds to refining the stratification by Harder-Narasimhan type, and one must then find a way to interpret this refinement in a sheaf-theoretic way. We do this for sheaves whose Harder-Narasimhan filtrations are of length 2, and give some results towards the length 3 case. We expect that one could generalise to arbitrary length using much the same ideas.

Chapter 6 of this thesis considers the moduli of projective algebraic curves. Here the appropriate notion of stability is the usual one: at worst nodal singularities and finite automorphism group. It is not clear, however, what the analogue of Harder-Narasimhan type should be: if we have a point in a Hilbert scheme corresponding to a curve embedded by a complete linear series, then a maximally destabilising (or Kempf) one parameter subgroup will give us a weighted filtration of the global

sections of the line bundle, but there is no sensible 'filtration by subcurves' to mimic the sheaves case exactly. What we can make sense of is the 'associated graded curve', and this leads us to define the notion of Rosenlicht-Serre type. This is a discrete invariant that one can associate to a singular curve, which plays the role that is played in the sheaves case by the Harder-Narasimhan type. We show that in the case of unibranch singularities this data is equivalent to the data of (the weights of) the semigroups of values of the curve's singularities, and give a construction of the quasiprojective moduli space of unibranch curves of fixed genus and Rosenlicht-Serre type. Throughout, we give indications of how this may be generalised to the multibranch case.

Conventions, Notation and Standing Assumptions

Our ground field is $\mathbb{C}$, though practically everything we say is valid over any algebraically closed field of characteristic zero. Throughout, unless otherwise specified, let X be a irreducible projective scheme carrying an ample line bundle L. When X is acted on by a linear algebraic group, we say that the action is linear, or linearised, if we have chosen a linearisation: that is, a lift of the action on X to the total space of L. We will denote linearisations themselves with cursive scripts, i.e. $\mathscr{L}$ is a linearisation with underlying bundle L. Unless otherwise stated, we will denote by G a reductive linear algebraic group, and by H an arbitrary linear algebraic group, which may be non-reductive. We fix a Levi factor R of H, so that $H = U \rtimes R$ where U is the unipotent radical of H. When given without qualification, point means closed point and sheaf means coherent sheaf. Following a now firmly established convention, we use the word 'unstable' to refer to points, sheaves, curves, &c, that are not *semi*stable. In some cases, where we wish to make the distinction explicit, the symbol $\simeq$ will relate objects that are non-canonically isomorphic, and $\cong$ will relate objects where there is a natural choice of isomorphism.

Chapter 2

Reductive GIT

> *The purpose of this book is to study two related problems: when does an orbit space of an algebraic scheme acted on by an algebraic group exist? And to construct moduli schemes for various types of algebraic objects. The second problem appears to be, in essence, a special and highly non-trivial case of the first problem.*
>
> David Mumford, *Geometric Invariant Theory*

In this chapter we review some of the key features of classical, i.e. reductive, Geometric Invariant Theory. The standard reference is [MFK94]. We begin with a few fundamental properties of reductive GIT quotients, paying particular attention to the Hilbert-Mumford criterion and the well-behaved dependence on the choice of linearisation. The latter is known as Variation of GIT, and we demonstrate in Section 3.4 how it may be carried over to the non-reductive setting. Finally, we recall the instability stratification,[1] which we can obtain from a linearised action of a reductive linear algebraic group on a projective scheme. In later chapters we will show how non-reductive GIT may be used to take quotients of the strata that appear in such stratifications, leading to moduli spaces of unstable objects.

[1] This stratification goes by many other names in the literature, usually formed by some combination of the names Hesselink, Kempf, Kirwan and Ness.

2.1 Mumford's GIT

When a reductive linear algebraic group G acts on a projective scheme X, Mumford's Geometric Invariant Theory allows us to construct quotients, in the following sense. We need to make a choice of linearisation $\mathscr{L}$; that is, an ample line bundle L on X, and a lift of the G-action to L. Then, by a theorem of Nagata [Nag63], the algebra $\bigoplus_{n=0}^{\infty} H^0(X, L^{\otimes n})^G$ of invariant sections is a finitely generated graded algebra. We call the associated projective scheme

$$X /\!\!/_{\mathscr{L}} G \coloneqq \operatorname{Proj} \bigoplus_{n=0}^{\infty} H^0(X, L^{\otimes n})^G$$

the GIT quotient of X by G with respect to $\mathscr{L}$. Further, there are open subschemes

$$X^s(\mathscr{L}) \subseteq X^{ss}(\mathscr{L}) \subseteq X,$$

respectively called the stable and semistable loci, such that the inclusion of the invariants into the algebra of sections induces a good quotient

$$X^{ss}(\mathscr{L}) \to X /\!\!/_{\mathscr{L}} G,$$

which restricts to a geometric quotient

$$X^s(\mathscr{L}) \to X^s(\mathscr{L})/G$$

onto an open subscheme of the GIT quotient [MFK94]. For the reader's convenience, we give the definitions of good and geometric quotient below.

2.1.0.1 Definition. Let X be a scheme acted on by an algebraic group G, and let $p : X \to Y$ be a G-invariant morphism of schemes. We say that $p : X \to Y$ is:

1. A *categorical quotient* if for every G-invariant map of schemes $f : X \to Z$, there is a unique morphism $g : Y \to Z$ making the triangle commute.
2. A *good quotient* if the following hold:
   1) p is surjective,
   2) p is affine,
   3) for all open $U \subseteq Y$, the map $\mathscr{O}_Y(U) \to \mathscr{O}_X(p^{-1}U)$ is an isomorphism onto the invariant sections, $\mathscr{O}_X(p^{-1}U)^G$,

4) if $V \subset X$ is closed and G-invariant, then $p(V) \subset Y$ is closed,

5) if V_1, $V_2 \subseteq X$ are closed disjoint G-invariant sets then $p(V_1)$ and $p(V_2)$ are disjoint.

3. A *geometric quotient*, if it is a good quotient and the fibre over each closed point of Y is a single orbit.

2.1.0.2 Remark. We note that geometric quotients are good, and good quotients are categorical, but the reverse implications do not hold.

2.1.0.3 Remark. When considering how the choice of linearisation affects the GIT quotient it is important to note that, since good quotients are in particular categorical, if two linearisations yield the same semistable locus then their associated quotients are canonically isomorphic.

Thus classical GIT provides us with G-invariant open subschemes of the scheme X and gives us the best possible quotients of those open subschemes by the G-action. A crucial feature of GIT is that those quotients are not too hard to compute: in order to determine the GIT quotient and its geometric restriction we need only determine the stable and semistable loci, and the following tells us that these are determined by essentially combinatorial data.

2.1.0.4 Theorem. *[MFK94]* ***(Hilbert-Mumford Criterion)*** *Let $T \subseteq G$ be a maximal torus, and choose coordinates that diagonalise the action of T.*

1) Stability and semistability for the T-action on X are determined purely combinatorially as follows. Take $x \in X$ and let $\Delta_x \subset \mathfrak{t}^*$ be the convex hull of the weights of the T-action corresponding to nonzero coordinates of x, and let Δ_x° be its interior. Then the semistable and loci are described as follows.

$$x \in X^{T-ss}(\mathscr{L}) \iff 0 \in \Delta_x;$$

$$x \in X^{T-s}(\mathscr{L}) \iff 0 \in \Delta_x^\circ.$$

2) A point $x \in X$ is semistable (resp. stable) for the G action with respect to $\mathscr{L}$ if and only if every point of its G-orbit is semistable (resp. stable) for T. That is, we have

$$X^{G-(s)s}(\mathscr{L}) = \bigcap_{g \in G} g \cdot X^{T-(s)s}(\mathscr{L}).$$

2.2 Variation of Reductive GIT

The study of how the GIT quotient depends on the linearisation, known as *Variation of GIT* or VGIT, was initiated by Thaddeus in [Tha96] and by Dolgachev and Hu in [DH98].

Let $\mathrm{Pic}^G(X)$ denote the set of G-linearised line bundles on X up to the natural notion of isomorphism. The tensor product of two G-linearised line bundles and the dual of a G-linearised line bundle both have canonical G-linearisations, and hence this set has an abelian group structure. The following definition comes from [Tha96].

2.2.0.1 Definition. We say that two linearisations, $\mathscr{L}_1$, $\mathscr{L}_2$, are G-algebraically equivalent if there is a connected scheme S with points s_1,$s_2 \in S$, and a linearisation $\mathscr{L}_S$ of the G-action on $S \times X$ induced from the second factor, such that $\mathscr{L}_S|_{s_i} \cong \mathscr{L}_i$ for $i = 1, 2$.

Thaddeus proved in [Tha96] that the semistable locus in X, and hence the GIT quotient of X by G, is invariant under the equivalence relation of G-algebraic equivalence, so we may pass to the quotient of $\mathrm{Pic}^G(X)$ by this equivalence relation, denoted $\mathrm{NS}^G(X)$. It follows from the Hilbert-Mumford criterion that the GIT quotient is unchanged if one replaces a given G-linearisation $\mathscr{L}$, of the line bundle L, with the induced linearisation $\mathscr{L}^n$ on $L^{\otimes n}$ for $n > 0$. So we may consider the rational vector space $NS^G(X) \otimes \mathbb{Q}$, which is shown to be finite-dimensional. This is the space of rational linearisations.

2.2.0.2 Definition. Let $A^G_{\mathbb{Q}}$ be the cone in $NS^G(X) \otimes \mathbb{Q}$ consisting of rational linearisation classes for which the underlying algebraic equivalence class of bundle is ample. Denote by $C^G \subset A^G_{\mathbb{Q}}$ the set of G-effective rational ample linearisations; that is, those linearisations in $A^G_{\mathbb{Q}}$ for which there is a semistable point in X.

The goal of VGIT is to understand how the quotient changes as we move around in this parameter space. This is described in [DH98] using a so-called wall-and-chamber structure on C^G: a subdivision of C^G into convex chambers by homogeneous walls which are locally polyhedral of codimension one in C^G. If we assume that generic points of X have 0-dimensional stabilisers, then a linearisation $\mathscr{L}$ lies on some wall if it has a strictly semistable point: that is, there is a point of X that is semistable but not stable with respect to $\mathscr{L}$. A connected component of the set of linearisations not lying on any wall is a *chamber*. A *cell* is a connected component of the set of points

of a given wall W that do not lie on any other wall except those walls containing the whole of W.

The main features of the dependence may be summarised as follows.

2.2.0.3 Theorem. *[DH98] Assume that there exists a linearisation such that semistability equals stability. Then:*

1) *The wall-and-chamber structure of C^G is finite, in the sense that there are only finitely many chambers, walls and cells, and polyhedral, in the sense that the closure of a chamber is a rational polyhedral cone in the interior of C^G.*

2) *Two linearisations in the same chamber give the same semistable locus, and hence the same GIT quotient, while the induced (rational) ample line bundle on the quotient varies in the chamber. The set of linearisations giving a particular semistable locus is either a chamber or a union of cells in the same wall.*

3) *As a wall is crossed, the quotients for the wall and for the chambers on either side are related by an explicit birational transformation, called a Thaddeus flip.*

2.2.0.4 Remark. The details of the birational transformation relating the quotients on either side of a wall will not concern us here. The definition may be found in [Tha96], where it is referred to simply as a flip.

2.2.0.5 Example. A simple example of this structure is that of $\mathbb{G}_m$ acting linearly on $\mathbb{P}^n$. Then C^G is the half-plane in $\mathbb{Q}^2$, with the direction perpendicular to $\mathrm{NS}\otimes\mathbb{Q} \cong \mathbb{Q}$ corresponding to a choice of rational character. Choosing coordinates $x_0, .., x_n$ to diagonalise the action, and letting $\alpha_i \in \mathbb{Z}$ denote the weight of x_i, we see that the chambers are bounded by rays passing through $(1, \alpha_i)$.

2.3 Instability Stratification

The results of Section 2.1, which of course are well known, may be interpreted as saying that, while we cannot usually expect to find a well-behaved quotient of the whole of a given G-scheme by its group action, GIT gives us a locus that does have a quotient, and an easy combinatorial way to compute it. Slightly less well known is the fact that this is not all we are given: in fact there is a stratification of the whole G-scheme into locally closed G-invariant strata, whose open piece is the semistable locus, which roughly speaking sorts points according to how unstable they are. In other words, instability in GIT is not a binary: a certain normalised version of the

Hilbert-Mumford function, together with some results of Kempf and others, allows us to see that instability is *structured*. From our point of view, this stratification tells us how we should split a given moduli problem up in order to have a hope of getting reasonable moduli spaces for unstable objects. We will now define this stratification, which will be crucial to what follows. For this, first recall the definition of the Hilbert–Mumford function.

2.3.0.1 Definition. Let $X \subset \mathbb{P}^r$ be a projective scheme and G a reductive linear algebraic group acting on X. Choosing a linearisation of this action with respect to a very ample line bundle, we may assume the action is via a representation

$$\varphi : G \to \mathrm{GL}(r+1).$$

Take a one-parameter subgroup $\rho : \mathbb{G}_m \to G$, choose coordinates to diagonalise the action of ρ, and let $w = (w_0, .., w_r)$ be its weight vector. For a point

$$x = [x_0 : .. : x_r] \in X$$

we define the Hilbert-Mumford function to be

$$\mu(x, \rho) = \max\{-w_i \mid x_i \neq 0\}.$$

For our purposes it will be convenient to work with a normalised version.

2.3.0.2 Definition. Fix an invariant inner product on the set of one-parameter subgroups of G: that is, take a positive definite Weyl-invariant integer-valued bilinear form on the set of one-parameter subgroups of any maximal torus T. Denote by $\| - \|$ the associated norm. Then define the normalised Hilbert-Mumford function to be

$$M(x, \rho) = \frac{\mu(x, \rho)}{\|\rho\|}.$$

2.3.0.3 Remark. If the group G is semisimple, for example if $G = \mathrm{SL}(r+1)$, we have a natural choice of invariant inner product: the Killing form.

2.3.0.4 Definition. We say that a one-parameter subgroup $\rho : \mathbb{G}_m \to G$ is a *maximally destabilising one-parameter subgroup*, or Kempf one-parameter subgroup, of a point $x \in X$, or that ρ is *adapted* to x, if $M(x, \rho)$ is minimal amongst all one-parameter subgroup's of G.

2.3.0.5 Example. If $G = T$ is a torus diagonalised by coordinates $x_0, .., x_r$ with weights $\alpha_0, .., \alpha_r$, then the minimal value of the normalised Hilbert-Mumford function for $x = [x_0 : .. : x_r]$ is the norm of the closest point to the origin, $B(x)$, of the convex hull of the T-weights in the state polytope of x, i.e. those w_i such that $x_i \neq 0$. There is, up to scaling, a unique maximally destabilising one-parameter subgroup for x, namely that obtained by exponentiating $B(x)$.

The following result tells us what to expect more generally.

2.3.0.6 Proposition. *[Kem78] Let G be a reductive group scheme acting linearly on a projective scheme X with ample linearisation $\mathscr{L}$. Let $x \in X$ be an unstable point. Then*

1) There exists $\rho : \mathbb{G}_m \to G$ which minimises the function $M(x, -)$ amongst all $1PS$'s of G.

2) There exists a parabolic subgroup P_x of G such that for any λ satisfying (1) we have

$$P_x = P(\lambda) := \{g \in G \mid \lim_{t \to 0} \lambda(t) g \lambda(t^{-1}) \text{ exists in } G\}.$$

3) Any $1PS$ satisfying (1) is unique up to replacing λ by λ^n and conjugating by an element of the subgroup P_x.

This enables us to associate to an unstable point a one-parameter subgroup that is 'most responsible' for that point's instability, unique in the sense described above. In applications to moduli problems, these are the subgroups that we think of as somehow best highlighting whatever intrinsic property of the object in question makes it moduli-unstable.

We now define the β-stratification associated to a linear action of a reductive linear algebraic group on a projective scheme.

2.3.0.7 Definition. Fix a maximal torus T of G acting on X as above, and a positive Weyl chamber $\mathfrak{t}^+$ in its Lie algebra. Define the index set for the stratification,

$$\mathscr{B} := \{\beta \in \mathfrak{t}^+ \mid \exists x \in \mathbb{P}^r \text{ s.t. } B(x) = \beta\}.$$

To put it another way, $\mathscr{B}$ is the set of points in the positive Weyl chamber that are the closest point to the origin of the convex hull of some non-empty subset of the T-weights. To each element of this index set we associate a one-parameter subgroup.

2.3.0.8 Definition. For an element $\beta \in \mathscr{B}$, denote by λ_β the one-parameter subgroup that is given by the weight vector $q\beta$, where q is the smallest positive rational number giving integral weights.

Now that we have all the combinatorial information we will need, we can begin to show how this data captures the instability of all points of X. Given $\beta \in \mathscr{B} \setminus \{0\}$, let

$$H_\beta = \{v \in \mathfrak{t} \mid v \cdot \beta = \|\beta\|^2\},$$

$$H_\beta^+ = \{v \in \mathfrak{t} \mid v \cdot \beta \geq \|\beta\|^2\}$$

denote respectively the hyperplanes through β meeting it perpendicularly, and the half-space on the opposite side of this hyperplane from the origin.

2.3.0.9 Definition. Following [Kir84], we define

$$Z_\beta \coloneqq \{x \in X \mid \text{ all weights of } x \text{ lie on } H_\beta\}$$

$$Y_\beta \coloneqq \{x \in X \mid \text{ all weights of } x \text{ lie in } H_\beta^+ \text{ and at least one weight lies on } H_\beta\}.$$

The first of these is a closed subscheme of X, and the second is a locally closed subscheme. Clearly both contain only unstable points for all $\beta \in \mathscr{B} \setminus \{0\}$. There is a natural surjective retraction $p_\beta : Y_\beta \to Z_\beta$, given by taking the limit as $t \to 0$ under λ_β. If X is nonsingular, the fibres are just affine spaces, by [BB73].

2.3.0.10 Definition. We further define

$$Y_\beta^{ss} = \{x \in Y_\beta \mid \lambda_\beta \text{ is adapted to } x\},$$

$$Z_\beta^{ss} = \{x \in Z_\beta \mid \lambda_\beta \text{ is adapted to } x\},$$

$$S_\beta = G \cdot Y_\beta^{ss}.$$

2.3.0.11 Definition. Denote by $\operatorname{Stab}\beta$ the stabiliser of β under the adjoint action of G, and note that Z_β is invariant under the action of this group. We restrict the linearisation to the action of $\operatorname{Stab}\beta$ on Z_β, and twist it by the character $-\beta$; this has the effect of shifting the weights by this vector. We denote this linearisation by $\mathscr{L}_\beta$, and we will call it the *canonical linearisation*.

2.3.0.12 Remark. The subscheme Z_β^{ss} is exactly the semistable locus for the action of $\operatorname{Stab}\beta$ on Z_β with respect to the twisted linearisation, and $Y_\beta^{ss} = p_\beta^{-1}(Z_\beta^{ss})$ ([Kir84], 12.21).

2.3.0.13 Definition. Given $\beta \in \mathscr{B}$, denote by P_β the parabolic subgroup of G associated to λ_β. Denote by $Y_\beta^{ss} \times^{P_\beta} G$ the quotient of $Y_\beta^{ss} \times G$ by the free action of P_β given by $h \cdot (g, x) \coloneqq (gh^{-1}, hx)$.

2.3.0.14 Remark. Though a priori an algebraic space, by [PV94] we know that $Y_\beta^{ss} \times^{P_\beta} G$ is in fact a scheme.

We can now summarise the relevant results of [Kir84] that give the stratification.

2.3.0.15 Theorem. *[Kir84] Let X be a projective scheme, acted on by a reductive linear algebraic group G, linearly with respect to an ample line bundle $\mathscr{L}$. Then, with notation as above, we have:*

1) There is a stratification of X into disjoint locally closed G-invariant subschemes,

$$X = \bigsqcup_{\beta \in \mathscr{B}} S_\beta,$$

where the open stratum, S_0, is X^{ss}, and the closure of a stratum S_β satisfies

$$\overline{S_\beta} \subset \bigcup_{\substack{\gamma \text{ s.t.} \\ \|\gamma\| \geq \|\beta\|}} S_\gamma.$$

2) There is, for every $\beta \in \mathscr{B}$, a G-equivariant isomorphism $Y_\beta^{ss} \times^{P_\beta} G \cong S_\beta$.

2.3.0.16 Definition. We will refer to the weight β associated to a given point $x \in X$ via the theorem above as the *instability type of x*.

We can try to find, for every $\beta \in \mathscr{B}$, a quotient of the stratum S_β by the G-action. By (2) in the theorem above, this is equivalent to the problem of finding a quotient for the action of P_β on Y_β^{ss}.

There is in fact a categorical quotient for the action of P_β on Y_β^{ss} given by the reductive GIT quotient of Z_β by $\operatorname{Stab}\beta$, where the linearisation is twisted by β (see Chapter 4). However, this categorical quotient is very far from being a geometric quotient. If we are to obtain quotients of the correct dimension, we must prevent points of Y_β^{ss} from being identified with their images under p_β by making Z_β^{ss} unstable. We therefore split each stratum into two pieces: roughly speaking, one piece corresponds to Z_β^{ss}, and one to the open subset of Y_β^{ss} not in the P_β sweep of Z_β^{ss}. The GIT quotient of Z_β by $\operatorname{Stab}\beta$ is a categorical quotient of the former, and we can use Theorem 3.3.1.6 and its variants to obtain quotients of open subschemes of the latter.

2.3.0.17 Remark. We can also define the instability stratification in the case where the group is not reductive, in an analogous way. We will do this in Chapter 6, when we consider moduli of singular curves.

Chapter 3

Non-Reductive Geometric Invariant Theory and its Variation

What immortal hand or eye,
Could frame thy fearful symmetry?
William Blake, *The Tyger*

As well as being indispensable for the problem under consideration, developing a good Non-Reductive Geometric Invariant Theory (NRGIT) is also an important problem in its own right. Unfortunately, the good properties of classical GIT as summarised in the previous chapter do not continue to hold in the non-reductive case, and the subject is considerably more complicated. In this chapter we will first explore some of the foundations of non-reductive GIT, in Section 3.1, before going on to give, in Section 3.3, an exposition of some recent and powerful results in non-reductive GIT, due to Bérczi, Doran, Hawes and Kirwan ([BDHK], [BDHK16]). These theorems, which we call the $\hat{U}$ theorems, will do a great deal of heavy lifting for us in the remainder of this thesis. We end this chapter, in Section 3.4, by giving some new results which explore the analogue of variation of GIT in the setting of the $\hat{U}$ theorems.

3.1 First steps in Non-Reductive GIT

We begin with an introduction to NRGIT, presenting some of the basic material of [DK07]. Section 3.1.2 follows the same path as for Reductive GIT, namely that of trying to construct a quotient from the gluing of suitable affine H-invariant pieces.

However, we will see that even the affine situation is more difficult in the non-reductive case. In Section 3.1.3 we explore what happens when we try to exploit previous results by converting the problem into a quotient by a reductive action. This done, we briefly demonstrate in Section 3.2 how the theory may be applied to a particular example.

3.1.1 Motivations

As in [DK07], we start by sketching the desiderata of any theory of quotients by non-reductive group actions. This is largely motivated by [MFK94], with the initial intention that NRGIT should resemble classical GIT as far as possible.[1]

1) The theory should take as input a (projective) scheme X together with an action of a general algebraic group H and the data of a H-linearisation of the action with respect to some (ample) line bundle.

2) As with classical GIT, the quotient should be constructed by considering the space associated to some subring of invariant functions B, and the quotient morphism should correspond to the inclusion of these invariants $B \hookrightarrow A$, where A is the ring associated to X.

3) It should output a quasiprojective scheme which is a good, or at least categorical, quotient of a suitably defined open subset X^{ss}. The quotient morphism should restrict to a geometric quotient on some second open subset X^s, which should coincide with X^{ss} in good cases.

4) There should be some analogue of the Hilbert-Mumford criterion which allows us to calculate X^{ss} and X^s with minimal difficulty, i.e. without having to find a generating set for B.

Many moduli problems involve taking a quotient by a non-reductive group action: we give two examples below to briefly illustrate the applications of NRGIT outside of the situation that we will be concerned with.

3.1.1.1 Example. Suppose that we want to construct a moduli space of hypersurface subschemes of a given scheme X. Then we will have to quotient by $\mathrm{Aut}(X)$, which is in general not reductive. For a specific example, consider the weighted projective space

$$X = \mathbb{P}(1, 1, 2).$$

[1] cf. [DK07] pp.2-3.

This is the projective scheme with homogeneous coordinate ring $\mathbb{C}[x_1, y_1, z_2]$ where x_1 and y_1 are considered to have degree 1, and z_2 is considered to have degree 2. The automorphism group $\mathrm{Aut}(X)$ is the group of graded automorphisms of graded rings

$$\mathbb{C}[x_1, y_1, z_2] \xrightarrow{\sim} \mathbb{C}[x_1, y_1, z_2]$$

modulo the rescaling action of $\mathbb{G}_m$, which acts via

$$t \cdot (x_1, y_1, z_2) = (tx_1, ty_1, t^2 z_2).$$

We find that x_1 and y_1 must both map to something of degree 1, so under any $\varphi \in \mathrm{Aut}(X)$ we see that

$$(x_1, y_1) \longmapsto (ax_1 + by_2, cx_1 + dy_2)$$

where

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{GL}(2).$$

Moreover, any automorphism must preserve the degree of z_2, so we must have

$$z_2 \longmapsto \alpha z_2 + \beta x_1^2 + \gamma x_1 y_1 + \delta y_1^2$$

for α, β, γ, $\delta \in \mathbb{C}$, $\alpha \neq 0$. Hence, once the $\mathbb{G}_m$ action is accounted for we have a Levi decomposition

$$\mathrm{Aut}(X) \cong (\mathbb{C}^+)^3 \rtimes \mathrm{GL}(2)$$

and in particular the unipotent radical $U \cong (\mathbb{C}^+)^3$ is non-trivial, so the group is not reductive.

3.1.1.2 Example. Another example is that of k-germs of holomorphic curves on a smooth variety (or complex manifold) X. A k-jet is an equivalence class of parameterised holomorphic curves

$$f : \mathbb{C} \to X$$

where two such curves f, g are declared equivalent if $f(0) = g(0)$ and the derivatives of f and g agree at 0 up to kth order. For many applications we wish to eliminate the redundancy arising from the choice of parameterisation of the curve f, and only consider the k-germ of the curve itself. This involves quotienting by the (non-reductive)

re-parameterisation group $\mathbb{G}_k$ of kth order germs of biholomorphisms $(\mathbb{C},0)\to(\mathbb{C},0)$:

$$\mathbb{G}_k \cong \left\{ \begin{pmatrix} \alpha_1 & \alpha_2 & \alpha_3 & \dots & \alpha_k \\ 0 & \alpha_1^2 & \dots & & \\ 0 & 0 & \alpha_1^3 & \dots & \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & \dots & \alpha_1^k \end{pmatrix} \mid \alpha_i \in \mathbb{C}\ \forall i, \alpha_1 \neq 0 \right\}.$$

3.1.2 Attempts to glue affines

As with the reductive case, we first consider the local situation. Let X be an affine H-scheme; then we get the dual action on $k[X]$ as before. Here we hit the first, and arguably most severe, stumbling block to NRGIT.

3.1.2.1 Problem (Nagata)**. The subalgebra of invariants $k[X]^H$ need not be finitely generated.**

Originally conceived as a counterexample to Hilbert's 14th Problem[2], in our setting the above tells us that if we want the quotient by H to be of finite type, even in the affine case we will have to do something more creative than use the whole subalgebra of invariants. This naturally has consequences for our putative definition of (semi)stability, the motivation for which is that we want the quotient to be at least locally of finite type.

3.1.2.2 Definition. Let X be an H-scheme with action linearised with respect to L.

1) Let

$$I^{ssfg} = \{f \in A_L(X)^H \mid \text{the distinguished open } X_f \text{ is affine and } k[X_f]^H \text{ is finitely generated}\}$$

We define the finitely generated semistable locus to be

$$X^{ssfg} = \bigcup_{f\in I^{ssfg}} X_f.$$

2) Let

$$I^{ns} = \{f \in I^{ssfg} \mid X_f \text{ is affine and } X_f \to \mathrm{Spec}((\mathscr{O}(X_f)^H) \text{ is a geometric quotient.}\}.$$

[2]The two most famous counterexamples are those of Roberts [Rob90] and Nagata [Nag60]; the latter has historical importance as the first known counterexample, and the former is notable for its simplicity.

We further define the naively stable locus to be

$$X^{ns} = \bigcup_{f \in I^{ns}} X_f.$$

3.1.2.3 Remark. If H is reductive (or more generally if the invariants are finitely generated) we have

$$X^{ssfg} = X^{ss}.$$

Accordingly, we generalise the reductive GIT quotient as follows.

3.1.2.4 Definition. The inner enveloping quotient is defined to be

$$X /\!\!/_e H \coloneqq \bigcup_{f \in I^{ssfg}} \operatorname{Spec} k[X_f]^H$$

where the union is glued according to the quotient map

$$q : X^{ssfg} \to X /\!\!/_e H$$

which is locally induced by the inclusion of invariant rings

$$k[X_f]^H \hookrightarrow k[X_f].$$

Almost tautologously, the quotient map restricts to give a geometric quotient of X^{ns}, and it agrees with previous definitions in the finitely generated case. Indeed, as in the reductive set-up, the best behaviour for non-reductive GIT occurs when X is projective and L is very ample. The difference is that here we also require that the invariants be finitely generated.

3.1.2.5 Proposition. *[DK07] Let X be projective, and L ample. Assume further that the subalgebra of invariants*

$$A_L(X)^H \subseteq A_L(X) \coloneqq \bigoplus_{k=0}^{\infty} H^0(X, L^k)$$

is finitely generated. Then the inner enveloping quotient is the projective variety $\operatorname{Proj}(A_L(X)^H)$*, and the quotient morphism is*

$$X = \operatorname{Proj}(A_L(X)) \dashrightarrow \operatorname{Proj}(A_L(X)^H) = X /\!\!/_e H$$

induced by the global inclusion of graded rings.

3.1.2.6 Remark. When the invariants are finitely generated, even in the non-reductive case we will drop the 'fg' and just write X^s, X^{ss}, &c.

Finite generation, however, is not the only issue.

3.1.2.7 Problem. *The quotient map $q : X^{ssfg} \to X /\!\!/_e H$ may not be surjective. This means that the image, $q(X^{ssfg})$, is some constructible subset. In particular, it may not be a scheme.*

3.1.2.8 Problem. *The enveloping quotient $X /\!\!/_e H$ may not be a scheme of finite type.*

3.1.2.9 Example. We give an affine example of failure of the surjectivity of the quotient map.[3] Let $G = \mathrm{SL}(2)$, and define $H \subset G$ as

$$H = \left\{ \begin{pmatrix} 1 & \alpha \\ 0 & 1 \end{pmatrix} \mid \alpha \in \mathbb{C} \right\} \cong \mathbb{C}^+.$$

Consider the natural action of H on G by left multiplication, equipped with the trivial linearisation on $\mathscr{O}_G$. Then by Weitzenbock's theorem [Wei32] the ring of invariants is finitely generated, and indeed we will calculate it shortly. So we have

$$G /\!\!/_e H \cong \mathrm{Spec} A(G)^H$$

with GIT quotient q induced by the inclusion of rings

$$i : A(G)^H \hookrightarrow A(G).$$

Consider the natural $G = \mathrm{SL}(2)$ action on $\mathbb{C}^2$. The point $(1, 0)$ has orbit $\mathbb{C}^2 \setminus \{(0, 0)\}$ and stabiliser H, so that the induced orbit map

$$G/H \times \{(1, 0)\} \to \mathbb{C}^2$$

yields an isomorphism

$$q(G) \cong G/H \cong \mathbb{C}^2 \setminus \{(0, 0)\}.$$

Indeed, a result of Matsushima says that if G is reductive then G/H is affine if and

[3] A projective example may be found in section 6 of [DK07], in which we can see the failure of the inner enveloped quotient to be a variety.

only if H is reductive[4]. Now

$$A(G)^H \cong A(G/H) \cong A(\mathbb{C}^2 \setminus \{(0,0)\}) \cong A(\mathbb{C}^2) \cong \mathbb{C}[x,y],$$

so that

$$G /\!\!/_e H \cong \mathbb{A}^2.$$

This shows that even finite generatation of invariants does not guarantee surjectivity of the quotient map.

3.1.3 Transferral to a reductive action

Given a scheme V and a non-reductive group H, one strategy for constructing a quotient might be to try to find another scheme V' with an action of some reductive group G such that $V /\!\!/_e H \cong V' /\!\!/ G$. This may be attempted in the following way.

For any group G of which H is a closed subgroup, i.e. for which there exists a closed embedding $H \hookrightarrow G$, we can consider

$$G \times^H X := (G \times X)/\sim \quad = \{[g,x] \mid (g,x) \sim (gh^{-1}, hx)\}.$$

This carries a natural action of G, from G acting on itself. The moral of this whole story is that quotienting by group actions in the category of varieties is difficult, so we certainly need to check:

3.1.3.1 Lemma. *(See [PPS+94], theorem 4.19) $G \times^H X$ is a scheme.*

We can regard X as a subscheme of $G \times^H X$ via the embedding

$$j : X \to G \times^H X$$

$$x \longmapsto [e,x].$$

The hope of being able to use the existing GIT for reductive groups leads one to consider something called a reductive envelope. For our purposes the precise definition, which may be found at [DK07] 5.2.4, need not concern us;[5] morally, it involves taking a projective completion

$$G \times^H X \subset \overline{G \times^H X}$$

[4] See [DK07] 5.1.1.

[5] In Section 3.2 we will be in the much simpler situation where the action extends to the action of a reductive group, and in our more substantial applications later on we will proceed using the $\hat{U}$ Theorems.

together with a linearisation having the property that enough H-invariants for $G\times^H X$ extend to G-invariants for $\overline{G\times^H X}$. Now by the 'Main Theorem' 5.3.1 of [DK07], the enveloping quotient $X /\!\!/_e H$ that we are interested in lives as an open subscheme of $\overline{G\times^H X} /\!\!/ G$ via the natural embedding above, and since the latter is a quotient by a reductive group action, it will be much easier to deal with. In particular, having taken a projective completion we can use the Hilbert-Mumford criterion.

In our example, as usual, both the resulting quotient and the semistable locus itself are dependent upon a choice of linearisation and ample line bundle. We will investigate different linearisations and see how this affects the quotient.

It sometimes happens that the action of H extends to an action of G. We then have the isomorphism of varieties

$$G\times^H X \to (G/H)\times X$$

$$[g,x]\to (gH,gx).$$

We can now apply all this to our case, but first we summarise the overall strategy.

1) Embed H in a reductive group G, preferably (as in the example we discuss below) such that the linear action of H on X extends to a linear action of G on X, and H-invariants on X extend to $G\times^H X$ in the appropriate way.

2) Since the action extends to G, we can look at the reductive GIT quotient $((G/H)\times X) /\!\!/ G$.

3) Find a G-equivariant projective completion of $G\times^H X$ with a suitable linearisation, preferably of the form $\overline{G/H}\times X$ after indentifying $G/H\times X$ with $G\times^H X$ as above.

4) We can then use the Hilbert-Mumford criterion on the projective GIT quotient $((\overline{G/H})\times X) /\!\!/ G$ to find the semistable locus.

5) By varying the linearisation of the G-action we can study how the quotient varies, using Hilbert-Mumford and the notion of S-equivalence.

In principle we need to ensure that every linearisation we look at is compatible with a suitable linearisation of the H-action on X, but in our case this will be clear.

3.2 Example: ordered points on $\mathbb{P}^1$ up to translation

Here we will go through in detail an elementary example of NRGIT, and illustrate how we can use the principle of transferring to a reductive action discussed above.

Suppose we are interested in parameterising choices of n ordered points on the projective line up to translation. A point of $(\mathbb{P}^1)^n$ consists of a choice of n ordered points on the projective line, which carries a natural action of the (non-reductive) group $\mathbb{C}^+$ via the faithful representation

$$\rho : \mathbb{C}^+ \to PGL(2)$$
$$\rho(\alpha) = \begin{pmatrix} 1 & \alpha \\ 0 & 1 \end{pmatrix}$$

If we view $\mathbb{P}^1$ as $\mathbb{C} \cup \{\infty\}$, then this action fixes the point at infinity and adds α to affine points. We are interested in the GIT quotient(s) of $(\mathbb{P}^1)^n$ by this action. Let

$$X = (\mathbb{P}^1)^n$$

$$G = \mathrm{SL}(2)$$

$$H = \mathbb{C}^+$$

with the embedding $H \hookrightarrow G$ given above. Then we have

$$G \times^H X \cong (G/H) \times X.$$

Moreover, the variety G/H can readily be identified. Consider, as we did above, the action of G on $\mathbb{C}^2$ in the natural way. Then the orbit of the point $(1, 0)$ is $\mathbb{C}^2 \setminus \{0\}$, and it has stabiliser H. Hence

$$(G/H) \cong \mathbb{C}^2 \setminus \{0\}.$$

This has a natural projective completion $\overline{\mathbb{C}^2 \setminus \{0\}} = \mathbb{P}^2$. So we are interested in

$$((\mathbb{P}^1)^n \times \mathbb{P}^2) /\!\!/ G.$$

3.2.0.1 Notation. It will relevant later to note that our original embedding

$$j : X \to G \times_H X$$

can be composed with all this to give another embedding (which by abuse of notation we also call j):

$$j : X \to (\mathbb{P}^1)^n \times \mathbb{P}^2$$
$$(P_1, ..., P_n) \longmapsto (P_1, ..., P_n, [1 : 0 : 1]).$$

3.2.1 The Linearisation

Until now we have largely neglected to mention linearisation of the G and H actions, and we address this here. Each factor $\mathbb{P}^1$ in the product $X = (\mathbb{P}^1)^n$ is already a projective space, so we have a natural choice of linearisation for each ample line bundle $\mathscr{O}_{\mathbb{P}^1}(k)$, for $k > 0$, since a linearisation on some L induces a linearisation on every tensor power of L. We embed X into projective space using a composition of Segre embeddings, whereby each choice of tuple $(k_1, ..., k_n)$ determines a linearisation on X. We consider linearisations of this form. We also need a linearisation of the G-action on $\mathbb{P}^2$, and here we use the embedding

$$G \hookrightarrow GL(3)$$
$$A \longmapsto \left(\begin{array}{c|c} A & 0 \\ \hline 0 & 1 \end{array}\right)$$

which similarly determines a linearisation for any positive integer m. So to sum up: we have a linearisation associated to the vector $(k_1, ..., k_n, m) \in \mathbb{N}^{n+1}$. These may be be considered as the Segre embedding applied after a linearisation with respect to $\mathscr{O}_{\mathbb{P}^1}(k_i)$ on the ith factor, or equivalently as performing a degree k_i Veronese embedding on the ith factor followed by the Segre embedding.

3.2.2 The 1-parameter subgroups

Recall that all the information about (semi)stability is contained in the data of the weights of any maximal torus. In this set-up we have a natural choice of maximal $T \subset G$, namely the diagonal matrices

$$T \coloneqq \left\{ \begin{pmatrix} t & 0 \\ 0 & t^{-1} \end{pmatrix} \mid t \in \mathbb{C} \right\} \cong \mathbb{G}_m.$$

In fact, all we need to know to determine T-stability of a point is where the origin lies in relation to the convex hull of its weights. Since T is one-dimensional, so is $\mathfrak{t}$, so the weights will just be points along a line. This tells us that we need only keep track of the extremal weights of each point, since these determine the convex hull.

3.2.3 The Embedding and the weights

We label the coordinates on $\mathbb{P}^1_i$ as $[x_i : y_i]$, and on $\mathbb{P}^2$ we take coordinates $[u : v : w]$. Then we see from the above that x_i has weight k_i, y_i has weight $-k_i$, and u, v, w have weights $-m$, m, 0, respectively.

In order to use Hilbert-Mumford, we must embed the whole product $(\mathbb{P}^1)^n \times \mathbb{P}^2$ into some projective space. Consideration of a Segre embedding of two projective spaces

$$\mathbb{P}^r \times \mathbb{P}^s \hookrightarrow \mathbb{P}^{(r+1)(s+1)-1} =: Y$$

on which the maximal torus acts with weights

$$\{a_i \mid 0 \leq i \leq r\}$$

$$\{b_j \mid 0 \leq j \leq s\}$$

shows that the weights of the action on the embedded product will be exactly the set

$$\{a_i + b_j \mid 0 \leq i \leq r,\ 0 \leq j \leq s\}.$$

We know that the ith factor in the $(\mathbb{P}^1)^n$ has maximal torus weights $\{k_i, -k_i\}$, and the torus acts on our $\mathbb{P}^2$ with weights $\{m, -m, 0\}$. So the set of weights of the action of the maximal torus will be

$$\Lambda = \Lambda^+ \cup \Lambda^- \cup \Lambda^0$$

$$\Lambda^0 = \{\sum_{i=1}^{n} \sigma_i k_i \mid \sigma_i \in \{1, -1\}\}$$

$$\Lambda^+ = \{m + (\sum_{i=1}^{n} \sigma_i k_i) \mid \sigma_i \in \{1, -1\}\}$$

$$\Lambda^- = \{-m + (\sum_{i=1}^{n} \sigma_i k_i) \mid \sigma_i \in \{1, -1\}\}.$$

A point $P = (P_1, ..., P_n, Q)$, where $P_i = [\alpha_i : \beta_i]$ and $Q = [u : v : w]$, will pick up the weights corresponding to those coordinates which are not zero.

3.2.3.1 Remark. We choose m large, so that the weights in Λ^+ are always all positive and the weights in Λ^- are always all negative. The signs of the individual weights of Λ^0 will depend on the k_i, but it will always contain both positive and negative weights.

3.2.3.2 Notation. For a point $(P_1, ..., P_n, Q) \in (\mathbb{P}^1)^n \times \mathbb{P}^2$ we define the *weighted sum of the points* to be

$$N = \sum_{i=1}^{n} k_i.$$

In addition, we say that a subset $I \subset \{1, ..., n\}$ contains *more than half of the points in the weighted sum* if $\sum_{i \in I} k_i \geq \frac{N}{2}$.

3.2.4 Calculation of the semistable locus

Recalling our previous results, we know that a point is semistable with respect to the action of G if and only if for every point of its G-orbit, the origin lies in the convex hull of the weights of our chosen maximal torus T picked up by that point of the orbit. A point is stable if the origin lies in the interior of that convex hull. So when, in our example, is a point (semi)stable? We first need to understand stability for the action of the maximal torus.

3.2.4.1 Lemma. *Choose a set of weights $(k_1, ..., k_n, m) \in \mathbb{Z}^{n+1}_{\geq 1}$, and let m be large. Let $\underline{P} = (P_1, ..., P_n, [a : b : 1]) \in (\mathbb{P}^1)^n \times \mathbb{P}^2$. Then*

1) *If $a \neq 0$ and $b \neq 0$ then $\underline{P}$ is stable for the action of the maximal torus T.*

2) *If $a = 0$, $b \neq 0$ then $\underline{P}$ is semistable for T if and only if fewer than half, or exactly half, of the points in the weighted sum coincide at infinity. It is stable for T if and only if this inequality is strict.*

3) *If $b = 0$, $a \neq 0$ then $\underline{P}$ is semistable for T if and only if fewer than half, or exactly half, of the points in the weighted sum coincide at zero. It is stable for T if and only if this inequality is strict.*

4) If $a = b = 0$ then $\underline{P}$ is semistable for T if and only if fewer than half, or exactly half, of the points in the weighted sum coincide at zero or that more than half of the points in the weighted sum coincide at infinity. It is stable for T if and only if both these inequalities are strict.

Proof. For (2), (3), and (4) we prove the statement about semistability; the stability statement follows in exactly the same way by changing 'positive' to 'strictly positive' and 'negative' to 'strictly negative'.

1) In this case we pick up weights from both Λ^+ and Λ^-, so since m is large we have both a positive and a negative weight and the point is stable for T.

2) If $b \neq 0$ then we certainly have a negative weight. Since $a = 0$ we have no Λ^+ weights so whether we have a positive weight will depend on the P_i. We have a positive weight from Λ^0 if and only if it is not the case that more than half the points in the weighted sum carry only negative weights, i.e. are equal to $[0 : 1]$. That is, if they coincide at zero.

3) This is analogous to (2): we have a positive weight, so we need a negative weight from Λ^0. This happens precisely when the given condition is satisfied.

4) In this case we pick up no Λ^+ or Λ^- weights, so we need both a positive and a negative weight from Λ^0. The rest is the same.

□

Having found necessary and sufficient conditions for (semi)stability with respect to the action of our chosen maximal torus T, we are ready to find conditions for (semi)stability for the full G action. We next make two observations.

3.2.4.2 Remark.

1) First recall that, in the case of the usual action of G on $(\mathbb{P}^1)^n$, a point is semistable if fewer than, or exactly, $\frac{n}{2}$ of the points coincide anywhere, and stable if this inequality is strict. This is shown in the same way as we deal with points of the form $(P_1, ..., P_n, [0 : 0 : 1])$ below. So we expect that semistability will be a requirement that 'not too many points coincide'.

2) The action of G on $\mathbb{P}^1$ is transitive, and intuitively it treats all points the same. However, for the translation action of $\mathbb{C}^+$ the point at infinity, $[1 : 0]$, is the only invariant point, and hence we would expect semistability to distinguish ∞ in some way.

Putting these two observations together, we can make an educated guess as to the stable and semistable loci for the points of the form $(P_1, ..., P_n, [1 : 0 : 1])$. This guess turns out to be correct, but of course for our purposes we will need a more complete understanding of the semistable locus than just those points coming from $j : X \to (\mathbb{P}^1)^n \times \mathbb{P}^2$.

3.2.4.3 Proposition. *Choose a set of weights $(k_1, ..., k_n) \in \mathbb{Z}^n_{\geq 1}$, and let $m >> 0$. Fix $\underline{P} = (P_1, ..., P_n, [a : b : 1]) \in (\mathbb{P}^1)^n \times \mathbb{P}^2$. Then for the given action of G on $(\mathbb{P}^1)^n \times \mathbb{P}^2$, weighted by $(k_1, ..., k_n, m)$,*

1) If $a = b = 0$, $\underline{P}$ is semistable if and only if fewer than half, or exactly half, of the points in the weighted sum coincide anywhere. It is stable if this inequality is strict.

2) If $(a, b) \neq (0, 0)$ then $\underline{P}$ is semistable if and only if fewer than half, or exactly half, of the points in the weighted sum coincide at $[a : b]$. It is stable if this inequality is strict.

Proof. 1) It is clear that points of the orbit of $\underline{P}$ will be all of the form

$$(P'_1, ..., P'_n, [0 : 0 : 1]),$$

and if not more than half the points in the weighted sum for $\underline{P}$ coincide anywhere then it will certainly hold that no point in the orbit has more than half of the points in the weighted sum at ∞ or at 0. Conversely, G acts transitively on the projective line and so any coinciding tuple of points with weights summing to greater than $\frac{N}{2}$ can be moved by G to ∞, for example.

2) If more than half of the points in the weighted sum for P coincide at $[a : b]$ then we can apply any element of G satisfying $[a : b] \mapsto \infty$ to get a point $\underline{P'} = (P'_1, ..., P'_n, [1 : 0 : 1])$ which is unstable for the maximal torus since more than half the point in its weighted sum now coincide at ∞. Conversely, we must check all the point of the orbit of $\underline{P}$ satisfying the hypotheses. Points $\underline{P'} = (P'_1, ..., P'_n, [s : t : 1])$ of the orbit of $\underline{P}$ are certainly stable if both s and t are non-zero, so it suffices to check separately the points where $s = 0$ and the points where $t = 0$, noting that, since $\underline{P'}$ is in the orbit of $\underline{P}$, we know that at most one of s, t can be zero at once. First let $s = 0$. Then $\underline{P'} = A \cdot \underline{P}$ where A sends $[1 : 0 : 1]$ to $[0 : t : 1]$, so under A we have $[a : b] \mapsto [0 : 1]$. Thus since it is not the case that more than half the points in the weighted sum for

$\underline{P}$ coincide at $[a : b]$, it is not the case that more than half the points in the weighted sum for $\underline{P'}$ coincide at zero, since A's action must be injective. Hence $\underline{P'}$ has a positive weight. And since $t \neq 0$, it has a negative weight. Now let $t = 0$. This is precisely analogous to the previous case, but with 0 replaced by ∞ and swapping the roles of positive and negative weights.

□

3.2.4.4 Corollary. *A point $(P_1, ..., P_n)$ of X is semistable (stable) if and only if not more than (resp. strictly less than) $\frac{N}{2}$ of the points P_i, considered with weights, lie at infinity.*

3.3 The $\hat{U}$ Theorems

Recently, it has been shown ([BDHK], [BDHK16]) that the good properties of GIT enjoyed in the reductive setting may be recovered for non-reductive linear algebraic group actions under certain circumstances, which are satisfied in a wide range of interesting applications. Here we present these results, collectively referred to as the $\hat{U}$ theorems. We first introduce some notation.

3.3.1 Statement of the theorems

3.3.1.1 Definition. Let H be an arbitrary linear algebraic group. Let U be its unipotent radical. We say that H has *graded unipotent radical* if there exists a one-parameter group of automorphisms $\mathbb{G}_m \leq \operatorname{Aut} H$ which acts with strictly positive weights on $\operatorname{Lie} U$, such that the induced action on $R = H/U$ is trivial. Such a $\mathbb{G}_m$ is called *admissible*, and we refer to any given fixed choice as the *distinguished* $\mathbb{G}_m$. We call an admissible $\mathbb{G}_m$ *internal* if it comes from a subgroup $\mathbb{G}_m \leq H$, with the automorphism being conjugation, and *external* otherwise. For the rest of this section we will assume that we have an internal admissible $\mathbb{G}_m$ (but see Remark 3.3.2.3 below). We will write $\hat{U} = U \rtimes \mathbb{G}_m$, or, for example, $\hat{U}_\lambda$ if we wish to emphasise a particular choice of homomorphism $\lambda : \mathbb{G}_m \to \operatorname{Aut}(H)$.

Up to conjugacy, we may assume that an internal admissible $\mathbb{G}_m$ is chosen from the Lie algebra $\mathfrak{z}$ of the centre of R, inside a maximal torus whose Lie algebra we denote by $\mathfrak{t}$. Then, choosing an invariant inner product, the admissible $\mathbb{G}_m$'s are precisely those whose pairings with all weights of the adjoint representation of $\mathfrak{t}$ on $\operatorname{Lie} U$ are strictly positive. This defines a cone $\mathfrak{C} \subseteq \mathfrak{z}^\vee$, which we call the *admissible cone*.

3.3.1.2 Remark. We emphasise that so far all of this is intrinsic to the group; it depends neither on X, nor on the action.

3.3.1.3 Definition. Now suppose H acts on X, a projective scheme. We say that a linearisation $\mathscr{L}$ of an action of H on X with respect to an ample line bundle L is *adapted* if the lowest weight of the action of the distinguished $\mathbb{G}_m$ on X is strictly negative, and the rest are strictly positive. Let $V_{\min}^{\mathbb{G}_m}$ be the minimal weight space of the linear representation of the distinguished $\mathbb{G}_m$. Let $Z_{\min}^{\mathbb{G}_m} = X \cap \mathbb{P}(V_{\min}^{\mathbb{G}_m})$, and let $X_{\min}^{0,\mathbb{G}_m}$ be the subscheme consisting of those $x \in X$ with at least one coordinate associated to the minimal weight space being non-zero: that is, the basin of attraction of $Z_{\min}^{\mathbb{G}_m}$ under the distinguished $\mathbb{G}_m$. Let

$$p : X_{\min}^{0} \to Z_{\min}$$

be the projection onto the minimal weight space, that corresponds to taking the limit under the distinguished $\mathbb{G}_m$ as $t \to 0$. We will observe later that $Z_{\min}^{\mathbb{G}_m}$ and $X_{\min}^{0,\mathbb{G}_m}$ are in fact independent of the choice of ample linearisation. If the choice of distinguished $\mathbb{G}_m$ is clear from context, we will simply write $Z_{\min}$ and $X_{\min}^{0}$.

3.3.1.4 Remark. If a given linearisation is not adapted, we can always twist the linearisation by a rational character to make it so, provided appropriate characters exist.

3.3.1.5 Remark. There is a technical point to be explained here, about the choice of linearisation. For the proofs of finite generation of invariants given in [BDHK, BDHK16] to work, we must twist the linearisation by a rational character so that it is not merely adapted, but within some sufficiently small $\varepsilon > 0$ of the lower wall of the adapted chamber. This can make the statement of the results that follow somewhat unwieldy, so we adopt the following convention used in [BDHK, BDHK16]: we say that a linearisation is *well adapted* if it is within distance $\varepsilon > 0$ of the relevant wall. More precisely, when in what follows we say that a property holds for a well adapted linearisation, we mean that for any adapted linearisation there exists an $\varepsilon > 0$ such that the relevant property holds after twisting by a rational character so that the adapted linearisation lies within an open ε-neighbourhood of the lower wall of the adapted chamber. This will only be relevant for finite generation of invariants; the stable and semistable loci and the associated quotients require simply that the linearisation should be adapted.

We are now ready to state the relevant versions of the $\hat{U}$ theorems.

3.3.1.6 Theorem. *[BDHK16]* ***($\hat{U}$-theorem when semistability coincides with stability for the unipotent radical)*** *Let H act on a projective scheme X, let $\hat{U}$ be formed as above by an admissible internal $\lambda : \mathbb{G}_m \to H$, and let $\mathscr{L}$ be a well adapted linearisation with respect to a very ample line bundle L. Assume that*

$$\mathrm{Stab}_U(x) = \{e\} \text{ for all } x \in Z_{\min}. \qquad (\star)$$

Then

1) *The $\hat{U}$-invariants are finitely generated, and the inclusion of the $\hat{U}$-invariant algebra induces a projective geometric quotient of an open subscheme $X^{\hat{U}-s} = X^{\hat{U}-ss}$ of X (the $\hat{U}$-(semi)stable locus):*

$$X^{\hat{U}-s} \to X /\!\!/_{\mathscr{L}} \hat{U}.$$

2) *Consequently the H-invariants are finitely generated, and the inclusion of the H-invariant subalgebra induces a good quotient of an open subscheme of X (the H-semistable locus):*

$$X^{H-ss} \to X /\!\!/_{\mathscr{L}} H,$$

where $X /\!\!/_{\mathscr{L}} H$ is the GIT quotient of $X /\!\!/_{\mathscr{L}} \hat{U}$ by the induced action of the reductive group $H/\hat{U} \cong R/\mathbb{G}_m$ with respect to the induced linearisation. The good quotient $X^{H-ss} \to X /\!\!/_{\mathscr{L}} H$ restricts to a geometric quotient $X^{s,H} \to X^{s,H}/H \subseteq X /\!\!/_{\mathscr{L}} H$ of the H-stable locus X^{H-s}.

3) ***(Non-Reductive Hilbert-Mumford Criterion)*** *$x \in X$ is (semi)stable for H if and only if it is (semi)stable for every 1PS of H. That is, if T denotes any maximal torus of H, we have*

$$X^{H-(s)s}(\mathscr{L}) = \bigcap_{h \in H} hX^{T-(s)s}(\mathscr{L}).$$

3.3.1.7 Remark. The condition $(\star)$ holds if and only if we have $\mathrm{Stab}_U(x) = \{e\}$ for all $x \in X^0_{\min}$.

If the relevant stabilisers are not trivial for all points in the minimal weight space, we can proceed by a process analogous to, and indeed a generalisation of, the partial desingularisation process of [Kir85] for the reductive setting, when a linearisation has strictly semistable points.

3.3.1.8 Theorem. *[BDHK16]* ***($\hat{U}$-theorem giving projective completions)*** *With notation as in Theorem 3.3.1.6, suppose now that ($\star$) may fail, but we have instead the weaker condition that*

$$\mathrm{Stab}_U(x) = \{e\} \text{ for generic } x \in X_{\min}. \qquad (\star\star)$$

Then

1) *There exists a sequence of blow-ups of X along H-invariant projective subvarieties resulting in a projective scheme $\widehat{X}$ (with blow-down map $\widehat{\pi} : \widehat{X} \to X$) for which the conditions of Theorem 3.3.1.6 hold for a suitable linearisation, $\widehat{\mathscr{L}}$, which can be taken to be an arbitrarily small perturbation of $\widehat{\pi}^*(\mathscr{L})$.*

2) *There exists a further sequence of blow-ups along H-invariant projective subvarieties, resulting in a projective scheme $\tilde{X}$ (with blow-down map $\widetilde{\pi} : \tilde{X} \to X$) such that the conditions of Theorem 3.3.1.6 still hold for a suitable linearisation, and such that the quotient given by that theorem is a geometric quotient of an open subscheme $\widetilde{X}^{H-s}$ of $\widetilde{X}$.*

3.3.1.9 Remark.

1) Although the blow-up process for $\hat{U}$ described above is analogous to the partial desingularization process of [Kir85] that one may carry out in the reductive setting, the analogy must not be pushed too far. In particular, it is important to note that, unlike in reductive GIT, in the non-reductive setting this process is necessary if one is to have any quotient at all. Consequently, the projective completion obtained thereby may not be thought of as a partial desingularization of some other more singular quotient.

2) This gives us an H-invariant open subscheme X^{H-s} of X with a geometric quotient by H which is an open subscheme of the projective variety $\widehat{X} /\!\!/_{\widehat{\mathscr{L}}} H$ (and also of $\tilde{X} /\!\!/_{\tilde{\mathscr{L}}} H$). Here X^{H-s} is the image under $\widehat{\pi}$ of the intersection of $\widehat{X}^{H-s}$ with the complement of the exceptional divisor in $\widehat{X}$.

3) The centres of the blow-ups used to obtain $\widehat{X}$ from X are determined by the dimensions of the stabilisers in U of the limits $\lim_{t\to 0} x$ for $x \in X^0_{\min}$. Another characterisation of $X^{\hat{U}-s}$ is as

$$X^{\hat{U}-s} = \{x \in \pi(\widetilde{X}^{\hat{U}-s}) \mid \dim \mathrm{Stab}_U(\lim_{t\to 0} x) = 0\}.$$

There are applications for which all points of X have non-trivial stabilisers in U, but we would still like to be able to perform quotients. There is a procedure that allows this, yielding a further version of the $\hat{U}$ theorem.

3.3.1.10 Theorem. *[BDHK16]* ***($\hat{U}$-theorem with positive-dimensional stabilisers in U)*** *Suppose that for the derived series*

$$U = U^1 > U^2 ... > U^s = \{e\}$$

of U, we have

$$\forall j \in \{1, \dots, s\} \; \exists d_j \in \mathbb{N} \text{ such that } \dim \mathrm{Stab}_{U^j}(x) = d_j \text{ for all } x \in X^0_{\min}. \quad (\star\star\star)$$

Then the conclusions of Theorem 3.3.1.6 hold.

3.3.1.11 Remark. In fact the theorem may be applied with any series

$$U = U^1 > U^2 ... > U^s = \{e\}$$

which is normalised by H and has successive quotients U^j/U^{j+1} abelian, provided that $(\star\star\star)$ holds.

Finally, we may not have constant-dimensional stabilisers across $X^0_{\min}$ for any such series. In this case there is a final version of the $\hat{U}$ theorem.

3.3.1.12 Theorem. *[BDHK16]* ***($\hat{U}$-theorem with positive-dimensional stabilisers in U, giving projective completions)***

There is a sequence of iterated blow-ups of X along H-invariant closed subvarieties resulting in $\pi : \widehat{X} \to X$ where $\widehat{X}$ is a projective scheme with an induced action of H, linearised by some $\widehat{\mathscr{L}}$, for which $(\star\star\star)$ holds for each $U^{(j)}$ in turn, and hence for which the conclusions of Theorem 3.3.1.10 hold.

3.3.2 The $\hat{U}$ blow-up process

We will make use of Theorem 3.3.1.12 in Chapter 5 to construct moduli spaces of unstable sheaves, so it is worthwhile giving here a few details of how the blow-up process is carried out in practice. Proofs of these statements, as far as they do not simply follow from the theorems quoted above, may be found in [BDHK16]. The

notation is a little cumbersome, but the idea is simple enough. For $d \in \mathbb{N}$, let

$$C_{j,d}(X) = \{x \in X^0_{\min} \mid \dim \mathrm{Stab}_{U^{s-j}}(x) = d\},$$

and define

$$d^j_{\max}(X) = \max\{\dim \mathrm{Stab}_{U^j}(x) \mid x \in X^0_{\min}\}.$$

With this notation, we first perform a series of blow-ups to achieve condition $(\star\star\star)$ for U^{s-1}- that is, to have a constant-dimensional stabiliser in U^{s-1} across all of $X^0_{\min}$. We begin by blowing up along the closure of the locus $C_{1,d^1_{\max}(X)}$. This gives us a new space $\pi^{(1)} : X^{(1)} \to X$, with exceptional locus $E^{(1)}$. Since the centre of the blow up was H-invariant, the new space has an H-action lifting the action on X, which we linearise by defining

$$\mathscr{L}^{(1)} \coloneqq (\pi^{(1)})^* \mathscr{L}^r \otimes \mathcal{O}(-E^{(1)}).$$

We choose $r \gg 0$, so that, since replacing a linearisation with a positive multiple makes no difference to (semi)stability, we may think of this linearisation as a perturbation of the original linearisation of the action on X by a sufficiently small multiple of the exceptional divisor.

3.3.2.1 Remark. Let us pause briefly to illustrate the properties of the linearisation on the blow up. For simplicity, assume that we have a $\mathbb{G}_m$-action on $\mathbb{P}^n$, with weights $w_i = \mathrm{wt}_{\mathbb{G}_m}(x_i)$, and that we are blowing up along a linear subspace $\mathbb{P}^k \subset \mathbb{P}^n$. Then the blow up is

$$\mathrm{Bl}_{\mathbb{P}^k} \mathbb{P}^n = \{([x_0 : \cdots : x_n], [\tilde{x}_{k+1} : \ldots \tilde{x}_n] \in \mathbb{P}^n \times \mathbb{P}^{n-k-1} \mid x_i \tilde{x}_j = x_j \tilde{x}_i\}.$$

With respect to the choice of linearisation given above, weights of the $\mathbb{G}_m$ action on the blow up will be given by summing r of the weights $w_0, \ldots, w_n$ and then perturbing by one of the weights $w_{k+1}, \ldots w_n$: that is, we replace each weight w_i by a small cluster of weights, and each weight in the cluster corresponds to a weight in the projectivised normal bundle of the centre of the blow up.

In general, we will be blowing up a possibly singular, non-reduced (&c) X along a possibly singular locus. However, our choice of ample linearisation gives an embedding $X \subset \mathbb{P}^n$, and we may perform the blow ups along the relevant loci considered in this space, which, by the arguments of [BDHK16], will be smooth because $\mathbb{P}^n$ is. This means that the identification of the exceptional divisor in $\mathrm{Bl}_{\mathbb{P}^k} \mathbb{P}^n$ with the projectivised normal bundle to the centre of the blow up holds, and we may consider the exceptional divisor in the blow up of X to be the intersection of the proper

transform of X with this, which will therefore have weights as described above. In addition, since the resulting semistable loci are given by a Hilbert-Mumford criterion, the result is unchanged under replacing our original linearisation with a positive multiple.

Let us denote the new minimal weight space for the distinguished $\lambda : \mathbb{G}_m \to H$ by $Z_{\min}(X^{(1)})$. Then let

$$(X^{(1)})^0_{\min} = \{x \in X^{(1)} \mid \lim_{t\to 0} \lambda(t)x \in Z_{\min}(X^{(1)})\}.$$

The arguments of [BDHK16] tell us that we now have

$$d^1_{\max}(X^{(1)}) \coloneqq \max\{\dim \mathrm{Stab}_{U^{s-1}}(x) \mid x \in X^{(1)}\} < d^1_{\max}.$$

So we repeat the process: we blow up $X^{(1)}$ along the closure of $C_{1,d^1_{\max}}(X^{(1)})$, obtaining $X^{(2)}$ with another action lifting that of X, which we again perturb by a small multiple of the exceptional divisor $E^{(2)}$.

Inductively, we see that this process will terminate with a space $\pi^{(l)} : X^{(l)} \to X$ such that the action of $U^{s-1} \leq U$ satisfies condition $(\star\star\star)$. We then repeat the whole process for U^{s-2}, blowing up along closures of loci of the above form, and replacing U^j with U^{j-1} until we have the $\pi : \widehat{X} \to X$ mentioned in Theorem 3.3.1.12. The results of [BDHK16], stated above, then tell us that there is an open subset $\widehat{X}^{H-ss} \subset \widehat{X}$, characterised by the Hilbert Mumford criterion 3.3.1.6, which admits a projective categorical quotient

$$\widehat{X}^{H-ss} \to \widehat{X} /\!\!/_{\widehat{\mathscr{L}}} H.$$

Of course, in many cases what one is really interested in is obtaining a quotient of an open subset of X. In the situation that will be of greatest concern to us, when we have posed a particular moduli problems using GIT, the H-orbits in X will have some moduli-theoretic interpretation as equivalence classes of objects to be classified, whereas the points of $\widehat{X}$ may not readily admit such an interpretation. As a result, we will mainly use Theorem 3.3.1.12 in the following form.

3.3.2.2 Corollary. *With notation as in Theorem 3.3.1.12, the blow down map $\pi : \widehat{X} \to X$ restricts to the complement of the exceptional divisor E, to give an H-equivariant identification $X \setminus \pi(E) \cong \widehat{X} \setminus E$, and consequently, we may define in X an open subset,*

$$X^{\widehat{ss}-H} \coloneqq \pi(\widehat{X}^{H-ss} \setminus E) \cong \widehat{X}^{H-ss} \setminus E.$$

This open set admits a quasi-projective categorical quotient, denoted

$$X^{\widehat{ss}-H} \to X \wr\!\wr_{\mathscr{L}} H,$$

of which $\widehat{X} /\!\!/ H$ is a projective completion. Moreover, because the linearisation for $\widehat{X}$ may be thought of as a small perturbation as the linearisation for X, this locus is determined within $X \setminus \pi(E)$ by the Hilbert-Mumford criterion.

Hence, the price that we must pay for the construction of moduli spaces when we do not have $(\star)$ for the unipotent radical is that the resultant spaces will not be compact, that is, they will be quasi-projective rather than projective, and the projective completions yielded by the above process will not necessarily readily admit modular interpretation.

As observed in [BDHK16], in order to determine $X^{\widehat{ss}-H}$, we must keep track of the minimal weight space and its basin of attraction at each stage of the process. To simplify notation slightly, let's suppose that we are at the very beginning of the process. We take the largest j such that the dimension of the stabilisers in $X^0_{\min}$ are not constant for $\hat{U}^{s-j}$, and then blow up along the closure of some locus $C_{j,d^j_{\max}}(X)$, to get $\pi^{(1)} : X^{(1)} \to X$ with exceptional locus $E^{(1)}$. Now there are two cases.

1) If $Z_{\min} \not\subseteq C_{j,d^j_{\max}(X)}$, then the minimal weight space in the blow up, $Z_{\min}(X^{(1)})$ is simply the proper transform of the old one, and $\pi^{(1)}$ gives an identification

$$(X^{(1)})^0_{\min} \setminus E \cong \pi^{(1)}((X^{(1)})^0_{\min} \setminus E) = \{x \in X^0_{\min} \mid \dim \mathrm{Stab}^{(j)}_U((p(x)) < d^j_{\max}\},$$

where we recall that $p : X^0_{\min} \to Z_{\min}$ is the map given by

$$x \longmapsto \lim_{t\to 0} \lambda(t)x$$

as defined at Definition 3.3.1.3.

2) If $Z_{\min} \subseteq C_{j,d^j_{\max}}(X)$, the situation is more complicated. Roughly speaking, since we are blowing up the whole minimal weight space, there will be nothing left of it, and consequently its place will be taken by that part of the exceptional divisor which lies over the old minimal weight space and corresponds to the lowest weight space that we haven't blown up entirely. More precisely, for $r \in \mathbb{Z}$, define Z_r to be the λ weight space of weight r, and then let

$$W^0_r = \{x \in X^0_{\min} \mid \lim_{t\to 0} \lambda^{-1}(t)x \in Z_r\}.$$

That is to say, W_r^0 is those points of $X_{\min}^0$ whose highest weight is r. Now define

$$r_{\min} = \min\{r \in \mathbb{Z} \mid W_r^0 \not\subseteq C_{j,d_{\max}^j}(X)\}.$$

Then the new minimal weight space is that subset of $(\pi^{(1)})^{-1}(Z_{\min})$ given by

$$p^{(1)} \circ (\pi^{(1)})^{-1}(W_r^0 \setminus C_{j,d_{\max}^j}(X)),$$

where $p^{(1)}$ is the map induced by p on the blow up.

It is clear from the above that, although Corollary 3.3.2.2 gives us an open subset of X that has a categorical quotient, determining this locus in practice can be non-trivial. Even more so: if we are to find an interpretation of this categorical quotient as a moduli space for certain objects, we must understand the whole blow-up process in terms of the objects parameterised by X, and such conditions as being in a certain $C_{j,d_{\max}^j}$ or W_r^0 should have some meaning intrinsic to the objects themselves. Overall, the task is analogous to that of interpreting a particular instability stratification in modular terms, and indeed the process outlined in [BHK17] encourages us to see these as two halves of the same problem. That is, given a moduli problem posed using reductive GIT, we first consider the instability stratification associated to this problem, each stratum of which then yields a non-reductive GIT problem, which we may have to perform blow-ups in order to solve. In [BHK17], the authors show that by applying Theorem 3.3.1.12 to the closures of the subvarieties where the dimensions of stabilisers take different values, and combining this with the partial desingularisation construction of [Kir85] for reductive GIT quotients, X can be stratified so that each stratum is a locally closed H-invariant subscheme of X with a geometric quotient by the action of H (cf. [BDHK16] §5). We will not explore the details of these refined stratifications in full generality, as for our moduli interpretations one may concentrate on the case where one begins with the instability stratification associated to some group action. There is, however, an algorithmic procedure given in [BDHK16], which shows that this can be done. This process is described explicitly in an appendix to [BDHK16], by Hamilton, which is given here in full, at the end of this thesis.

A particularly nice case is when the minimal weight space is never contained in the centre of the blow up, at any stage of the whole process. In this case, the blow ups achieving $(\star\star\star)$ for $U^{(j)}$ terminate in a space such that the complement of the exceptional locus of the map to X may be identified with the open subset of X given

by

$$\{x \in X^0_{\min} \mid \dim \mathrm{Stab}^{(j)}_U((p(x)) = d^j_{\min}\},$$

where

$$d^j_{\min}(X) = \min\{\dim \mathrm{Stab}_{U^j}(x) \mid x \in X^0_{\min}\}.$$

We may then look inside this locus, and determine $X^{H-\widehat{ss}}$ by applying the Hilbert-Mumford criterion. This is the situation we will find ourselves in when we consider moduli of sheaves of Harder-Narasimhan length $\ell = 2$, in Section 5.3.2. On the other hand, for $\ell > 2$ we will sooner or later find ourselves in case (2) of the above, and the situation is rather more complicated.

3.3.2.3 Remark. Although we will not consider the external case in detail here, one may use all of the $\hat{U}$ theorems to perform the quotient of X by H using an external $\mathbb{G}_m$, in the following sense. We apply one of the theorems above with H replaced by $\hat{H} = H \rtimes \mathbb{G}_m$, and X replaced by $X \times \mathbb{P}^1$, linearising the action by tensoring $\mathscr{L}$ with $\mathscr{O}_{\mathbb{P}^1}(N)$ for $N \gg 0$. This yields a projective scheme, which if semistability equals stability for the unipotent radical is just $(X \times \mathbb{P}^1) /\!\!/ \hat{H}$, containing an open subscheme which is a geometric quotient of a certain H-invariant open subset $X^{H-\hat{s}} \subset X$ by H. This approach is of course particularly useful if no admissible internal $\mathbb{G}_m$ exists, for example if H is unipotent.

3.4 Variation of Non-Reductive GIT

3.4.1 Change of distinguished $\mathbb{G}_m$

Our investigation of the dependence of the quotient on the linearisation and one-parameter subgroup grading the unipotent radical begins with the choice of $\mathbb{G}_m$. For now we fix a choice of ample linearisation $\mathscr{L}$, and observe what happens when we vary the $\mathbb{G}_m$ amongst those to which $\mathscr{L}$ is (well) adapted. Throughout this section we assume that the condition $(\star)$ holds (semistability coincides with stability for the unipotent radical) for the linearisation $\mathscr{L}$. For simplicity of notation we will also assume that H is connected; this involves no loss of generality since GIT for finite group actions is independent of the choice of linearisation. The following is immediate from the Hilbert-Mumford criterion, Theorem 3.3.1.6.

3.4.1.1 Proposition. *$X /\!\!/_{\mathscr{L}} H$ is independent of the choice of internal admissible $\mathbb{G}_m$ to which $\mathscr{L}$ is (well) adapted.*

In the external case, we have no reason to suppose the quotient $(X \times \mathbb{P}^1)//\hat{H}$ to be independent of the choice of $\mathbb{G}_m$ with $\hat{H} = H \rtimes \mathbb{G}_m$, since we are quotienting by a group action that is dependent on that choice. However, we can use an argument very much in the spirit of [Tha96] to relate the two quotients. The plan is to reduce a change of external $\mathbb{G}_m$ to a change of linearisation for a certain group action on an auxiliary space.

In order to compare the quotients associated to different choices of external $\mathbb{G}_m$, we first wish to find an abelian group that simultaneously contains every choice of admissible external $\mathbb{G}_m$ up to an appropriate automorphism of X. Then given two choices of admissible external $\mathbb{G}_m$ represented by one-parameter subgroups of this abelian group, we will set up an action of $(\mathbb{G}_m)^2$ such that the H-quotient with respect to each of these choices of external $\mathbb{G}_m$ can be viewed as a quotient of $X \times (\mathbb{P}^1)^2$ by $H \rtimes (\mathbb{G}_m)^2$ with respect to some linearisation. We construct such an abelian group as follows.

By assumption H is connected, so the action of H on X yields a homomorphism $\varphi : H \to \mathrm{Aut}^0(X)$ into the connected component $\mathrm{Aut}^0(X)$ of the identity in the group $\mathrm{Aut}(X)$ of automorphisms of X. Any admissible $\mathbb{G}_m$ must by hypothesis act on X via $\psi : \mathbb{G}_m \to \mathrm{Aut}^0_H(X)$, where

$$\mathrm{Aut}^0_H(X) \coloneqq \{g \in \mathrm{Aut}^0(X) \mid g \text{ normalises } \varphi(H), \text{with trivial action on } \varphi(H)/\varphi(U)\}.$$

Moreover the existence of an admissible $\mathbb{G}_m$ implies that $\varphi|_U$ is injective.

Now $\mathrm{Aut}^0(X)$ is an algebraic group (see e.g. [Bri17a]), and hence so is $\mathrm{Aut}^0_H(X)$, but $\mathrm{Aut}^0_H(X)$ may not be linear. However, by Theorem 1 of [Bri17b], there exists a smallest normal subgroup $N_{X,H} \trianglelefteq \mathrm{Aut}^0_H(X)$ such that the quotient $\mathrm{Aut}^0_H(X)/N_{X,H}$ is affine, and moreoever $N_{X,H}$ is contained in the centre of $\mathrm{Aut}^0_H(X)$. Choose a maximal torus $T_{X,H}$ of this linear algebraic group $\mathrm{Aut}^0_H(X)/N_{X,H}$.

Given any one-parameter subgroup $\lambda : \mathbb{G}_m \to \mathrm{Aut}^0_H(X)$ of $\mathrm{Aut}^0_H(X)$, we can conjugate by some element of $\mathrm{Aut}_H(X)$ so that $\lambda(\mathbb{G}_m) \leq \pi^{-1}(T_{X,H})$, where

$$\pi : \mathrm{Aut}^0_H(X) \to \mathrm{Aut}^0_H(X)/N_{X,H}$$

is the quotient map. Thus $\pi^{-1}(T_{X,H})$ is an abelian subgroup of $\mathrm{Aut}^\circ_H(X)$, such that any admissible $\mathbb{G}_m$ is conjugate to one in $\pi^{-1}(T_{X,H})$ via an element of $\mathrm{Aut}^\circ_H(X)$.

This means that for any admissible one-parameter subgroup $\lambda : \mathbb{G}_m \to \mathrm{Aut}_H(X)$, there is an automorphism $\alpha : X \to X$ conjugating λ to $\mu : \mathbb{G}_m \to \pi^{-1}(T_{X,H})$, along which any linearisation $\mathscr{L}$ of the $\hat{U}_\lambda$ action may be pulled back, so that α induces an

isomorphism

$$X /\!\!/_{\mathscr{L}} \hat{U}_\lambda \cong X /\!\!/_{\alpha^*\mathscr{L}} \hat{U}_\mu.$$

Finally, we observe that no admissible $\mathbb{G}_m$ can lie in $N_{X,H}$: since the latter is contained in the centre of $\mathrm{Aut}^0_H(X)$, any $\mathbb{G}_m$ it contained would commute with the H-action, and hence have trivial conjugation action on U. This means that any admissible $\mathbb{G}_m$ may, up to automorphism of the whole picture, be considered to be a subgroup of $T_{X,H}$.

We can now investigate the dependence on the choice of external $\mathbb{G}_m$. Consider the action of $\hat{H}$ on $X \times \mathbb{P}^1$, linearised so as to weight the $\mathbb{P}^1$ factor heavily, as in Remark 3.3.2.3. We will show that, under certain circumstances, changing the one-parameter group $\mathbb{G}_m$ for the action of $H \rtimes \mathbb{G}_m$ on X may equivalently be viewed as changing the linearisation for an action

$$(H \rtimes (\mathbb{G}_m \times \mathbb{G}_m)) \times (X \times \mathbb{P}^1 \times \mathbb{P}^1) \to X \times \mathbb{P}^1 \times \mathbb{P}^1.$$

We begin with a linearisation $\mathscr{L}'$ of the H action on X, with respect to an underlying line bundle L'. Let T_λ and T_μ be two copies of $\mathbb{G}_m$ equipped with actions on X and H, admissible respectively for linearisations $\mathscr{L}_\lambda$ and $\mathscr{L}_\mu$ of the actions of $\hat{H}_\lambda \coloneqq H \rtimes T_\lambda$ and $\hat{H}_\mu \coloneqq H \rtimes T_\mu$ on $X \times \mathbb{P}^1$, with respect to the same underlying bundle $L = L' \boxtimes \mathscr{O}_{\mathbb{P}^1}(N)$, so that the induced linearisations for the H-action are the same. By the above we may assume that both $\mathbb{G}_m$'s lie in $T_{X,H}$ and hence commute, so that we get an induced action of their product on X. The goal is to produce an action of $T_\lambda \times T_\mu$ on $\mathbb{P}^1 \times \mathbb{P}^1$ and two linearisations $\hat{\mathscr{L}_\lambda}$ and $\hat{\mathscr{L}_\mu}$ for the resultant action of $H \rtimes (T_\lambda \times T_\mu)$ on $X \times \mathbb{P}^1 \times \mathbb{P}^1$, such that we have

$$(X \times \mathbb{P}^1 \times \mathbb{P}^1) /\!\!/_{\hat{\mathscr{L}_\lambda}} (H \rtimes (T_\lambda \times T_\mu)) \cong (X \times \mathbb{P}^1) /\!\!/_{\mathscr{L}_\lambda} \hat{H}_\lambda, \tag{3.1}$$

$$(X \times \mathbb{P}^1 \times \mathbb{P}^1) /\!\!/_{\hat{\mathscr{L}_\mu}} (H \rtimes (T_\lambda \times T_\mu)) \cong (X \times \mathbb{P}^1) /\!\!/_{\mathscr{L}_\mu} \hat{H}_\mu. \tag{3.2}$$

Having done this, it will only remain to explore the dependence of the quotient on the choice of linearisation.

3.4.1.2 Proposition. *Adopt the notation above. Then there are linearisations $\hat{\mathscr{L}_\lambda}$, $\hat{\mathscr{L}_\mu}$ of the action*

$$(H \times T_\lambda \times T_\mu) \times (X \times \mathbb{P}^1 \times \mathbb{P}^1) \to (X \times \mathbb{P}^1 \times \mathbb{P}^1)$$

extending the respective linearisations on $X \times \mathbb{P}^1$, giving isomorphisms (3.1) and (3.2).

Proof. Extend the actions of $H \rtimes T_\lambda$ and $H \rtimes T_\mu$ to actions on $X \rtimes \mathbb{P}^1_\lambda \rtimes \mathbb{P}^1_\mu$ by declaring that T_λ acts trivially on $\mathbb{P}^1_\mu$ and vice-versa. We thus obtain two linearisations of the action of $H \rtimes (T_\lambda \times T_\mu)$ with repect to the line bundle

$$\hat{L} \coloneqq L \boxtimes \mathscr{O}_{\mathbb{P}^1_\lambda}(N) \boxtimes \mathscr{O}_{\mathbb{P}^1_\mu}(N)$$

on $X \times \mathbb{P}^1_\lambda \times \mathbb{P}^1_\mu$: one by tensoring $\mathscr{L}_\lambda$ with the pullback of the trivial linearisation on $\mathscr{O}_{\mathbb{P}^1_\mu}$, and the other with the roles of λ and μ reversed. We tensor these together, and call the result $\widehat{\mathscr{L}}$.

Now consider the weight lattice of $T_\lambda \times T_\mu$ with respect to this linearisation. Using a Segre embedding into a projective space $X \times \mathbb{P}^1_\lambda \times \mathbb{P}^1_\mu \hookrightarrow \mathbb{P}^s$, we find that for the linearisation decribed above the weights are arranged in four clusters, each placed at one of the four corners of a square. Each cluster is a small copy of the weights of X, and the weights in it correspond to taking the coordinates of X in conjunction with one of the four pairs $y_i z_j$ for $i, j = 0, 1$, where $(y, z) \in \mathbb{P}^1_\lambda \times \mathbb{P}^1_\mu$. Let $\widehat{\mathscr{L}_\lambda}$ be the twist of $\widehat{\mathscr{L}}$ by the $T_\lambda \times T_\mu$-character $(0, N + r_\lambda - \varepsilon)$, and let $\widehat{\mathscr{L}_\lambda}$ be the twist of $\widehat{\mathscr{L}}$ by the character $(N + r_\mu - \varepsilon, 0)$, where r_λ and r_μ are respectively the minimal weights on $\mathscr{L}_\lambda$ and $\mathscr{L}_\mu$ for their distinguished choices of $\mathbb{G}_m$.

Consider the quotient with respect to $\hat{\mathscr{L}_\lambda}$; the quotient with respect to $\hat{\mathscr{L}_\mu}$ works in the same way. By Proposition 3.4.1.1, the internal quotient is independent of the choice of 1PS used to form $\hat{U}$, so we may choose $\hat{U}$ to be $\hat{U}_\lambda = U \rtimes T_\lambda$. Then $\hat{U}_\lambda$ acts trivially on the third factor, and by our choice of character the quotient is

$$(X \times \mathbb{P}^1_\lambda \times \mathbb{P}^1_\mu) /\!\!/_{\hat{\mathscr{L}_\lambda}} \hat{U}_\lambda \cong ((X \times \mathbb{P}^1_\lambda) /\!\!/_{\mathscr{L}_\lambda} \hat{U}_\lambda) \times \mathbb{P}^1_\mu.$$

We follow this by the residual reductive quotient by $R \times T_\mu$, which we may factor as first a quotient by T_μ, then a quotient by R. We observe that the T_μ quotient is just the first factor of the product, since the stable (and semistable) locus is

$$(((X \times \mathbb{P}^1_\lambda) /\!\!/_{\mathscr{L}_\lambda} \hat{U}_\lambda) \times \mathbb{P}^1_\mu)^{s, T_\mu} = ((X \times \mathbb{P}^1_\lambda) /\!\!/_{\mathscr{L}_\lambda} \hat{U}_\lambda) \times (\mathbb{A}^1_\mu \setminus \{0\}).$$

In addition, the induced linearisation of the R-action on this first factor is equal to (a positive scalar multiplie of) the R-linearisation induced by $\mathscr{L}_\lambda$. So we have

$$((X /\!\!/_{\mathscr{L}_\lambda} \hat{U}_\lambda) \times (\mathbb{P}^1_\mu)) /\!\!/ (R \times T_\mu) \cong (X /\!\!/_{\mathscr{L}_\lambda} \hat{U}_\lambda) /\!\!/ R \cong (X \times \mathbb{P}^1) /\!\!/_{\mathscr{L}_\lambda} \hat{H}_\lambda.$$

This gives us the isomorphism (3.1); the isomorphism (3.2) follows in the same way. □

3.4.1.3 Corollary. *Changing between two admissible external $\mathbb{G}_m$'s, for a fixed linearisation well adapted to both, is equivalent to a change of linearisation.*

3.4.2 Change of linearisation

Continuing with our analysis, we now tackle the other half of the question. We continue to assume that we have semistability coincides with stability for the unipotent radical. Fix an admissible one-parameter group $\mathbb{G}_m$, internal or external, and consider the space of linearisations well adapted to it. The following observation is immediate from Theorem 3.3.1.6(3).

3.4.2.1 Proposition. *Let $\mathscr{L}_1$ and $\mathscr{L}_2$ be two well adapted linearisations of the H-action on X, and suppose that their restrictions to a maximal torus T of H give the same T-semistable locus. Then $X^{H-ss}(\mathscr{L}_1) = X^{H-ss}(\mathscr{L}_2)$, and the quotients $X /\!\!/_{\mathscr{L}_1} H$ and $X /\!\!/_{\mathscr{L}_2} H$ can be canonically identified.*

This gives us part of the story: if we consider the structure of the space of linearisations for the action of the maximal torus T of H, and identify a linearisation of the H-action with the linearisation of the T-action that it induces, then by Proposition 3.4.2.1, if two linearisations lie in the same chamber (or more generally are GIT-equivalent for the action of T in the sense of Dolgachev & Hu), then they must produce the same H-quotient. What we have not yet determined is how the quotient may change when we cross a wall (or more generally move from one GIT-equivalence class to another). We consider first the situation when $H = \hat{U}$.

3.4.2.2 Proposition. *The quotient $X /\!\!/_{\mathscr{L}} \hat{U}$ is independent of the choice of ample well adapted linearisation $\mathscr{L}$.*

Proof. Recall that the $\hat{U}$ quotient admits a description as the geometric quotient of

$$X^{\hat{U}-s} = X \setminus U(X \setminus X^{s-\mathbb{G}_m})$$

by $\hat{U}$, where $X^{\hat{U}-s}$ is an open subscheme of X dependent only on the $\hat{U}$-action and the stable locus for our chosen one parameter group $\lambda : \mathbb{G}_m \to \hat{U}$ which grades U, and any two choices for λ are conjugate in $\hat{U}$. It suffices to show that $X^{s,\mathbb{G}_m}$ is independent of the choice of linearisation $\mathscr{L}$.

One way to argue this is the following (compare Theorem 4.3 of [BB73] on the Bialynicki-Birula stratification). The one-parameter subgroup $\lambda : \mathbb{G}_m \to \hat{U}$ determines a map of sets $p : X \to X$ defined by

$$p(x) \coloneqq \lim_{t\to 0} \lambda(t)x \in X$$

taking each point to its limit under the flow given by the action of $\lambda : \mathbb{G}_m \to \hat{U}$, and $p(X) = X^{\mathbb{G}_m}$ is the fixed point set for the action of λ.

Let $Z_1, \ldots, Z_n$ be the connected components of $X^{\mathbb{G}_m}$. By assumption X is irreducible, so $Z_{\min}$ is the union of those Z_j for which $p^{-1}(Z_j)$ is open in X, and $X^{s,\mathbb{G}_m} = p^{-1}(Z_{\min}) \setminus Z_{\min}$ is independent of $\mathscr{L}$. □

In general, two different H-linearisations will induce different linearisations of the residual $\bar{R} \coloneqq R/\mathbb{G}_m$ action on $X /\!\!/ \hat{U}$, and hence the two H quotients for change of linearisation will be related by a reductive VGIT picture applied to $X /\!\!/ \hat{U}$. This allows us to recover, for fixed $\mathbb{G}_m$, a wall-and-chamber structure to the space of well adapted linearisations, namely that of reductive VGIT applied to the T-action, so that two quotients are related by Thaddeus flips of $X /\!\!/ \hat{U}$.

3.4.3 When semistability does not coincide with stability

So far we have been considering the situation covered by Theorems 3.3.1.6 and 3.3.1.10 when we can construct quotients of the scheme X (or $X \times \mathbb{P}^1$ in the external case) without needing to do the blow-up process of Theorems 3.3.1.8 and 3.3.1.12. As described in [BDHK16], this process begins by defining

$$d_0 = \max\{\dim \mathrm{Stab}_U(x) \mid x \in X^{0,\mathbb{G}_m}_{\min}\},$$

and then blowing up along the closure in X of

$$C_0 = \{x \in X^{0,\mathbb{G}_m}_{\min} \mid \dim \mathrm{Stab}_U(x) = d_0\}.$$

Iteratively, one then sets $X_{(0)} = X$ and defines $X_{(j)}$ to be the blow-up along the closure in $X_{(j-1)}$ of

$$C_j = \{x \in X^{0,\mathbb{G}_m}_{\min} \mid \dim \mathrm{Stab}_U(x) = d_j\},$$

where

$$d_j = \max\{\dim \mathrm{Stab}_U(x) \mid x \in (X_j)^{0,\mathbb{G}_m}_{\min}\}.$$

This procedure is clearly dependent on the choice of $\mathbb{G}_m$, as written. However, one may describe it in a different way so that the results above will apply to these sorts of quotients, too. Instead of blowing up along the loci described, we let

$$X^0 = \bigcup_{\mathbb{G}_m \in \mathfrak{C}} X^{0,\mathbb{G}_m}_{\min}.$$

This is the set of points in X that are in the basin of attraction to the minimal weight space for some admissible $\mathbb{G}_m$. Then we can apply the procedure described in [BDHK16] with $X^{0,\mathbb{G}_m}_{\min}$ replaced by X^0 to achieve a 'universal' blow-up, $\check{X}$, that is not dependent on any choice. For each admissible one-parameter group $\mathbb{G}_m$ in $\mathfrak{C}$, there is an open subscheme of $\check{X}$ that is isomorphic to the open subset $(\hat{X}_{\mathbb{G}_m})^{0,\mathbb{G}_m}_{\min} \subset \hat{X}_{\mathbb{G}_m}$, where the latter is the space obtained from X by applying the procedure above using $X^{0,\mathbb{G}_m}_{\min}$. All quotients obtained from the blow-up process using different choices of $\mathbb{G}_m$ can be regarded as quotients of the 'master space' $\check{X}$ using different choices of $\mathbb{G}_m$, and hence the quotients are related as already described in the cases covered by Theorems 3.3.1.6 and 3.3.1.10.

3.4.4 Summary of NRVGIT results

As we have seen, the main difference between the parameter spaces for reductive and non-reductive GIT is the additional choice of a one-parameter group $\mathbb{G}_m$. This can be chosen from the admissible cone $\mathfrak{C} \subset \mathfrak{z}$, as defined at Definition 3.3.1.1, in the internal case, or from a corresponding cone $\mathfrak{C} = \mathfrak{C}_{X,H}$ in the external case. Without loss of generality we may assume that semistability coincides with stability for the unipotent radical; if not we can work with the universal blow-up described in Section 3.4.3. Having made a choice of one-parameter group $\mathbb{G}_m$, we must choose a linearisation well adapted to this choice.

3.4.4.1 Definition. Given a choice of admissible internal $\lambda : \mathbb{G}_m \to H$, let $\partial_\lambda \subset \mathfrak{C}$ be the set of linearisations that are well adapted to λ. Let

$$\partial = \bigcup_{\lambda \text{ admissible}} \partial_\lambda$$

be the set of linearisations that are well adapted for some admissible internal $\mathbb{G}_m$.

The set ∂_λ is a suitably small open region along the boundary of the intersection of $\mathfrak{C}$ with a region bounded by several hyperplanes: one giving the negativity condition for the minimal weight space for $\lambda : \mathbb{G}_m \to H$ (which by Proposition 3.4.2.2 is the same for all ample linearisations), and the others giving the positivity conditions for the other weight spaces. In particular, ∂_λ is the intersection with a convex set of an open neighbourhood of its boundary in $\mathfrak{C}$. While this region evidently depends on the choice of $\mathbb{G}_m$, we can recover ∂ as the union of ∂_λ for only finitely many λ.

We may summarise our results as follows.

3.4.4.2 Theorem. ***(Internal Case)*** *Let X be a projective scheme, acted on by a linear algebraic group H.*

1) *Suppose that condition $(\star)$ or $(\star\star\star)$ is satisfied. Then the quotient $X /\!\!/_{\mathscr{L}} H$ is independent of the choice of admissible internal $\mathbb{G}_m$, and may be formed using any linearisation $\mathscr{L} \in \partial$ which is well adapted to some admissible $\mathbb{G}_m$. Moreover if two linearisations are GIT-equivalent for the action of a maximal torus $T \leq H$ then the associated quotients are isomorphic. Further, any two quotients with respect to different linearisations, which are both well adapted for the same one-parameter subgroup, are related by the usual reductive variation of GIT picture induced via the $\hat{U}$ quotient map for the relevant actions of $R/\mathbb{G}_m$ on the quotient $X /\!\!/ \hat{U}$ (where $X /\!\!/ \hat{U}$ is independent of the choice of linearisation).*

2) *If we do not have $(\star)$ or $(\star\star\star)$, then there is a sequence of H-equivariant blow ups of X along closed H-invariant subvarieties resulting in a scheme $\widetilde{X}$ such that either $(\star)$ or $(\star\star\star)$ holds, as desired, and the results of the first two parts of this Theorem apply.*

Note in particular that there are only finitely many possible GIT quotients yielded by the $\hat{U}$ theorems with internal $\mathbb{G}_m$, and we may move between any two of these by a sequence of changes of linearisation, if necessary changing the $\mathbb{G}_m$ to maintain well adaptedness (though without changing the quotient). We therefore interpolate between any two H-quotients by a sequence of 'vertical' and 'horizontal' moves in $\mathfrak{C} \times \partial$, giving a sequence of points such that any two consecutive points have either the linearisation or the one-parameter subgroup $\mathbb{G}_m$ in common.

3.4.4.3 Theorem. ***(External case)*** *Let X be a projective scheme acted on by a linear algebraic group H. Form the non-reductive GIT quotient*

$$X \widehat{/\!\!/} H = (X \times \mathbb{P}^1) /\!\!/ \hat{H}$$

with respect to a choice of admissible external $\mathbb{G}_m \leq \operatorname{Aut} H$ *and a well adapted ample linearisation* $\mathscr{L}$ *by letting* $\hat{H} = H \rtimes \mathbb{G}_m$ *act on* $X \times \mathbb{P}^1$*, linearised with respect to* $\mathscr{L} \boxtimes \mathscr{O}_{\mathbb{P}^1}(N)$ *for* $N \gg 0$*.*

1) *There is an abelian subgroup* $\pi^{-1}(T_{X,H})$ *of* $\operatorname{Aut}(X)$ *such that any admissible external* $\mathbb{G}_m$ *is GIT-equivalent to a one-parameter subgroup* $\mathbb{G}_m$ *of* $\pi^{-1}(T_{X,H})$*. Furthermore external one-parameter groups* $\mathbb{G}_m$ *yield the same non-reductive quotient* $X \hat{/\!/} H$ *for the same linearisation if they have the same minimal weight space in* X*.*

2) *The dependence of the quotient* $X \hat{/\!/} H$ *on the linearisation can be described as in the previous theorem.*

3.4.5 Points and lines in $\mathbb{P}^2$

We conclude with a brief description of a concrete example, to illustrate these ideas in practice. We consider configurations of points and lines in the projective plane up to an action of a certain non-reductive subgroup of SL(3). We take V to be a 2-dimensional complex vector space, and let

$$X = \mathbb{P}(V)^p \times \mathbb{P}(V^*)^q,$$

so that a point of X corresponds to a choice of p points and q lines in the projective plane.

Our first choice for the group in question is

$$H = \left\{ \begin{pmatrix} a & b & c \\ 0 & d & 0 \\ 0 & 0 & f \end{pmatrix} \in \mathrm{SL}(3) \right\},$$

and the unipotent radical and our chosen maximal torus are, respectively,

$$U = \left\{ \begin{pmatrix} 1 & b & c \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \in \mathrm{SL}(3) \right\}$$

$$T = \left\{ \begin{pmatrix} t & 0 & 0 \\ 0 & s & 0 \\ 0 & 0 & (st)^{-1} \end{pmatrix} \in \mathrm{SL}(3) \right\}.$$

Note that in this example H is just the semidirect product of a torus and the unipotent part, so the torus is a quotient group and hence all its characters extend to the whole group H. The action of H on X is that induced by the natural action on V, and the dual action on V^*. Accordingly, the weight polytope is a scaled version of the hexagon familiar from the study of the representation theory of SL(3). Take $p = 2$, $q = 1$ for concreteness.

The admissible cone, shaded in blue, is larger than the standard positive Weyl chamber in SL(3), because we do not have the full Borel subgroup, and so we only need positivity with two roots (cf. [BDHK16] 7.1). The well adapted linearisations correspond to the section of the boundary indicated in black. We observe also that for this H there is a natural choice for the 1PS, given by the solid blue arrow: a universal choice in the sense that any linearisation which is allowable for some admissible $\mathbb{G}_m$ will be allowable for this one. It should be noted that this does not hold in general. For example, if instead we impose also the additional condition that $b = 0$, so that H is the semidirect product of $\mathbb{G}_a$ and the maximal torus, the admissible cone is enlarged and there is no longer a universal choice of 1PS.

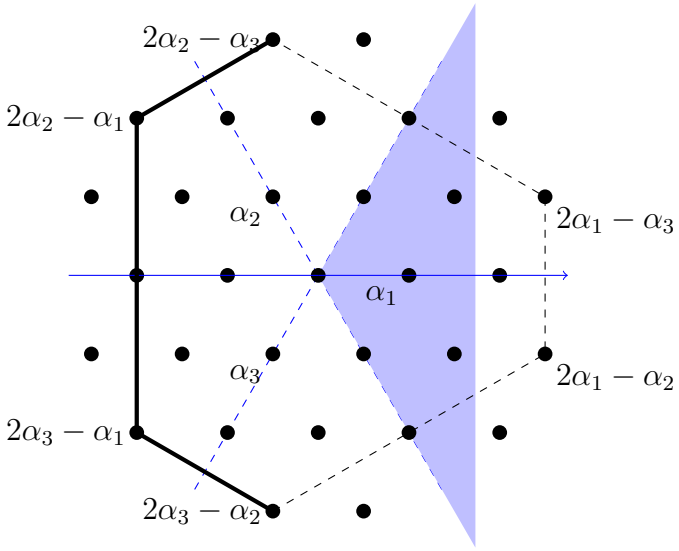


Figure 3.1: Weight diagram and admissible cone for the action of H on $(\mathbb{P}^2)^2 \times (\mathbb{P}^2)^*$

Chapter 4

Moduli Spaces of Unstable Objects

We will now show that the results of the previous two chapters may be used to construct moduli spaces of unstable objects. Mathematically, this merely involves pulling together the things we have said so far into an explicit overall strategy. Nevertheless this chapter is the conceptual heart of the thesis, and forms the basis for what we will say in Chapters 5 and 6, where we will see how this strategy will be carried out in two major examples.

The set-up we will consider is as follows. Suppose we have set up some scheme X with a G-action[1] whose orbits correspond to isomorphism classes of whatever objects we wish to construct a moduli space for. After a choice of linearisation, Mumford's Geometric Invariant Theory (or a non-reductive equivalent) gives us a quotient, in the appropriate sense, of an open subset consisting of semistable points. From a certain point of view, however, this is unsatisfactory. The moduli classification thus obtained says nothing at all about those objects whose corresponding points were unstable, and while there are good geometric reasons not to try to include these objects in the same moduli space as the semistable ones, there is usually no good reason to ignore them altogether. Rather than simply throw away the unstable points, a more natural thing to do is to regard the classical quotient (that is, the moduli space of semistable objects) as merely the first step in the solution of the moduli problem, in the following way. Just as GIT tells us, in effect, which objects we can include in the moduli space and still get a reasonable[2] answer, it also tells us, via the instability stratification, how we should divide the problem further in order to make sense of the unstable part. Then we can try to find, for every instability index $\beta \in \mathscr{B}$, a

[1] We will assume for the moment that G is reductive, though, as we discussed in Remark 2.3.0.17, we may also attempt this in the non-reductive case - and indeed, we will do so in Chapter 6.

[2] Finite type, separated, projective &c ...

quotient of the stratum S_β by the G-action. By Theorem 2.3.0.15, this is equivalent to the problem of finding a quotient for the action of P_β on Y_β^{ss}. However, P_β *is in general non-reductive*. Hence, we will need to make use of the recent progress in non-reductive GIT described in Chapter 3, and of the $\hat{U}$ theorems of Section 3.3 in particular. These theorems were proved with moduli of unstable objects very much in mind, and for this reason they can readily be applied in this setting. To summarise, the strategy is as follows:

I) Pose the given moduli problem as a GIT problem.[3]

II) Interpret the instability stratification associated to this GIT problem in an intrinsic way. That is: convert the purely geometric-invariant-theoretic notion of instability type into a moduli-theoretic instability type, analogous to the Harder-Narasimhan type in the case of sheaves, for the objects under consideration. This gives a discrete classification, analogous to fixing the Hilbert polynomial or similar invariant.

III) Use non-reductive GIT to quotient the unstable strata, obtaining moduli spaces of objects of fixed instability type. Depending on the particular problem, and which of the $\hat{U}$ Theorems we are led to use, this may be straightforward or somewhat challenging. In particular, we may have to further refine the instability stratification using the blow up process described in Section 3.3.2.

The output of this is a moduli problem that has been split into (usually an infinite number of) manageable pieces, each with its own moduli space. Each isomorphism class of object (curve, sheaf &c) should appear in one, and only one, of these pieces.

4.1 Reductive quotients of the unstable strata

We begin by motivating the use the $\hat{U}$ theorems in this setting, by illustrating what happens if we try to quotient the unstable strata using reductive GIT. This was the approach taken in [Hos11], before an adequate non-reductive GIT was available. Recall that in the best setting for GIT we have a reductive group G acting on a

[3]Of course, the sets of objects in which we are interested will usually come in unbounded families, and in any given GIT-setup we will only see finitely many strata. So strictly speaking what we really need is an infinite family of GIT problems that somehow approximate our moduli problems asymptotically, so that we can work with an asymptotic instability stratification.

projective scheme X with ample linearisation $\mathscr{L}$, where the quotient morphism

$$X^{ss} \to X /\!\!/ G$$

is induced via the Proj functor by the inclusion of rings

$$\bigoplus_{r\geq 0} H^0(X, \mathscr{L}^{\otimes r})^G \hookrightarrow \bigoplus_{r\geq 0} H^0(X, \mathscr{L}^{\otimes r}).$$

Outside of this setting, say when L is not ample or G is not reductive, we can still consider the morphism of schemes resulting from this inclusion, but this will not necessarily coincide with the usual GIT quotient in the sense of Mumford. Moroever, if the invariants are not finitely generated, the scheme obtained thereby will not be of finite type. A more tractable quotient may be obtained, following [Hos11], by taking a finitely generated subalgebra

$$R \subseteq \bigoplus_{r\geq 0} H^0(X, \mathscr{L}^{\otimes r})^G$$

and then defining the 'generalised GIT quotient' to be the variety

$$X /\!\!/_{\mathscr{L},R} G \coloneqq \operatorname{Proj} R.$$

This comes equipped with a natural rational map from the scheme X, defined by the corresponding inclusion of subrings.

Let $\overline{Y}_\beta$ be the closure of Y_β inside the ambient projective space, and define

$$\hat{S}_\beta = G \times^{P_\beta} \overline{Y}_\beta.$$

In [Hos11], Hoskins constructs a linearisation on $\hat{S}_\beta$, which is also denoted by $\mathscr{L}_\beta$ and also called the canonical linearisation (cf. Definition 2.3.0.11). This has the property that there are isomorphisms of algebras

$$\bigoplus_{n\geq 0} H^0(\hat{S}_\beta, L_\beta^{\otimes n})^G \cong \bigoplus_{n\geq 0} H^0(\overline{Y}_\beta, L_\beta^{\otimes n})^{P_\beta} \cong \bigoplus_{n\geq 0} H^0(Z_\beta, L_\beta^{\otimes n})^{\mathrm{Stab}\beta}.$$

By the reductivity of $\mathrm{Stab}\beta$, we know that the invariant sections on $\hat{S}_\beta$ are finitely generated. The canonical linearisation on $\hat{S}_\beta$ is therefore particularly amenable to the strategy of 'generalised GIT'– we can take R to be, in each case, the whole subalgebra of invariants. (Let us abuse notation and denote them all by R, since they

are isomorphic.) On the other hand, this linearisation has the drawback that it is not usually ample.[4] At the heart of the problem lies the fact that P_β is not reductive, and in attempting to get by without using non-reductive GIT we have been forced to consider non-ample linearisations, which we know do not behave well. The result of this approach is summarised as follows.

4.1.0.1 Theorem. *([Hos11], 2.2)*

1) The semistable locus for the P_β-action on $\overline{Y}_\beta$ is Y_β^{ss}, and we have isomorphisms of generalised GIT quotients

$$\hat{S}_\beta /\!\!/_{(\mathscr{L}_\beta,R)} G \cong \overline{Y}_\beta /\!\!/_{(\mathscr{L}_\beta,R)} P_\beta \cong Z_\beta /\!\!/_{\mathscr{L}_\beta} Stab\beta.$$

2) The quotient map

$$Y_\beta^{ss} \to \overline{Y}_\beta /\!\!/_{(\mathscr{L}_\beta,R)} P_\beta$$

factors through the morphism $p_\beta : Y_\beta \to Z_\beta$.

3) The morphism $S_\beta \to \overline{Y}_\beta /\!\!/_{(\mathscr{L}_\beta,R)} P_\beta$ is a categorical quotient for the action of G on S_β.

The first point tells us the good news: in taking a quotient with respect to the canonical linearisation on S_β we can essentially treat the problem as if it were classical GIT. That is, the quotient we are looking for is equivalent to a quotient of the action of the reductive group Stabβ on Z_β with respect to the ample canonical linearisation.

However, the second point shows that each point of Y_β^{ss} is identified in the quotient with its image under the retraction $p_\beta : Y_\beta \to Z_\beta$. This is generally undesirable, and leads to GIT quotients being of much smaller dimension than they ought to be. For example, in the setting of moduli of unstable sheaves, this results in a sheaf being identified with the associated graded of its Harder-Narasimhan filtration; this means that the moduli spaces of unstable sheaves constructed in this way would be isomorphic to products of moduli spaces of semistable sheaves, with no information about the extensions that glue them together; this is plainly not satisfactory. A more straightforward illustration, where the moduli spaces resulting from this method of construction are just points, is given in the example below, and we will see in Section 4.1.1 how we can use non-reductive GIT to improve on this.

If we are to find a quotient morphism that does not factor through p_β, we must somehow make Z_β unstable, otherwise we will have the problem that too many orbit

[4] See [Hos11], Example 2.2.5

closures meet there. Hence, as suggested in [Hos11], a good approach would be to perturb the canonical linearisation, thus splitting each stratum into two pieces: a piece corresponding to Z_β, whose associated moduli space is the reductive quotient by Stab β above, and the rest,[5] which we will quotient using non-reductive GIT with respect to this perturbation of the canonical linearisation on $\overline{Y}_\beta$. The problem arises, however, that perturbing the linearisation may potentially destroy the resemblance to a classical GIT problem noted in the theorem above, and creates a need for the methods of non-reductive GIT.

4.1.1 Toy Example: unordered points on $\mathbb{P}^1$

Our aim here is to illustrate ideas by means of an elementary example from classical GIT, namely the moduli space of unordered points on the projective line up to automorphisms of the line. We proceed as follows, which provides an outline for how such problems would be tackled in general.

1) The first step is to apply the results of the first chapter. We choose a maximal torus $T \subset G$ and find the weights of the action of T. Using the Hilbert-Mumford function we find the semistable locus.

2) Next we find the indexing set $\mathscr{B}$ for the stratification, and for each $\beta \in \mathscr{B}$ the sets Z_β and Y_β, together with the 1PS λ_β and the parabolic subgroup P_β.

3) We can then identify S_β as the G-sweep of Y_β^{ss}, giving us the stratification. By the isomorphism
$$S_\beta \cong Y_\beta^{ss} \times^{P_\beta} G,$$
a quotient of a given stratum S_β by the G-action is equivalent to a quotient of the corresponding Y_β by the parabolic P_β.

4) We demonstrate the approach of [Hos11] in twisting the linearisation on Z_β by the character $-\beta$ to produce the canonical linearisation, and show that this has undesirable effects. This illustrates the need for a perturbation of the canonical linearisation.

From now until Section 4.2, we consider $G \coloneqq SL(2)$ and the natural action on
$$X \coloneqq Sym^n(\mathbb{P}^1) \cong \mathbb{P}^n.$$

[5] This will be $Y_\beta^{ss} \setminus U_\beta Z_\beta^{ss}$; compare to $X^0_{\min} \setminus UZ_{\min}$ in Section 3.3.

This admits the following description: consider X as the space of degree n homogeneous forms $F(X,Y)$ on $\mathbb{P}^1$, with the zeros of the form giving the choice of points. This comes equipped with the action of G induced by the dual action on $\mathbb{P}^1$. That is,

$$g \cdot F(X,Y) = F(g^{-1} \cdot (X,Y)).$$

The linearisation $\mathscr{L}$ is that induced from the standard $\mathscr{O}(1)$-linearisation on $\mathbb{P}^1$. Our chosen maximal torus T is identified with $\mathbb{G}_m$ via the isomorphism

$$\begin{aligned} T &\to \mathbb{G}_m \\ \begin{pmatrix} t & 0 \\ 0 & t^{-1} \end{pmatrix} &\longmapsto t, \end{aligned}$$

acting on X diagonally with respect to our chosen coordinates:

$$t \cdot X^i Y^{n-i} = t^{n-2i} X^i Y^{n-i}.$$

The weights of T are

$$\{\alpha_i \coloneqq n - 2i \in \mathbb{Z} \mid i = 0, ...n\}.$$

We observe that if more than half of the points are at zero, there will be no negative weight, and if more that half are at infinity, there will be no positive weight. We therefore find that the first level of our stratification, the G-semistable locus, is:

$$S_0 = X^{ss} = \{x \in X \mid \text{not more than half of the } n \text{ points coincide anywhere}\}.$$

We choose the positive Weyl chamber $\mathfrak{t}^+ \cong \mathbb{R}_{\geq 0}$, and so the indexing set is:

$$\mathscr{B} = \Big\{0\Big\} \cup \Big\{2i - n \mid \frac{n}{2} < i \leq n\Big\}.$$

Immediately from the definition we compute that the parabolics and stabilisers are the same for every member of the indexing set, so that for all i we have:

$$P_{n-2i} = P \coloneqq \left\{ \begin{pmatrix} a & b \\ 0 & a^{-1} \end{pmatrix} \mid a \in \mathbb{C}^*, b \in \mathbb{C} \right\}$$

$$\mathrm{Stab}(n-2i) = T.$$

The Z_β's each consist of a Stabβ-invariant single point:

$$Z_{n-2i} = \{x \in X \mid \text{ Exactly } i \text{ of the points are at } \infty \text{ and the rest are at } 0\},$$

and we find:

$$Y_{n-2i} = \{x \in X \mid \text{ Exactly } i \text{ of the points are at } \infty\},$$

which we can check directly is P-invariant. Now we consider the shifted linearisation $\mathscr{L}_\beta$ on a specific stratum, as used in [Hos11]. For a fixed i we have the T-character

$$\chi_{n-2i} : T \to \mathbb{G}_m$$

$$\begin{pmatrix} t & 0 \\ 0 & t^{-1} \end{pmatrix} \longmapsto t^{n-2i}$$

and we define $\mathscr{L}_{n-2i}$ to be the linearisation of the G-action given by twisting $\mathscr{L}$ by χ_{n-2i}. The semistable locus of Z_β with respect to $\mathscr{L}_{n-2i}$ is

$$Z^{ss}_{n-2i} = Z_{n-2i} = \{\text{pt}\}$$

since the weight of the Stab$(n-2i) = T$ action on the fibre over the single point is now zero. Hence also, for each $\beta_i = (n-2i) \in \mathscr{B}$, we have

$$Y^{ss}_{n-2i} = Y_{n-2i}.$$

Taking the projective closure yields:

$$\overline{Y}_{n-2i} = \{x \in X \mid \text{At least } i \text{ of the points are at } \infty\} \cong \mathbb{P}^{n-i},$$

so, taking the G-sweep of the Y^{ss}_{n-2i}, the stratification is:

$$S_{n-2i} = \{x \in X \mid \text{Exactly } i \text{ of the points coincide.}\}$$

Recall that the quotient morphism for each β-stratum factors through the retraction

$$\rho : Y^{ss}_\beta \to Z^{ss}_\beta.$$

Hence we see that following the method in [Hos11], the generalised GIT quotient of each stratum is just a point. Clearly, this is not a satisfactory quotient. We will

return to this example later, and apply some of the results of Non-reductive GIT to see if the situation can be improved.

4.2 Quotients using Non-Reductive GIT

Let us now show that we can get satisfactory quotients of the unstable strata by means of the $\hat{U}$ theorems. Let G be a reductive group acting on a projective scheme X with respect to an ample linearisation $\mathscr{L}$, so that as in Section 2.3, associated to a choice of norm $||\cdot||$ on $\operatorname{Lie} G$ there is an instability stratification

$$X = \bigsqcup_{\beta \in \mathcal{B}} S_\beta.$$

First, we verify that the non-reductive group P_β has graded unipotent radical.

4.2.0.1 Lemma. *The parabolic subgroup $P_\beta = U_\beta \rtimes L_\beta$ has graded unipotent radical U_β, where the grading $\mathbb{G}_m$ is given by the central 1-PS $\lambda_\beta : \mathbb{G}_m \to L_\beta$.*

Proof. By definition $P_\beta = P(\lambda_\beta)$ as in 2.3. It follows that the weights of the conjugation action of $\lambda_\beta(\mathbb{G}_m)$ on $\operatorname{Lie} P_\beta$ are non-negative and are precisely zero on the Levi factor L_β. Hence, all the weights of $\lambda_\beta(\mathbb{G}_m)$ on $\operatorname{Lie} U_\beta$ are strictly positive. □

In order to apply the $\hat{U}$ theorems to the P_β-action on Y_β^{ss}, we first need to compactify Y_β^{ss}. Let $\overline{Y_\beta}$ denote the closure of the locally closed subscheme $Y_\beta \subset X$. We recall that we gave explicit descriptions of Y_β and Z_β at Section 2.3 in terms of torus weights.

4.2.0.2 Proposition. *Let $\beta \in \mathcal{B}$ be a non-zero index of an instability stratum $S_\beta \subset Y$. For the P_β-action on $\overline{Y_\beta}$ with respect to the linearisation $\mathscr{L}$ and the grading given by $\lambda_\beta : \mathbb{G}_m \to L_\beta$ we have*

$$Z(\overline{Y_\beta})_{\min} = Z_\beta \quad \textit{and} \quad (\overline{Y_\beta})^0_{\min} = Y_\beta.$$

Moreover, for $\epsilon \in \mathbb{Q}_{>0}$ sufficiently small, the twisted linearisation $\mathscr{L}^{\chi-(1+\epsilon)\beta}$ is adapted for the graded unipotent group $\widehat{U}_\beta = U_\beta \rtimes_{\lambda_\beta} \mathbb{G}_m$.

Proof. Since for a G-equivariant closed immersion $i : X \hookrightarrow \mathbb{P}^n$ such that $\mathscr{L} = i^*\mathcal{O}_{\mathbb{P}^n}(1)$, we have $Z(Y)_{\min} = X \times_{\mathbb{P}^n} Z(\mathbb{P}^n)_{\min}$ and similarly for the instability strata, it suffices to prove that $Z(\overline{Y_\beta})_{\min} = Z_\beta$ when $X = \mathbb{P}^n$ and $\mathscr{L} = \mathcal{O}_{\mathbb{P}^n}(1)$. We can

choose a maximal torus T containing $\lambda_\beta(\mathbb{G}_m)$ and coordinates on $\mathbb{P}^n$ such that the action of the fixed maximal torus $T \subset G$ is given by

$$t \mapsto \mathrm{diag}(\alpha_0(t), \ldots, \alpha_n(t))$$

for characters $\alpha_i : T \to \mathbb{G}_m$. Recall that we identify characters and co-characters of T via our invariant inner product on $\mathfrak{t}$. The weights of the $\lambda_\beta(\mathbb{G}_m)$-action on $H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(1))^*$ are given by the projections $\frac{\alpha_i \cdot \beta}{||\beta||^2}\beta$ of the T-weights $\alpha_i \in \mathfrak{t}$ to the line spanned by $\beta \in \mathfrak{t}$. Thus to determine the minimum $\lambda_\beta(\mathbb{G}_m)$-weight $\omega_{\min}$, we want to minimise $\alpha_i \cdot \beta$ over all weights α_i of points in Y_β. As $\beta \in \mathcal{B}$, we have $Z_\beta \neq \emptyset$ and so from the descriptions of $\overline{Y_\beta}$ and Z_β we see that there is a point in Z_β with at least one T-weight α_i satisfying $\alpha_i \cdot \beta = ||\beta||^2$. The weights α_j of points in $\overline{Y_\beta}$ all satisfy $\alpha_j \cdot \beta \geq ||\beta||^2$. Hence, the minimal $\lambda_\beta(\mathbb{G}_m)$-weight corresponds to $\beta \in \mathfrak{t}$ and by definition $Z(\overline{Y_\beta})_{\min} \subset X^{\mathbb{G}_m}$ is the minimal $\lambda_\beta(\mathbb{G}_m)$-weight space (cf. Definition 3.3.1.3). That is,

$$Z(\overline{Y_\beta})_{\min} = \{[p_0 : \cdots : p_n] : p_i = 0 \text{ if } \alpha_i \cdot \beta \neq ||\beta||^2\},$$

which is equal to Z_β. From the weight description of Y_β, we see that Y_β is the open subscheme of $\overline{Y_\beta}$ consisting of points y such that $\lim_{t\to 0} \lambda_\beta(t)\cdot y \in Z_\beta$, and so it follows that $(\overline{Y_\beta})^0_{\min} = Y_\beta$. We note that if we twist $\mathscr{L}$ by any scalar multiple of the character corresponding to β, then the subsets $(\overline{Y_\beta})^0_{\min}$ and $Z(\overline{Y_\beta})_{\min}$ remain unchanged.

For the final statement, we recall that the (rational) character

$$\chi_{-(1+\epsilon)\beta} : P_\beta \to \mathbb{G}_m$$

merely shifts the $\lambda_\beta(\mathbb{G}_m)$ weights by $-(1+\epsilon)||\beta||$. In order for the $\widehat{U}_\beta$-action with respect to this twisted linearisation to be adapted, we need the origin to separate the minimal weight from all the other weights. Since the minimum weight for the original linearisation $\mathscr{L}$ corresponds to β, we need to shift this weight just beyond the origin, so that this weight is negative and all the other weights are positive. This is achieved by taking $\epsilon \in \mathbb{Q}_{>0}$ sufficiently small. □

4.2.0.3 Definition. We fix $\epsilon \in \mathbb{Q}$ sufficiently small as above, so the perturbation

$$\mathscr{L}_\beta^{\mathrm{per}} := \mathscr{L}^{\chi_{-(1+\epsilon)\beta}}$$

of the canonical linearisation $\mathscr{L}_\beta$ on $\overline{Y_\beta}$ is adapted for the $\widehat{U}_\beta$-action.

4.2.1 When we have condition ($\star$)

Having established that the unipotent radical is graded and we can find a well-adapted linearisation, we now know that some version of the $\hat{U}$ theorems will apply, and we will be able to obtain geometric quotients by the $\hat{U}$ action of suitable open subsets of X. These can then be followed up by quotients by the residual reductive action of the Levi subgroup L_β, yielding categorical quasi-projective quotients of open subsets of X by P_β. Whether these quotients are in fact geometric, or projective, will depend on whether we have 'semistability coincides with stability'[6] for $\hat{U}$ and the action of L_β, the former determining which version of the $\hat{U}$ theorem we are required to use. Let us consider the simplest case, in which we have condition ($\star$).

4.2.1.1 Theorem. *With the notation and assumptions of Proposition 4.2.0.2, suppose additionally that the $\widehat{U}_\beta$-action on $\overline{Y_\beta}$ with respect to the adapted linearisation $\mathscr{L}_\beta^{per}$ satisfies condition ($\star$). Then the following statements hold.*

i) $\overline{Y_\beta}^{\widehat{U}_\beta - s} = Y_\beta - U_\beta Z_\beta$ admits a geometric and projective $\widehat{U}_\beta$-quotient $\overline{Y_\beta}^{\widehat{U}_\beta - s}/\widehat{U}_\beta$.

ii) We have $\overline{Y}_\beta^{P_\beta - ss} = Y_\beta^{ss} - U_\beta Z_\beta^{ss}$, which admits a good and projective P_β-quotient.

iii) $U_\beta Z_\beta \subset Y_\beta$ is a closed subscheme and the map $p_\beta : U_\beta Z_\beta \to Z_\beta$ is a geometric $\widehat{U}_\beta$-quotient. The reductive quotient $Z_\beta /\!\!/ L_\beta$ is a good P_β-quotient of $U_\beta Z_\beta^{ss}$.

Proof. The first statement follows directly from Theorem 3.3.1.6 and Proposition 4.2.0.2. For the second statement, we need to determine the semistable locus for the action of the reductive group $\overline{L_\beta} := L_\beta/\lambda_\beta(\mathbb{G}_m)$ on Z_β with respect to the perturbed linearisation $\mathscr{L}_\beta^{\mathrm{per}}$, as

$$\overline{Y}_\beta^{P_\beta - ss} = \{y \in \overline{Y_\beta}^{\widehat{U}_\beta - s} \mid p_\beta(y) \in Z_\beta^{\overline{L_\beta} - ss}(\mathscr{L}_\beta^{\mathrm{per}})\}.$$

We claim that

$$Z_\beta^{\overline{L_\beta} - ss}(\mathscr{L}_\beta^{\mathrm{per}}) = Z_\beta^{\overline{L_\beta} - ss}(\mathscr{L}_\beta) = Z_\beta^{L_\beta - ss}(\mathscr{L}_\beta) = Z_\beta^{ss}. \tag{4.1}$$

We can calculate these semistable loci for the actions of $\overline{L_\beta}$ and L_β from their torus weights using the Hilbert–Mumford criterion. Let $T \subset L_\beta$ be a maximal torus containing $\lambda_\beta(\mathbb{G}_m)$ and write $\overline{T} := T/\lambda_\beta(\mathbb{G}_m)$. As always, we will use the norm $||-||$ to identify characters and co-characters and consider them both to lie in $\mathfrak{t} := \mathrm{Lie}\, T$.

[6] In the case of $\hat{U}$, this is meant in the sense of ($\star$) and its variants.

If $\alpha_1, \ldots, \alpha_r \in \mathfrak{t} := \operatorname{Lie} T$ denote the T-weights on Z_β with respect to the original linearisation $\mathscr{L}$; then $\alpha_i - \beta$ (resp. $\alpha_i - (1+\epsilon)\beta$) for $1 \leq i \leq r$ denote the T-weights for the canonical linearisation $\mathscr{L}_\beta$ (resp. the perturbed linearisation $\mathscr{L}_\beta^{\mathrm{per}}$). We recall that the 1-PS λ_β corresponds to the ray through β, thus to obtain the $\overline{T}$-weights from the T-weights we project these weights onto the perpendicular hyperplane to the ray through β; this hyperplane corresponds to $\bar{\mathfrak{t}} := \operatorname{Lie} \overline{T}$. Under this projection, both β and $(1+\epsilon)\beta$ are sent to the origin; this explains the first equality in (4.1). The second equality follows, as all points in Z_β have T-weights lying on the perpendicular hyperplane through β, by the weight decription of Z_β, and the final equality follows by definition. Therefore, the claim holds and consequently we obtain that $\overline{Y_\beta}^{P_\beta - ss} = Y_\beta^{ss} - UZ_\beta^{ss}$ and this has a good and projective quotient by Theorem 3.3.1.6. The final statement follows from Theorem 4.2.0.2 in addition to the universal property of categorical quotients. □

4.2.1.2 Remark.

1) As Z_β is unstable for the $\lambda_\beta(\mathbb{G}_m)$-action on $\overline{Y_\beta}$ linearised by $\mathscr{L}_\beta^{\mathrm{per}}$, we see that a semistable point $y \in Y_\beta$ cannot be identified with its limiting flow $p_\beta(y) \in Z_\beta$, which thus avoids the collapsing described in Section 4.2.2.

2) Even if condition $(\star)$ does not hold, then one can construct a quotient of the P_β-action on an open subset of $\overline{Y}_\beta$ by performing a sequence of blow-ups as in 3.3. We will see an example of this in Chapter 5. In particular, we note that the proof above used only the Hilbert-Mumford criterion, so similar statements are valid in any setting where it can be shown that such a criterion determines the stable and semistable loci. This process is described in detail in the appendix.

4.2.2 Toy Example revisited

Let us return to the example of Section 4.1.1.

4.2.2.1 Remark. The caculation of the semistable locus for a parabolic action on

$$X = \operatorname{Sym}^m(\mathbb{P}^1)$$

is done in detail in [Haw15], using the technique of reductive envelopes described above. Hawes considers linearisations of the form

$$\mathscr{L}_{k,r,N} = (\mathscr{O}_X(k) \otimes \mathscr{O}_\chi) \boxtimes \mathscr{O}_{\mathbb{P}^2}(N),$$

i.e. the natural $\mathscr{O}_X(k)$-linearisation twisted by the character χ_r of SL(2) and tensored with $\mathscr{O}_{\mathbb{P}^2}(N)$, for some $N \gg 0$. Our situation in this section corresponds to his in the case

$$m = n - i$$

$$k = 1$$

$$r = (1+\varepsilon)(2i - n)$$

but we will use the $\hat{U}$ theorem to determine the semistable and stable loci.

Recall that $P_\beta = P$ was the same for every level of the stratification. Indeed, our parabolic is just the Borel subgroup of SL(2), i.e.

$$P = \left\{ \begin{pmatrix} a & b \\ 0 & a^{-1} \end{pmatrix} \mid a \in \mathbb{C}^*, b \in \mathbb{C} \right\} \cong \mathbb{G}_a \rtimes \mathbb{G}_m.$$

Hence we can apply the $\hat{U}$ theorem to find quotients:

$$S_\beta /\!\!/ G \cong Y_\beta /\!\!/ P$$

once we have verified its hypotheses. The subgroup P fixes the point $\infty \in \mathbb{P}^1$, so P acts on $\overline{Y}^{ss}_{n-2i}$ as stated above. Indeed, $\overline{Y}^{ss}_{n-2i} \cong \mathbb{P}^{n-i}$ and the action is the natural one, so it extends to an SL(2)-action on each $\mathbb{P}^{n-i}$. Now take

$$V = \begin{pmatrix} 0 & v \\ 0 & 0 \end{pmatrix} \in \mathrm{Lie}(\mathbb{G}_a) \cong \mathbb{C}$$

$$g_\lambda = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix}.$$

The adjoint action of $\mathbb{G}_m$ on the Lie algebra of $\mathbb{G}_a$ is via

$$g_\lambda \cdot X \cdot g_\lambda^{-1} = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix} \begin{pmatrix} 0 & v \\ 0 & 0 \end{pmatrix} \begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix}^{-1} = \begin{pmatrix} 0 & \lambda^2 v \\ 0 & 0 \end{pmatrix}.$$

So the $\mathbb{G}_m$ in the Levy decomposition of P acts on the Lie algebra $\mathrm{Lie}(\mathbb{G}_a)$ of the $\mathbb{G}_a$ with sole weight $2 > 0$.

Now we turn to to weights of the $\mathbb{G}_m$-action from our linearisation. We began with the linearisation $\mathscr{L}$, with respect to which the weights on $\overline{Y}_{n-2i}$ are $\{(2i-n), ...n\}$.

To make the $\hat{U}$ theorem apply, we twist by the rational character associated to

$$(1+\varepsilon)(n-2i),$$

where $\varepsilon > 0$ is a small rational number. All of this is just a specific instance of Proposition 4.2.0.2.

It remain to verify the final hypothesis, to determine whether $(\star)$ holds. The $\mathbb{G}_m$-action is just multiplication on the affine points and fixes the point at infinity, so

$$(\overline{Y}^{ss}_{n-2i})_{\mathbb{G}_m} = \{\text{all points are at either infinity or zero}\}$$

The lowest weight for $\overline{Y}^{ss}_{n-2i}$ corresponds to every point being at 0, so that

$$Z(\overline{Y}^{ss}_{n-2i})_{min} = Z_{n-2i}$$

and it suffices to note that no point of Z_{n-2i} is fixed under the translating action of the $\mathbb{G}_a$, unless $i = n$ and we are considering the case where all points coincide at ∞. This Y_β corresponds to the S_β stratum where all points coincide, and it is clear in this case that the quotient ought to be a point. At all other levels of the stratification, the hypotheses are verified and we can apply the Theorem 3.3.1.6.

The theorem tells us that the (semi)stable locus with respect to our peturbed linearisation is

$$(\overline{Y}_{n-2i})^{ss} = (\overline{Y}_{n-2i})_{min} \setminus (\mathbb{G}_a \cdot Z_{n-2i}).$$

Now we recognise the $\mathbb{G}_a$-sweep of Z_{n-2i} as being the locus where i points are at ∞ and the remaining points are affine and all coincide somewhere. So:

$(\overline{Y}_{n-2i})^{ss} = (\overline{Y}_{n-2i})^{s} = \{\text{Exactly } i \text{ points at } \infty, \text{ the remaining points don't all coincide}\}.$

Since Theorem 3.3.1.6 gives us a geometric quotient of this locus, we see that using non-reductive GIT gives a far better answer than the methods of Section .

4.2.3 Comparison of the two applications

The remainder of this thesis will comprise an attempt to apply the methods discussed above to two important moduli problems, both with well-known reductive GIT formulations: the moduli of semistable coherent sheaves on a projective scheme, and the

moduli of stable curves. Though we have grouped them under the same umbrella, carrying out the above in practice turns out to pose very different challenges.

For sheaves, an intrinsic measure of stability has long been known: the Harder-Narasimhan type, which records the Hilbert polynomials of the subquotients in the Harder-Narasimhan filtration, is the obvious candidate. The task, then, was to show that this type coincided, or could be made to coincide, with the instability type in some GIT formulation of the problem; this was completed by Hoskins and Kirwan in [HK12]. On the other hand, for curves there is a priori no precise idea of what information the instability stratification should contain. The instability of a curve may be said in a certain sense to be caused by its singularities, and so one would expect the instability type to record certain discrete invariants associated to those, and hence for Kempf one parameter subgroup to be one that best 'picks out' the singularities. The search for exactly what discrete data the instability type of a curve should record leads us to the notion of *Rosenlicht-Serre type*, which we discuss further in Chapter 6.

For the second part of the task, the situation is reversed. For sheaves the task of actually performing the quotients is non-trivial, even with the $\hat{U}$ theorem, while for curves (once one has the right set-up) it is relatively easy. This difference comes down to the presence or absence in the unipotent radical $U_\beta \leq P_\beta$ of non-trivial stabilisers of points in Y_β. In order to know which of the $\hat{U}$ theorems we may use, and hence how straightforward the quotient process will be, we must ask whether the points corresponding to generic sheaves, or generic curves, have such stabilisers in the relevant parameter space; this question, in turn may ultimately be interpreted as asking: do the objects those points represent have unipotent automorphisms? It turns out that the answers are 'almost always' (for sheaves: Section 5.3), and 'almost never' (for curves: [Pop78]).

Chapter 5

Moduli Spaces of Sheaves of Fixed Harder-Narasimhan type

Here we present the paradigmatic application of the approach to moduli problems outlined in the previous chapters. The problem in question is that of classifying, i.e. constructing moduli spaces for, coherent sheaves over a projective base scheme. Being the more modern descendant of the moduli of holomorphic vector bundles on a smooth projective curve, classical approaches to this problem have so far proceeded along natural generalisations of this special case. That is, by imposing a stability condition for sheaves analogous to the slope condition on sub-bundles originally used for vector bundles on a curve, and constructing, via reductive GIT, coarse moduli spaces classifying (semi)stable coherent sheaves. The reductive GIT construction we will consider, which gives rise to the instability stratification to which we will apply the results of previous chapters, is the well-known one due to Simpson [Sim94], valid for any projective base scheme. This generalises that of Maruyama [Mar77, Mar78] which works for torsion free sheaves on integral smooth projective schemes.

Simpson constructs the moduli space of semistable sheaves as the quotient by a special linear group of a certain subscheme of a Quot scheme. The instability stratification associated to this GIT set-up was studied by Hoskins and Kirwan in [HK12], who compared it to another stratification of the same space, by Harder-Narsimhan type. Since both stratifications may be said to measure the failure of a coherent sheaf to be semistable, one might hope that they match up somehow, and indeed they do: Hoskins and Kirwan demonstrated that each Harder-Narasimhan stratum is, in an appropriate sense, a closed subscheme of an instability stratum. The tools necessary to quotient these unstable strata, and hence construct moduli spaces of sheaves of

fixed Harder-Narasimhan type, were presented in Section 3.3 and Chapter 4, and here we will use them to carry out this work. However, there is a sense in which this is not a very well behaved example: we have condition ($\star$) only in very special cases. Consequently, we will need the full force of Theorem 3.3.1.12. The natural way to implement this from a moduli point of view is to introduce a refinement of the Harder-Narasimhan stratification, thereby obtaining moduli of semistable sheaves of fixed refined Harder-Narasimhan type. We will complete this task in the case of Harder-Narasimhan length $\ell = 2$, and offer some results towards a construction in the case $\ell = 3$.

5.1 Moduli Spaces of Semistable Sheaves

Let X be a complex projective scheme with a fixed ample line bundle $\mathcal{O}_X(1)$, and suppose one wishes to study the moduli of coherent sheaves on X up to isomorphism. For a coherent sheaf $\mathcal{F}$ on X, we let $P(\mathcal{F}) \in \mathbb{Q}[t]$ denote the Hilbert polynomial of $\mathcal{F}$ with respect to $\mathcal{O}_X(1)$. As the Hilbert polynomial of such a sheaf is locally constant in flat families, it is natural to fix this discrete invariant. Hence, we fix $P \in \mathbb{Q}[t]$, and from now on, unless specified otherwise, all coherent sheaves will have this Hilbert polynomial.

5.1.1 Simpson's construction

The starting point of our work is Simpson's construction of the classical moduli space, which classifies semistable coherent sheaves with Hilbert polynomial P over X, given in [Sim94]: this is the reductive GIT set-up whose instability-strata we will quotient using Non-Reductive GIT, as outlined in Chapter 4 above. We will give a brief summary of Simpson's construction here, beginning by giving the stability condition. First, define a partial ordering $\preccurlyeq$ on $\mathbb{Q}[t]$ by

$$P \preccurlyeq Q \iff \frac{P(n)}{P(m)} \leq \frac{Q(n)}{Q(m)} \qquad \text{for} \quad m \gg n \gg 0,$$

and similarly define a strict partial order $\prec$ by replacing $\leq$ with $<$.

5.1.1.1 Definition. We say $\mathcal{F}$ is semistable if for all non-zero subsheaves $\mathcal{F}' \subset \mathcal{F}$, we have $P(\mathcal{F}') \preccurlyeq P(\mathcal{F})$.

Every semistable sheaf has a Jordan-Hölder filtration: that is, a filtration

$$0 \subset \mathcal{F}_1 \subset \cdots \subset \ldots \mathcal{F}_r = \mathcal{F}$$

such that every successive subquotient $\mathcal{F}_i/\mathcal{F}_{i-1}$ is stable and, placing $\mathcal{F}$ and $\mathcal{F}_i/\mathcal{F}_{i-1}$ on either side of the above inequality, we have equality. This filtration, unlike the Harder-Narasimhan filtration we will consider later, is not canonical. However, its associated graded sheaf is determined up to isomorphism by $\mathcal{F}$.

5.1.1.2 Definition. We say two semistable sheaves are *S-equivalent* if the associated graded sheaves from their Jordan-Hölder filtrations are isomorphic.

5.1.1.3 Remark. The moduli-theoretic motivation for this definition is as follows. If a semistable sheaf $\mathcal{F}$ can be expressed as an extension

$$0 \to \mathcal{F}_1 \to \mathcal{F} \to \mathcal{F}_2 \to 0$$

of two other sheaves, $\mathcal{F}_1$, $\mathcal{F}_2$ also semistable, then once can construct a family $\mathscr{F}$ (parameterised, for example by a suitable line in $\mathrm{Ext}^1(\mathcal{F}_2, \mathcal{F}_1)$) of sheaves over $\mathbb{A}^1$, such that for $t \neq 0$ we have $\mathscr{F}_t \cong \mathcal{F}$, and for $t = 0$ we have $\mathscr{F}_0 \cong \mathcal{F}_1 \oplus \mathcal{F}_2$. This is an example of so-called jump phenomena, and shows us that any separated moduli space including all semistable sheaves must identify $\mathcal{F}$ and $\mathcal{F}_1 \oplus \mathcal{F}_2$. In particular, it shows that any such moduli space must identify all sheaves in the same S-equivalence class.

5.1.1.4 Remark. The definition of semistability above is due to Rudakov [Rud97], and is equivalent to the notion of Gieseker-semistability used by Simpson. In particular, to be semistable it is necessary that a sheaf be pure; that is, the dimensions of the supports of all nonzero subsheaves must be equal. Hence, if X is a smooth projective curve, the above is also equivalent to the notion of slope semistability for vector bundles on a curve. Our reason for using this formulation of semistability is that it allows us to assign Harder-Narasimhan filtrations to any coherent sheaf, not just the pure ones that would be allowed under Gieseker's definition.

Next, we must find the right parameter space and group action. Here we run up against the first obstacle to posing moduli of coherent sheaves as a single GIT problem: that of boundedness, that is, putting all objects of interest into a single family over a scheme of finite type. There is no such family of all coherent sheaves (of

fixed Hilbert polynomial), even in very simple cases.[1] Later we will get around this problem essentially by focusing, at any given time, on at most finitely many Harder-Narasimhan types. A special case of this, found in the construction of Simpson, is to look only at semistable sheaves (i.e. those whose Harder-Narasimhan type is just given by P). For these sheaves, one can find a parameter scheme essentially by reducing the problem to linear algebra, in the following way. By the Le Potier estimates ([Sim94] Theorem 1.1), the family of semistable sheaves over X with Hilbert polynomial P is bounded. Hence, for $n \gg 0$, all such sheaves are n-regular in the sense of Castelnuovo-Mumford ([Mum66] Lecture 14). Recall that, defining

$$\mathcal{F}(n) \coloneqq \mathcal{F} \otimes \mathscr{O}_X(n),$$

a coherent sheaf $\mathcal{F}$ on X is said to be n-regular if we have

$$H^i(X, \mathcal{F}(m-i)) = 0 \text{ for all } i > 0.$$

As observed by Mumford (*loc. cit.*), n-regularity implies that

(i) The sheaf $\mathcal{F}(n)$ is generated by global sections, so that the evaluation map

$$H^0(\mathcal{F}(n)) \otimes \mathcal{O}_X(-n) \to \mathcal{F}$$

is surjective, by right exactness of tensor product;

(ii) All the higher cohomology groups of $\mathcal{F}(n)$, $H^i(\mathcal{F}(n))$ for $i > 0$, vanish, and so $\dim H^0(\mathcal{F}(n)) = P(\mathcal{F}, n)$.

This means that, if we fix a complex vector space V of dimension $P(n)$, and are prepared to make a choice of isomorphism $H^0(\mathcal{F}(n)) \simeq V$, we can instantiate any n-regular sheaf $\mathcal{F}$ with Hilbert polynomial P as the image of a quotient map

$$q : V \otimes \mathcal{O}_X(-n) \to \mathcal{F}.$$

Hence we can associate to any such sheaf a closed point in the Quot scheme

$$\mathrm{Quot}_n \coloneqq \mathrm{Quot}(V \otimes \mathcal{O}_X(-n), P).$$

[1] This can be seen by thinking about the vector bundles $\mathscr{O}(a) \oplus \mathscr{O}(-a)$ on $\mathbb{P}^1$: all of these are rank two and degree zero, but clearly there is no single n that will make them all n-regular.

This scheme, first constructed by Grothendieck[2], is a generalisation of the Grassmannian, parameterising quotient sheaves of $V \otimes \mathcal{O}_X(-n)$ with Hilbert polynomial P. By a quotient sheaf we mean nothing more than an equivalence class of surjections $q : V \otimes \mathcal{O}_X(-n) \to \mathcal{F}$. Two such maps $q_i : V \otimes \mathcal{O}_X(-n) \to \mathcal{F}_i$, $i = 1, 2$, are regarded as equivalent, and hence correspond to the same point in the Quot scheme, if they have the same kernel; or, equivalently, if there exists an isomorphism of sheaves $\varphi : \mathcal{F}_1 \to \mathcal{F}_2$ making the below diagram commute:

$$\begin{array}{ccc} V \otimes \mathcal{O}_X(-n) & = & V \otimes \mathcal{O}_X(-n) \\ \downarrow{\scriptstyle q_1} & & \downarrow{\scriptstyle q_2} \\ \mathcal{F}_1 & \xrightarrow{\quad\varphi\quad} & \mathcal{F}_2. \end{array} \tag{5.1}$$

The points associated to n-regular sheaves will all lie in the open subscheme

$$Q = Q_n \subset \mathrm{Quot}_n$$

defined to be that locus of quotient $q : V \otimes \mathcal{O}_X(-n) \to \mathcal{F}$ such that $H^0(q(n))$ is an isomorphism.

The point of Q associated to a sheaf is not canonical: it depends on the choice of isomorphism $H^0(\mathcal{F}(n)) \simeq V$ above. In other words, a point of the Quot scheme corresponds to a sheaf $\mathcal{F}$, together with a choice of basis of $H^0(\mathcal{F}(n))$[3]. This gives rise to a natural action of $\mathrm{GL}(V)$ on Quot_n, induced by the action on V, which preserves Q. For

$$[q : V \otimes \mathcal{O}_X(-n) \to \mathcal{F}] \in \mathrm{Quot}_n$$

and $g \in SL(V)$, we define $g \cdot [q]$ to be the quotient map given by the composition

$$V \otimes \mathcal{O}_X(-n) \xrightarrow{g^{-1} \otimes \mathrm{id}} V \otimes \mathcal{O}_X(-n) \xrightarrow{\quad q \quad} \mathcal{F}.$$

We note the following properties of this action on Q, which will be important in what follows.

5.1.1.5 Proposition. *([HL10], 4.3.2)*

1) *Orbits of closed points of Q under the $\mathrm{GL}(V)$-action correspond to isomorphism classes of sheaves.*

[2]See [FGI+05] for a very readable account.

[3]In an attempt to avoid cumbersome language, we will occasionally speak of *the sheaf corresponding to a point on the Quot scheme.* By this we mean the image sheaf of the corresponding map, up to isomorphism.

2) For a point $[q] \in Q$ corresponding to a globally generated sheaf $\mathcal{F}$, there is an isomorphism $\mathrm{Stab}_{\mathrm{GL}(V)}(q) \cong \mathrm{Aut}(\mathcal{F})$.

Proof. The first statement is evident from the discussion above. For the second statement we note that any automorphism $\varphi : \mathcal{F} \to \mathcal{F}$ is determined by the map of vector spaces

$$H^0(\varphi(n)) : H^0(\mathcal{F}(n)) \to H^0(\mathcal{F}(n)).$$

Conversely, the isomorphism $H^0(q(n)) : V \to H^0(\mathcal{F}(n))$ coming from the point $q \in Q$ allows us to view an element of $g \in \mathrm{GL}(V)$ as an isomorphism

$$H^0(q(n)) \circ g \circ H^0(q(n))^{-1} : H^0(\mathcal{F}(n)) \to H^0(\mathcal{F}(n)).$$

By the definition of equivalence of points in the Quot scheme given by diagram (5.1) above, this isomorphism of twisted global sections induces an isomorphism of $\mathcal{F}$ if and only if $g \in \mathrm{Stab}(q)$. □

5.1.1.6 Remark.

1) From the above proposition, it is clear that the centre $\mathcal{Z}(\mathrm{GL}(V))$ of $\mathrm{GL}(V)$ acts trivially, and so in our GIT set-up we will pass to considering an action of $SL(V)$. Up to replacing the linearisation with a tensor power, which we are certainly prepared to do, this is equivalent to considering the action of $\mathrm{GL}(V)/\mathcal{Z}(\mathrm{GL}(V)) = \mathrm{PGL}(V)$. See [HL10] 4.3 for more discussion of this.

2) The above properties may be interpreted stack-theoretically, as saying that we have a presentation of $\mathcal{C}oh_{X,P}$, the stack of coherent sheaves on X with Hilbert Polynomial P, as an infinite increasing union of quotient stacks. For this of course it is necessary that we work with $\mathrm{GL}(V)$, since we must record all automorphisms of the sheaves, including the $\mathbb{C}^*$ common to all.

The final ingredient necessary to turn this set-up into a reductive GIT problem is a linearisation. Take $m > n$ and write $H \coloneqq H^0(\mathscr{O}_X(m-n))$, following [Hos11]. Now, tensoring with $\mathcal{O}_X(m)$ and taking global sections gives a map

$$\mathrm{Quot}_n \to \mathrm{Gr}(V \otimes H, P(m))$$

to the Grassmannian of $P(m)$-dimensional vector space quotients

$$V \otimes H \to W.$$

This map is given by

$$[q : V \otimes \mathscr{O}_X(-n) \to \mathcal{F}] \longmapsto [V \otimes H \twoheadrightarrow H^0(\mathcal{F}(m))].$$

Furthermore, because $\mathrm{Ker}(q)$ will be m-regular for m large enough, if we take $m \gg n \gg 0$, this map is an embedding ([HL10], 2.2.4). Following this up with the usual Plücker embedding of the Grassmannian

$$\mathrm{Gr}(V \otimes H, P(m)) \hookrightarrow \mathbb{P}\Big(\bigwedge^{P(m)} (V \otimes H)^*\Big)$$

we obtain an equivariant embedding of Quot_n, and hence Q, into projective space. Denote by $\overline{Q}$ the closure of Q in this projective space. Since Quot_n is projective ([FGI+05]), we have $\overline{Q} \subset \mathrm{Quot}_n$. Let us denote by $\mathscr{L}_{m,n}$ the very ample line bundle on Quot_n obtained from this embedding.

We are now in the best setting for reductive GIT: we have a projective scheme $\overline{Q}$ carrying an action of a reductive G, and the very ample line bundle corresponding to this embedding has a natural linearisation of the G-action. This is the set-up that Simpson uses to prove the following theorem.

5.1.1.7 Theorem. *[Sim94] Let X be a projective scheme, and fix a Hilbert Polynomial $P \in \mathbb{Q}[t]$. Then*

1) *If $m \gg n \gg 0$, the semistable locus for the GIT problem above is*

   $$\overline{Q}^{ss} = \{[q : V \otimes \mathscr{O}_X(-n) \to \mathcal{F}] \in Q \mid \mathcal{F} \text{ is a semistable sheaf in the sense of 5.1.1.1}\}.$$

   Similarly, the GIT stable locus consists of those points of Q corresponding to stable sheaves.

2) *The GIT quotient $M_{P,ss} \coloneqq \overline{Q} /\!\!/_{\mathscr{L}_{m,n}} G$ is a coarse moduli space for S-equivalence classes of semistable coherent sheaves on X with Hilbert polynomial P.*

This moduli space, the natural successor to the moduli space of semistable vector bundles on a smooth curve, is by now a central object in algebraic geometry, and is very well studied from many different perspectives. From our point of view, however, it is just a starting point. Our goal is to extend the classification of coherent sheaves on X to include the unstable ones as well, so that Simpson's moduli space may be thought of as a special case: that of Harder-Narasimhan length one.

5.2 Stratifications for Sheaves

In Section 5.3 we will use the $\hat{U}$ Theorems to perform quotients of the instability strata for Simpson's GIT construction. Before this, following the strategy laid out in Chapter 4, one first needs an intrinsic sheaf-theoretic interpretation of the instability type for this stratification.

5.2.1 Harder-Narasimhan Stratification

To any sheaf, one may associate a canonical filtration by subsheaves, first introduced by Harder and Narasimhan in [HN75]. A proof of its existence and uniqueness in our setting may be found in [HL10].

5.2.1.1 Definition. Let $\mathcal{F}$ be a sheaf on X.

1) The Harder-Narasimhan (HN) filtration of $\mathcal{F}$ is the unique filtration by subsheaves
$$0 = \mathcal{F}^0 \subsetneq \mathcal{F}^1 \subsetneq \cdots \subsetneq \mathcal{F}^\ell = \mathcal{F}$$
such that the subquotients $\mathcal{F}_i \coloneqq \mathcal{F}^i/\mathcal{F}^{i-1}$ are semistable with strictly decreasing Hilbert polynomials $P(\mathcal{F}_1) \succ P(\mathcal{F}_2) \succ \cdots \succ P(\mathcal{F}_\ell)$.

2) We say that F has HN type $\tau = (P(\mathcal{F}_1), \ldots, P(\mathcal{F}_\ell))$, of length ℓ.

3) We say that a Harder-Narasimhan type $\tau = (P_1, \ldots, P_\ell)$ is coprime if, for all $i = 1, \ldots \ell$, all semistable sheaves on X of Hilbert Polynomial P_i are stable. [4]

4) The associated graded of $\mathcal{F}$ is $\mathrm{gr}^{HN}(\mathcal{F}) \coloneqq \oplus_{i=1}^{\ell} \mathcal{F}_i$.

5) When we speak of the HN type of a point $[q : V \otimes \mathcal{O}_X(-n) \to \mathcal{F}] \in \mathrm{Quot}_n$, we simply mean the HN type of the image sheaf $\mathcal{F}$.

5.2.1.2 Notation. From now on, to prevent needless repetition, given a sheaf $\mathcal{F}$ the terms in the Harder-Narasimhan filtration of $\mathcal{F}$ will always be denoted by $\mathcal{F}^i$, and the subquotients will always be denoted by $\mathcal{F}_i$.

There is a so-called 'HN type function', which, given any family of sheaves $\mathscr{F} \to S$, and a closed point $s \in S$, gives the Harder-Narasimhan type of $\mathscr{F}_s$. This function,

[4] The reason for this choice of term is that coprimality generalises the situation of semistability and stability coinciding for vector bundles on a curve, which happens when the rank and degree, two positive integers, are coprime in the usual sense.

by results of Shatz and Nitsure ([Sha77, Nit11]), is upper semi-continuous, and so, if we define

$$\mathrm{Quot}_n^{\tau} \coloneqq \{[q : V \otimes \mathcal{O}_X(-n) \to \mathcal{F}] \in \mathrm{Quot}_n \mid \mathcal{F} \text{ has HN type } \tau\},$$

one has a HN stratification[5] into locally closed subschemes

$$\mathrm{Quot}_n = \bigsqcup_{\tau} \mathrm{Quot}_n^{\tau}.$$

This in turn gives a stratification

$$Q = \bigsqcup_{\tau} Q_n^{\tau}$$

defined in the same way.

5.2.1.3 Remark. All of the above can be translated into the language of stacks. It was shown in [Hos] that there is a Harder-Narasimhan type stratification on the stack, which may be thought of, in an appropriate sense, as a kind of colimit of the Harder-Narasimhan stratifications of the Quot schemes, just as the stack itself is in some sense their colimit.

5.2.2 Instability stratification for sheaves

Recall the notation of 2.3. In addition to the stratification by Harder-Narasimhan type one also has, according to the general principles laid in 2.3, an instability stratification of the Quot scheme,

$$Q = \bigsqcup_{\beta \in \mathcal{B}_{n,m}} S_{\beta}$$

coming from Simpson's GIT set-up.[6] The natural thing to do, then, is to compare the two. Indeed, since they both measure the failure of a sheaf to be semistable, one might hope that they would coincide. The relationship between these stratifications was studied by Hoskins and Kirwan in [Hos, HK12]. First, in [HK12, §6], it was shown that for a pure HN type[7] τ and for $m \gg n \gg 0$ depending on τ, there is a conjugacy class $\beta_{n,m}(\tau)$ of rational one parameter subgroups of G such that the HN

[5] Sometimes also called the Shatz stratification.

[6] This stratification depends in a non-trivial way upon the parameters m and n, so strictly speaking the first question one should ask is whether there is any meaningful sense in which these stratifications can be said to stabilise. Fortunately, as we shall see, there is.

[7] That is, a HN type of a pure sheaf.

stratum Q^τ is contained (set theoretically) in the instability stratum $S_{\beta_{n,m}}(\tau)$. In [Hos], this result was extended to a scheme theoretic statement for all HN types of coherent sheaves. Theorem 5.2.2.4 summarises the relevant results from these papers. First, we define the GIT index $\beta_{n,m}(\tau)$ corresponding to a HN type τ.

5.2.2.1 Definition. *For a tuple of Hilbert polynomials $\tau = (P_1, \ldots, P_l)$ with*

$$\sum_{i=1}^{l} P_i = P$$

and $(n, m) \in \mathbb{N}^2$, we let $\beta_{n,m}(\tau)$ denote the conjugacy class of the following rational 1-PS of $\mathrm{SL}_{P(n)}$

$$\lambda_{\beta_{n,m}(\tau)}(t) = diag \quad (\underbrace{t^{\beta_1}, \ldots, t^{\beta_1}}_{P_1(n)}, \quad \ldots \quad \underbrace{t^{\beta_l}, \ldots, t^{\beta_l}}_{P_l(n)}) \qquad where \quad \beta_i := \frac{P(m)}{P(n)} - \frac{P_i(m)}{P_i(n)}.$$

5.2.2.2 Remark. The origin of this one-parameter subgroup is in fact not so mysterious. As is always the case, a one-parameter subgroup $\lambda : \mathbb{G}_m \to \mathrm{SL}(V)$ is the same thing as the data of its weight filtration

$$\{0\} \subseteq \cdots \subseteq V^i \subseteq V^{i-1} \subseteq \cdots \subseteq V$$

where V^i denotes the span of those vectors on which λ acts with weight at least k (so that only finitely many of the V^i/V^{i+1} are non-zero). And given a point

$$[q : V \otimes \mathcal{O}_X(-n) \to \mathcal{F}] \in \mathrm{Quot}_n,$$

such a weight filtration induces a filtration of $\mathcal{F}$, by defining

$$\mathcal{F}^i := q(V^i \otimes \mathscr{O}_X(-n)).$$

One can then guess the $\beta_{n,m}(\tau)$ associated to $\mathcal{F}$, by simply maximising (with Lagrange multipliers) the normalised Hilbert-Mumford function over all weight filtrations that induce the Harder-Narasimhan filtration of $\mathcal{F}$, and then show that this guess is correct.

5.2.2.3 Definition. Let $P_\tau := P(\lambda_{\beta_{n,m}(\tau)}) = U_\tau \rtimes L_\tau < G = \mathrm{SL}(V)$ be the parabolic subgroup determined by $\lambda_{\beta_{n,m}(\tau)}$. Thus we have, with blocks taken with respect to

the filtration of V induced by $\lambda_{\beta_{n,m}}(\tau)$,

$$P_\tau = \left\{ \begin{pmatrix} A_{11} & A_{12} & \cdots & \cdots & A_{1l} \\ 0 & A_{22} & A_{23} & \cdots & A_{2l} \\ \vdots & \ddots & \ddots & & \vdots \\ \vdots & & 0 & \ddots & A_{l-1\,l} \\ 0 & \cdots & \cdots & 0 & A_{ll} \end{pmatrix} \in G : A_{ij} \in \mathrm{Mat}_{P_i(n)\times P_j(n)} \right\}.$$

The unipotent radical, $U_\tau \triangleleft P_\tau$ is the groups of matrices that are strictly upper triangular with respect to this filtration

$$U_\tau = \{A \in P_\tau : A_{ii} = I_{P_i(n)} \text{ for } 1 \leq i \leq l\}.$$

5.2.2.4 Theorem ([Hos, HK12])**.** *Let $\tau = (P_1, \ldots, P_l)$ be a HN type for sheaves on X with Hilbert polynomial P; then, for $m \gg\gg n \gg 0$, the following statements hold for $\beta := [\beta_{n,m}(\tau)]$.*

i) All sheaves over X with HN type τ are n-regular.

ii) If S_τ denotes the open subscheme of Q^τ consisting of quotients

$$[q : V \otimes \mathcal{O}_X(-n) \to \mathcal{F}]$$

such that $H^0(q(n))$ is an isomorphism, then the G-orbits of closed points in S_τ are in bijection with the isomorphism classes of sheaves with HN type τ.

iii) $\beta \in \mathcal{B}$ is an instability index for the G-action on Q with respect to $\mathscr{L}_{n,m}$ and S_τ is a closed subscheme of the corresponding instability stratum S_β.

iv) There is a closed L_β-invariant subscheme $Z_\tau^{ss} \subset Z_\beta^{ss}$ such that if $Y_\tau^{ss} := p_\beta^{-1}(Z_\tau^{ss})$; then

$$S_\tau = G \cdot Y_\tau^{ss} \cong G \times^{P_\tau} Y_\tau^{ss}.$$

Moreover, by [Hos, HK12], we have an isomorphism

$$Z_\tau^{ss} \cong Q_n^{ss}(P_1) \times \cdots \times Q_n^{ss}(P_l), \tag{5.2}$$

where $Q_n^{ss}(P_i)$ is the open subscheme of

$$Q(P_i) := \mathrm{Quot}(\mathbb{C}^{P_i(m)} \otimes \mathcal{O}_X(-n), P_i)$$

consisting of quotient sheaves $[q_i : V \otimes \mathcal{O}_X(-n) \to \mathcal{F}_i]$ such that $\mathcal{F}_i$ is semistable and $H^0(q_i(n))$ is an isomorphism. In fact, if $R_n(P_i)$ denotes the closure of $Q_n^{ss}(P_i)$ in $Q(P_i)$, then Z_τ^{ss} is the open subscheme of

$$Z_\tau \cong R_n(P_1) \times \cdots \times R_n(P_l) \tag{5.3}$$

equal to the GIT semistable set for the action of the Levi subgroup $L_\tau < P_\tau$ with respect to the canonical linearisation $\mathscr{L}_\tau := \mathscr{L}_{\beta_{m,n}(\tau)}$. The closed points of Y_τ^{ss} are exactly those quotients

$$q : V \otimes \mathcal{O}_X(-n) \to \mathcal{F}$$

such that $H^0(q(n))$ is an isomorphism, $\mathcal{F}$ has HN type τ and the weight filtration of the one parameter subgroup $\lambda_{\beta_{n,m}(\tau)}$ induces the HN filtration of $\mathcal{F}$.

5.2.2.5 Remark. As stated in Section 2.3, the retraction $p_\tau : Y_\tau^{ss} \to Z_\tau^{ss}$ sends a quotient sheaf $[q : V \otimes \mathcal{O}_X(-n) \to \mathcal{F}]$ to the associated graded sheaf for its HN filtration. Thus, if we are to obtain reasonable moduli spaces for sheaves of fixed Harder-Narasimhan type, and prevent sheaves from being identified with their associated graded sheaves, we must make the associated graded unstable. Otherwise, as in [Hos11], the resultant moduli spaces will have too small a dimension: they will parameterise only the graded sheaves, and will be isomorphic to a product of Simpson's moduli spaces. Hence, we will make use of the $\hat{U}$ Theorems, as stated in Section 3.3, according to the strategy described in Section 4.2.

5.3 Moduli of Unstable Sheaves

Having recapped the necessary background, we now turn our attention to the construction of moduli spaces of coherent sheaves of fixed Harder-Narasimhan type. By Theorem 5.2.2.4 we can do this, at least in principle, by quotienting, by quotienting (subschemes of) the unstable strata arising from Simpson's construction for moduli of semistable sheaves. As explained above, the condition $(\star)$ fails very badly here. Not only do we usually have positive-dimensional stabilisers in the unipotent radical, but these dimensions can also vary along Brill-Noether-like loci in the Quot scheme. Consequently, we will need the most general of the $\hat{U}$ Theorems 3.3.1.12, whose associated blow up process is described at Section 3.3.2, which gives us a way to refine the instability stratification in such a way that the refined strata have categorical quotients by the P_τ action. In the case of Harder-Narasimhan length 2 we will show

that this refined stratification, like the Harder-Narasimhan stratification of which it is a refinement, admits an intrinsic sheaf-theoretic description, which we give below, leading to moduli spaces of sheaves of refined Harder-Narasimhan type, of length 2. We will also give a few results towards the length 3 case, to give an idea of how these ideas might be generalised to cover $\ell > 2$.

5.3.1 Stabilisers in U_τ

From now on, let us fix a Harder-Narasimhan type τ. In order to quotient a Harder-Narasimhan stratum using Theorem 3.3.1.12, we need to understand the stabiliser groups for the subgroups in the derived series of U_τ. In our notation, the derived series of U_τ is

$$U_\tau^{(l)} = \{I_{P(n)}\} < U_\tau^{(l-1)} < \cdots < U_\tau^{(k)} < \cdots < U_\tau^{(1)} = U_\tau,$$

$$U_\tau^{(k)} := \{A \in U_\tau : A_{ij} = 0 \text{ if } 0 < j - i < k\}.$$

Thus U_τ is abelian for HN types of length $l = 2$ and non-abelian for $l > 2$.

Before we can interpret these stabilisers intrinsically, we need a few definitions. They are stated for sheaves equipped with a filtration, but we will mainly use them in the case where this is the Harder-Narasimhan filtration.

5.3.1.1 Definition. Let $\mathcal{E}$ and $\mathcal{F}$ be sheaves on X, carrying length ℓ filtrations

$$0 \subset \mathcal{E}^1 \subset \cdots \subset \mathcal{E}^\ell = \mathcal{E}$$

$$0 \subset \mathcal{F}^1 \subset \cdots \subset \mathcal{F}^\ell = \mathcal{F}$$

with successive subquotients

$$\mathcal{E}_i := {\mathcal{E}^i}/{\mathcal{E}^{i-1}}$$

$$\mathcal{F}_i := {\mathcal{F}^i}/{\mathcal{F}^{i-1}}.$$

For each $k \geq 0$, we define the sheaf of k-shifted filtered homomorphisms from $\mathcal{E}$ to $\mathcal{F}$, to be that sheaf whose sections over an open $U \subset X$ are given by

$$\mathcal{H}om_{-k}(U) := \{\varphi \in \mathcal{H}om(\mathcal{E}, \mathcal{F}) \mid \varphi(\mathcal{E}^i) \subset \mathcal{F}^{i-k}\}.$$

We refer to the global sections of this sheaf as the group of k-shifted homomorphisms from $\mathcal{E}$ to $\mathcal{F}$, denoted $Hom_{-k}(\mathcal{E}, \mathcal{F})$. Similarly, in the case $\mathcal{E} = \mathcal{F}$, we define

$\mathrm{End}_{-k}(\mathcal{F})$, the k-shifted endomorphism sheaf of $\mathcal{F}$, its global sections, the k-shifted endomorphism group of $\mathcal{F}$, denoted $\mathrm{End}_{-k}(\mathcal{F})$; and similarly for automorphisms.

5.3.1.2 Remark. By setting $k = 0$ in the above, we recover, the sheaf of homomorphisms from $\mathcal{E}$ to $\mathcal{F}$ that respect their filtrations, and, in the case $\mathcal{F} = \mathcal{F}$, the sheaf of filtration-preserving endomorphisms.

5.3.1.3 Proposition. *Let $\mathcal{E}$ and $\mathcal{F}$ be sheaves over X equipped with length l filtrations as above. Then the following hold.*

i) For $k \geq l$, we have $\mathcal{H}om_{-k}(\mathcal{E}, \mathcal{F}) = 0$.

ii) For $k = \ell - 1$, we have $\mathcal{H}om_{-(\ell-1)}(\mathcal{E}, \mathcal{F}) \cong \mathcal{H}om(\mathcal{E}_l, \mathcal{F}_1)$.

iii) For $0 \leq k \leq l - 2$, we have an exact sequences of sheaves over X

$$0 \to \mathcal{H}om_{-(k+1)}(\mathcal{E}, \mathcal{F}) \to \mathcal{H}om_{-k}(\mathcal{E}, \mathcal{F}) \to \bigoplus_{j=i+k} \mathcal{H}om(\mathcal{E}_j, \mathcal{F}_i). \tag{5.4}$$

iv) If the filtrations of the sheaves $\mathcal{E}$ and $\mathcal{F}$ are both completely split, so that $\mathcal{E} \cong \oplus_i \mathcal{E}_i$ and $\mathcal{F} \cong \oplus_i \mathcal{F}_i$, then all of the above sequences (5.4) are exact on the right, and split. Thus

$$\mathcal{H}om_{-k}(\mathcal{E}, \mathcal{F}) = \bigoplus_{j \geq i+k} \mathcal{H}om(\mathcal{E}_j, \mathcal{F}_i).$$

v) If $\mathcal{E}$ and $\mathcal{F}$ are locally free, whether or not their filtrations split, the above sequences (5.4) are also exact on the right, giving short exact sequences.

Proof.

i) This statement is clear.

ii) Observe that over any open $U \subset X$, an element $\varphi \in \mathcal{H}om_{-(l-1)}(\mathcal{E}, \mathcal{F})(U)$ is the same thing as a map $\mathcal{E}^\ell(U) \to \mathcal{F}_1(U)$ with kernel containing $\mathcal{E}^{\ell-1}(U)$, which then induces a map

$$\mathcal{E}_\ell(U) = {}^{\mathcal{E}^\ell}\!/_{\mathcal{E}^{\ell-1}}(U) \to \mathcal{F}_1(U).$$

Conversely any map $\mathcal{E}_\ell(U) \to \mathcal{F}_1(U)$ gives rise to an element of $\mathcal{H}om_{-(l-1)}(\mathcal{E}, \mathcal{F})(U)$.

iii) In the same way, for any $0 \leq j \leq \ell$, an element of $\mathcal{H}om_{-k}(\mathcal{E}, \mathcal{F})(U)$ may be restricted to a homomorphism

$$\varphi : \mathcal{E}^j(U) \to \mathcal{F}^{j-k}(U),$$

and since $\varphi(\mathcal{E}^{j-1}(U)) \subset \mathcal{F}^{j-k-1}(U)$, we see that φ gives rise to a map

$$\mathcal{E}_j(U) = {\mathcal{E}^j}/{\mathcal{E}^{j-1}}(U) \to {\mathcal{F}^{j-k}}/{\mathcal{F}^{j-k-1}}(U) = \mathcal{F}_{j-k}(U),$$

which will be zero if and only if $\varphi(\mathcal{E}^j(U)) \subset \mathcal{F}^{j-k-1}(U)$. Combining all of these maps together, we have defined a map of sheaves

$$\mathcal{H}om_{-k}(\mathcal{E}, \mathcal{F}) \to \bigoplus_{j=i+k} \mathcal{H}om(\mathcal{E}_j, \mathcal{F}_i),$$

with kernel $\mathcal{H}om_{-(k+1)}(\mathcal{E}, \mathcal{F})$.

iv) This follows easily from definition 5.3.1.1.

v) Since $\mathcal{E}$ and $\mathcal{F}$ are locally free, and it suffices to prove that the sequence is exact on stalks, this follows from the usual equivalence of categories with vector bundles.

□

We can now express the relevant stabiliser groups in U_τ, in terms of filtered endomorphism groups.

5.3.1.4 Proposition.

1) Isomorphism classes of filtered sheaves of Harder-Narasimhan type τ correspond to orbits of Y_τ^{ss} under the action of P_τ.

2) Let $[q : V \otimes \mathcal{O}_X(-n) \to \mathcal{F}] \in Y_\tau^{ss}$. Consider $\mathcal{F}$ as a filtered sheaf, equipped with its Harder-Narasimhan filtration. Then:

i) $\mathrm{Stab}_{P_\tau}(q) \cong \mathrm{Aut}(\mathcal{F})$

ii) $\mathrm{Lie}\,\mathrm{Stab}_{U_\tau}(q) \cong \mathrm{End}_{-1}(\mathcal{F})$

iii) For $1 \leq j \leq \ell$, we have $\mathrm{Stab}_{U_\tau^{(j)}}(q) \cong \mathrm{End}_{-j}(\mathcal{F})$.

Proof. For the first statement, we begin by noting that any automorphism of a sheaf must preserve its Harder-Narasimhan filtration. The rest of the proposition is a straightforward consequence of the analogue in Simpson's construction, as expressed by Proposition 5.1.1.5. One simply needs to follow what happens under the isomorphism between the stabiliser in the Quot scheme and the automorphism group of the sheaf, and compare the filtrations of both sides: induced, on one side, by the weight filtration of $\lambda_{\beta_{m,n}(\tau)}$ and, on the other, by the Harder-Narasimhan filtration of $\mathcal{F}$. □

5.3.2 The case $\ell = 2$

Let us now fix, from now until Section 5.3.3, a HN type $\tau = (P_1, P_2)$ of length 2. We will study the refined stratification in this case.

5.3.2.1 Definition. Let $\mathcal{F}$ be a sheaf of Harder-Narasimhan length 2. The *Harder-Narasimhan sequence* of $\mathcal{F}$ is the short exact sequence

$$0 \to \mathcal{F}_1 \to \mathcal{F} \to \mathcal{F}_2 \to 0.$$

We say that $\mathcal{F}$ is HN-split if this sequence splits, i.e. if $\mathcal{F} \cong \operatorname{gr}\mathcal{F}$, and that $\mathcal{F}$ is Harder-Narasimhan non-split otherwise.

Analysis of the length 2 case is made significantly easier by the following property, which follows from the results of Section 5.3.1, and fails in higher length.

5.3.2.2 Corollary. *For any sheaf* $[q : V \otimes \mathcal{O}_X(-n) \to \mathcal{F}] \in Y_\tau^{ss}$*, we have*

$$\operatorname{Stab}_{U_\tau}(q) \cong Hom(\mathcal{F}_2, \mathcal{F}_1).$$

Hence, in particular, the dimension of the stabiliser in the unipotent radical is the same for a sheaf and for its graded.

5.3.2.3 Remark. As a consequence of this, in performing the blow-up procedure for the unipotent radical given in [BDHK16], we never have to blow up the whole minimal weight space. This means that we are in case (1) of 3.3.2, which makes determining the semistable locus much easier.

We can now define the moduli functors for the moduli problem in length two.

5.3.2.4 Definition.

1) For a Harder-Narasimhan type τ and a scheme S, a flat family $\mathscr{F}$ of sheaves on $S \times X$ is said to have Harder-Narasimhan type τ if for all points $s \in S$, the fibre $\mathscr{F}_s$ has Harder-Narasimhan type τ. Similarly we say that the family is Harder-Narasimhan non-split if none of the fibres are Harder-Narasimhan split, and the family is Harder-Narasimhan split if all of the fibres are.

2) Define a functor

$$\mathcal{M}_{\tau,d} : \mathrm{Sch}_{\mathbb{C}} \to \mathrm{Set}^{op}$$

to be the sheafification of the presheaf that associates to a scheme S the set of isomorphism classes of families $\mathscr{F}$ of sheaves over S such that

i) The family $\mathscr{F}$ is of Harder-Narasimhan type τ, and Harder-Narasimhan non-split

ii) For all $s \in S$, if $\mathscr{F}_s \cong \mathcal{F}$ then $\dim \mathrm{Hom}(\mathcal{F}_2, \mathcal{F}_1) = d$.

3) We futher define a functor

$$\mathcal{M}_{\tau,d}^{split} : \mathrm{Sch}_{\mathbb{C}} \to \mathrm{Set}^{op}$$

to be the sheafification of the presheaf that associates to a scheme S the set of isomorphism classes of families $\mathscr{F}$ of schemes over S such that

i) The family $\mathcal{F}$ is of Harder-Narasimhan type τ, and is Harder-Narasimhan split.

ii) For all $s \in S$, if $\mathscr{F}_s \cong \mathcal{F}$ then $\dim \mathrm{Hom}(\mathcal{F}_2, \mathcal{F}_1) = d$.

5.3.2.5 Theorem. *Let $(X, \mathscr{O}_X(1))$ be a projective scheme with an ample line bundle. Fix a Harder-Narasimhan type $\tau = (\tau_1, \tau_2)$ of length $\ell = 2$, and fix $d \in \mathbb{N}$. Then:*

i) The functor $\mathcal{M}_{\tau,d}$ is co-represented by a quasi-projective scheme $M_{\tau,d}$. Hence, $M_{\tau,d}$ is a coarse moduli space for coherent sheaves $\mathcal{F}$ on X of type τ, such that $\mathcal{F} \ncong \mathrm{gr}\,\mathcal{F}$ and $\dim \mathrm{Hom}(\mathcal{F}_2, \mathcal{F}_1) = d$.

ii) Similarly, the functor $\mathcal{M}_{\tau,d}^{split}$ is co-represented by a quasi-projective scheme $M_{\tau,d}^{split}$, and we have

$$M_{\tau,d}^{split} \cong \{(\mathcal{F}_1, \mathcal{F}_2) \in M_{P_1,ss} \times M_{P_2,ss} \mid \dim \mathrm{Hom}(\mathcal{F}_2, \mathcal{F}_1) = d\}$$

iii) If τ is a coprime type, as defined in Definition 5.2.1.1, the GIT quotients in question are all geometric, points of these moduli spaces correspond to isomorphism classes.

iv) Furthermore, both moduli spaces $M_{\tau,d}$ and $M_{\tau,d}^{split}$ have projective completions

$$M_{\tau,d} \subset \overline{M}_{\tau,d}$$

$$M_{\tau,d}^{split} \subset \overline{M}_{\tau,d}^{split}$$

constructed via non-reductive GIT, such that there is a map $\overline{\mathrm{gr}} : \overline{M}_{\tau,d} \to \overline{M}_{\tau,d}^{split}$, induced by $p_\tau : Y_\tau^{ss} \to Z_\tau^{ss}$, restricting to a map $\mathrm{gr} : M_{\tau,d} \to M_{\tau,d}^{split}$ which takes a sheaf to its associated graded.

The proof of the theorem is rather lengthy, but the basic idea is this: beginning with a suitable subset

$$Y' \subset \overline{Y_\tau}$$

of the projective closure of Y_τ, one performs the blow-up process described above. This corresponds to refining the instability stratification, and hence to refining the Harder-Narasimhan stratification of the Quot scheme. This results in a space $\pi : \widehat{Y}' \to Y'$ satisfying the conditions of Theorem 3.3.1.10. An open subset of $\widehat{Y}'$ therefore has a projective categorical quotient by the P_τ action, and these quotients will be the projective completions mentioned above. To prove the theorem, we must show that certain loci in these projective completion can be identified with the moduli spaces above, and essentially this will be done via the isomorphism over the complement of the exceptional divisor given by restricting the blow-down map. Before moving on to the proof of the theorem, let us make a few remarks.

5.3.2.6 Remark.

1) It may happen that a particular moduli space mentioned above is in fact empty. For example, $M_{\tau,d}$ will be empty if X is a smooth projective curve and we choose P_1, P_2 and d to be such that $\mathrm{Ext}^1(\mathcal{F}_2, \mathcal{F}_1)$ is forced to vanish for degree reasons. Nevertheless, the above statements still hold in such cases.

2) As observed above, the situation is particularly good if the Harder-Narasimhan type in question is coprime. Similarly to the reductive-semistable case, this is a consequence of the fact that, for these HN types, semistability and stability (in the sense of GIT) coincide for the action of $\overline{L_\beta}$. And, again as in the reductive case [Kir85] if this does not happen one can perform a sequence of blow-ups according to the dimensions of reductive stabilisers, to obtain a space on which this condition is satisfied, yielding partial desingularisations of the projective completions in the theorem above, and hence alternative compactifications of $M_{\tau,d}$ and $N^{dec}_{\tau,d}$.

Proof of Theorem 5.3.2.5. We will proceed according to the process laid out in [BDHK16, BHK17]. First, note that since the unipotent radical is abelian, we do not need to take a series of subgroups in U as in Theorem 3.3.1.10, and can blow up according to the dimensions of stabilisers in U directly. Now, let $d_{\min}$ and $d_{\max}$ be respectively the minimal and maximal values of $\dim \mathrm{Stab}_U(y)$ across $y \in Y_\tau$. By Corollary 5.3.2.2, if $y \in Y_\tau^{ss}$, corresponding to a quotient $q : V \otimes \mathscr{O}(-n) \to \mathcal{F}$, we may identify this

stabiliser with the group $\mathrm{Hom}(\mathcal{F}_2, \mathcal{F}_1)$.

Case: $d = d_{\min}$

Consider the locus

$$(Y_\tau^{ss})^{d_{\min}} := \{y \in Y_\tau^{ss} \mid \dim \mathrm{Stab}_U(y) = d_{\min}\}.$$

By Proposition 5.3.1.4 and Corollary 5.3.2.2 orbits under the P_τ action on this space correspond to isomorphism classes of sheaves $\mathcal{F}$ of Harder-Narasimhan type τ, with $\dim \mathrm{Hom}(\mathcal{F}_2, \mathcal{F}_1) = d_{\min}$. Furthermore, by the results of Section 5.2.2 the sheaves in

$$U(Z_\tau^{ss})^{d_{\min}} = UZ_\tau^{ss} \cap (Y_\tau^{ss})^{d_{\min}}$$

are precisely those sheaves which are HN split. Hence we must exhibit categorical quotients of the loci $U(Z_\tau^{ss})^{d_{\min}}$ and $(Y_\tau^{ss})^{d_{\min}} \setminus U(Z_\tau^{ss})^{d_{\min}}$ by their P_τ actions.

Taking the projective closure $\overline{Y}_\tau$ in the ambient projective space, we first blow up $\overline{Y_\tau}$ along the closure of the subset

$$Y_\tau^{d_{\max}} := \{y \in Y_\tau \mid \dim \mathrm{Stab}_U(y) = d_{max}\}.$$

This results in a space $\overline{Y_\tau}^{(1)}$ with an action of P_τ lifting that on $\overline{Y_\tau}$. Since the centre of the blow up did not contain Z_τ, the minimal weight space for λ_τ acting on $\overline{Y_\tau}^{(1)}$, which we denote by $Z_\tau^{(1)}$, is the proper transform of the minimal weight space Z_τ of $\overline{Y_\tau}$, as in Section 3.3.2 case (1). Let us simplify notation by denoting by $Y_\tau^{(1)}$ the open subset

$$(\overline{Y_\tau}^{(1)})^0_{\min} \subset \overline{Y_\tau}^{(1)}$$

consisting of the points flowing down to $Z_\tau^{(1)}$ under λ_τ as $t \to 0$. Finally, we linearise the P_τ action on $\overline{Y_\tau}^{(1)}$ by replacing the linearisation $\mathscr{L}_\tau^{per}$ with its perturbation by a small multiple of the exceptional divisor of the blow up. Because all of these blow ups may be performed in the ambient space, we may assume that at each stage of the blow up process the exceptional divisor is just the projectivised normal bundle to the centre of the blow up, and the corresponding weights are just certain small perturbations of the weights for Y, in the way described in Remark 3.3.2.1.

This process strictly reduces the dimensions of the stabilisers in U over the centre of

the blow-up (see [BDHK16]), thus guaranteeing that the dimensions

$$\dim \mathrm{Stab}_U(y),\ y \in Y_\tau^{(1)}$$

are all strictly less than $d_{\max}$. We now blow up $\overline{Y_\tau}^{(1)}$ along the closure of the locus of points where the stabiliser above is maximal, and we repeat this process until it terminates with a space $\pi : \widehat{Y}_\tau \to \overline{Y_\tau}$. This space carries a P_τ action lifting that on $\overline{Y_\tau}$, and this action has a well-adapted linearisation $\widehat{\mathscr{L}_\tau^{per}}$ obtained as described, by perturbing at each stage of the blow-up process with a small multiple of the exceptional divisor. Let us denote the minimal λ_τ weight space in $\widehat{Y}_\tau$ by $\widehat{Z}_\tau$. This has basin of attraction $(\widehat{Y}_\tau)^0_{\min}$. Then we have arrived at a situation where

$$\dim \mathrm{Stab}_U(y) = d_{\min}\ \forall\ y \in (\widehat{Y}_\tau)^0_{\min}.$$

In other words, the conditions of Theorem 3.3.1.10 are satisfied, so there is a projective categorical quotient of the locus $\widehat{Y}_\tau^{P_\tau - ss}$, given by composing a geometric quotient by $\hat{U}$ with a categorical quotient by the residual action of the reductive group L_τ. If this quotient is non-empty, we define

$$\overline{M}_{\tau, d_{\min}} \coloneqq \widehat{Y}_\tau /\!\!/_{\widehat{\mathscr{L}_\tau^{per}}} P_\tau.$$

In addition, there is an open subset $U\widehat{Z}_{\min}$ that has a categorical quotient by the P_τ action. Indeed, the geometric quotient of $U\widehat{Z}_{\min}$ by $\hat{U}$ is of course $\widehat{Z}_{\min}$, and the open subset $\widehat{Z}_{\min}^{L_\tau - ss}$ has a categorical quotient given by the usual GIT quotient with respect to the canonical linearisation, as described in Section 4.1. Hence we define

$$\overline{M}^{split}_{\tau, d_{\min}} \coloneqq \widehat{Z}_{\min} /\!\!/_{\mathscr{L}_\tau} L_\tau.$$

It remains to identify suitable subsets of these projective varieties with the moduli spaces above.

Let E be the exceptional locus of the map $\pi : \widehat{Y}_\tau \to \overline{Y_\tau}$; then π gives an identification

$$\widehat{Y}_\tau \setminus E \cong \overline{Y_\tau} \setminus \pi(E).$$

The locus $(Y_\tau^{ss})^{d_{\min}}$ is untouched by the blow-up, and hence this identification restricts to an identification of $(Y_\tau^{ss})^{d_{\min}}$ with a subset of $\widehat{Y}_\tau \setminus E$. And, since at every stage $\pi_i : \overline{Y_\tau}^{(i)} \to \overline{Y_\tau}^{(i)}$ of the blow-up process the centre of the blow up did not contain the minimal λ_τ weight space of $\overline{Y_\tau}^{(i)}$, the blow down map $\pi : \widehat{Y}_\tau \to \overline{Y_\tau}$ also restricts to

an identification $\widehat{Z}_\tau \setminus E \cong (Z_\tau^{ss})^{d_{\min}}$ with a subset of $\widehat{Z}_\tau \setminus E$. Futhermore, all of these identifications are P_τ-equivariant, since π is P_τ equivariant.

Proceeding analogously to Section 4.2, let $\widehat{Z}_\tau^{ss} \subset \widehat{Z}_\tau$ denote the semistable locus for the action of L_τ on $\widehat{Z}_\tau$ with respect to the canonical linearisation $\mathscr{L}_\tau$ for $\widehat{Y}_\tau$. Then let $\widehat{Y}_\tau^{ss}$ be the inverse image $\widehat{p_\tau}^{-1}(\widehat{Z}_\tau^{ss})$ under the flow

$$\widehat{p_\tau} : (\widehat{Y}_\tau)^0_{\min} \to \widehat{Z}_\tau$$

given by taking $t \to 0$ in the λ_τ action. By the proof[8] of Theorem 4.2.1.1 (ii) for the blown up space $\widehat{Y}_\tau$, one has

$$\widehat{Y}_\tau^{P_\tau - ss} = \widehat{Y}_\tau^{ss} \setminus U_\tau \widehat{Z}_\tau^{ss}.$$

Then, by the Hilbert-Mumford criterion of Section 3.3, we may identify $\widehat{Z}_\tau^{ss} \setminus E$ with $(Z_\tau^{ss})^{d_{\min}}$, via π. It follows that we may identify $(Y_\tau^{ss})^{d_{\min}}$ with $\widehat{Y}_\tau^{ss} \setminus E$. This tells us that we can restrict the quotient map

$$\widehat{Y}_\tau^{P_\tau - ss} \to \widehat{Y} /\!\!/ P_\tau$$

to get quotients by the P_τ actions :

$$M^{split}_{\tau, d_{\min}} := U(Z_\tau^{ss})^{d_{\min}} /\!\!/_{\widehat{\mathscr{L}}_\tau} P_\tau$$

$$M_{\tau, d_{\min}} := [(Y_\tau^{ss})^{d_{\min}} \setminus U(Z_\tau^{ss})^{d_{\min}}] /\!\!/_{\widehat{\mathscr{L}}_\tau^{per}} P_\tau.$$

Moreover, the closure of every orbit in these loci is disjoint from the exceptional locus of the blow up, since it is disjoint from the centre of the blow-up downstairs. It follows from standard properties of GIT quotients that these quotients are good in the sense of [New78], and hence these quotients are quasi-projective and categorical. The statements about co-representation of functors then follows from the properties of categorical quotients, as in [HL10] 4.3.

Case: $d > d_{min}$

Supposing inductively that the theorem holds for some $d-1$, we return to our original

[8] Note that the proof applies whenever there is a Hilbert-Mumford characterisation of the stable and semistable loci, as in this case there is.

$\overline{Y}_\tau$. By upper semi-continuity of stabiliser dimension, the subset

$$\overline{Y_\tau^d} := \overline{Y}_\tau \setminus \{y \in \overline{Y_\tau} \mid \dim \mathrm{Stab}_U(y) < d\}$$

is closed in $\overline{Y}_\tau$. We then simply repeat the process described above, replacing $\overline{Y}_\tau$ by $\overline{Y_\tau^d}$ and $d_{\min}$ by d.

□

In particular, as a special case we immediately recover a result obtained in [BPMN] by totally different means. It is interesting to note that the same discrete invariant appears there as when one approaches the problem using non-reductive GIT.

5.3.2.7 Corollary. *[BPMN] Let X be a nonsingular projective curve, and $\tau = (P_1, P_2)$ a pure HN type corresponding to coherent sheaves of rank 2. Then the space $M_{\tau,d}$ above is a quasi-projective moduli space of indecomposable coherent sheaves $\mathcal{F}$ of Harder-Narasimhan type τ with* $\dim \mathrm{End}(\mathcal{F}) = d + 1$.

Proof. It suffices to note the following, which is proved in [BPMN]: for a coherent sheaf of HN rank 2, the non-splitting of the HN sequence

$$0 \to \mathcal{F}_1 \to \mathcal{F} \to \mathcal{F}_2 \to 0$$

is equivalent to indecomposability of $\mathcal{F}$. One direction is clear. For the other, note that if $\mathcal{F} \cong \mathcal{F}_1 \oplus \mathcal{F}_2$, where E_i are necessarily line bundles, we must have without loss of generality that $\deg\mathcal{F}_1 > \deg\mathcal{F}_2$. Hence the sequence

$$0 \to \mathcal{F}_1 \to \mathcal{F} \to \mathcal{F}_2 \to 0,$$

which splits, must be the Harder-Narasimhan sequence of $\mathcal{F}$. □

5.3.2.8 Remark. In higher rank, this result does not hold because it is no longer the case that non-splitting of HN filtration implies indecomposability. For example, we could take $\mathcal{E} = \mathcal{F} \oplus \mathcal{F}$ where $\mathcal{F}$ is any sheaf of HN length two.

5.3.3 Towards $\ell = 3$

Having dealt with the length two case, let us now fix a Harder-Narasimhan type $\tau = (\tau_1, \tau_2, \tau_3)$ of length three. Passing from length two to length three involves a

significant step up in complexity, for two main reasons. First, since the unipotent radical U_τ is no longer abelian, one must consider the derived series

$$U = U^1 > U^2 > \{e\}.$$

To achieve the situation of Theorem 3.3.1.10, one first blows up according to the dimension the stabiliser in U^2, which for points of Y_τ^{ss} we may identify with $\mathrm{Hom}(\mathcal{F}_3, \mathcal{F}_1)$, by Proposition 5.3.1.3. Up to this point, the process is similar to the length two case. Having achieved the necessary condition for U^2, one then must perform a series of blow-ups for the dimensions of stabilisers in U as a whole. The problem is that for length $l > 2$, the analogue of Corollary 5.3.2.2 fails to hold, and in fact this stabiliser is not readily identifiable with any homomorphism group.

5.3.3.1 Example. Now suppose X to be a smooth projective curve, and τ a pure Harder-Narasimhan type of length 3, so that we consider vector bundles on X. Fix a point $q \in Y_\tau^{ss}$ corresponding to a vector bundle $\mathcal{F}$. By Proposition 5.3.1.3 one has an exact sequence

$$0 \to \mathcal{H}om(\mathcal{F}_3, \mathcal{F}_1) \to \mathrm{Stab}_U(q) \to \mathcal{H}om(\mathcal{F}_3, \mathcal{F}_2) \oplus \mathcal{H}om(\mathcal{F}_2, \mathcal{F}_1) \to 0$$

From the associated long exact sequence, we obtain maps

$$\delta_{32} : \mathrm{Hom}(\mathcal{F}_3, \mathcal{F}_2) \to \mathrm{Ext}^1(\mathcal{F}_3, \mathcal{F}_1)$$

$$\delta_{21} : \mathrm{Hom}(\mathcal{F}_2, \mathcal{F}_1) \to \mathrm{Ext}^1(\mathcal{F}_3, \mathcal{F}_1).$$

By the proof of Proposition 5.3.1.3, these maps are nothing but the boundary maps in the long exact sequences one obtains by applying the functors $\mathrm{Hom}(\mathcal{F}_3, -)$, and $\mathrm{Hom}(-, \mathcal{F}_1)$, respectively, to the short exact sequences

$$0 \to \mathcal{F}_1 \to \mathcal{F}^2 \to \mathcal{F}_2 \to 0 \tag{5.5}$$

$$0 \to \mathcal{F}_2 \to \mathcal{F}/\mathcal{F}_1 \to \mathcal{F}_3 \to 0. \tag{5.6}$$

Consequently, for a point in $q \in Y_\tau^{ss}$ corresponding to a bundle $\mathcal{F}$ one has

$$\dim \mathrm{Stab}_U(y) = \sum_{1 \le i < j \le 3} \dim \mathrm{Hom}(\mathcal{F}_j, \mathcal{F}_i) - \mathrm{rank}(\delta_{32}) - \mathrm{rank}(\delta_{21}). \tag{5.7}$$

The ranks of the maps δ_{ji} will naturally depend on the sheaf in question, and will both be zero if and only if the corresponding extensions (5.5) and (5.6) are split. This means that the stabiliser in U_τ of a point $q \in Y_\tau^{ss}$ can in general be of different dimension than that of $p_\tau(q) \in Z_\tau^{ss}$, with the result that at some stage in the blow-up process one expects to have to blow up the whole minimal weight space. This makes the final minimal weight space, and the set of points that flow down to it, more difficult to determine.

Let us return to the case where X is any projective scheme. Recalling the discussion of the blow up process of Section 3.3.2, we will need to interpret what it means for a given $q \in Y_\tau^{ss}$ to have or not have certain weights for the λ_τ action. The associated graded sheaf, being the limit under λ_τ, of course only has one weight, and so we might expect the presence of certain weights for a given q to be related to the question of to what extent the Harder-Narasimhan filtration of the associated sheaf can be said to split. This turns out to be correct, as the following proposition makes precise.

5.3.3.2 Proposition. *Recall the embedding*

$$Y_\tau^{ss} \subset \mathrm{Quot}_n \subset \mathrm{Gr}(V \otimes H, P(m)) \subset \mathbb{P}\Big(\bigwedge^{P(m)} (V \otimes H)^* \Big),$$

and the linearisation of the action of G, discussed in Section 5.1.1. Let $\lambda_\tau = \lambda_{\beta_{m,n}(\tau)} : \mathbb{G}_m \to G$, as defined at Definition 5.2.2.1, be the one parameter subgroup inducing the Harder-Narasimhan filtration of all $y \in Y_\tau^{ss}$, so that

$$\lambda(t) = diag \quad (\underbrace{t^{\beta_1}, \ldots, t^{\beta_1}}_{P_1(n)}, \quad \ldots \quad \underbrace{t^{\beta_l}, \ldots, t^{\beta_\ell}}_{P_\ell(n)}) \qquad where \quad \beta_i := \frac{P(m)}{P(n)} - \frac{P_i(m)}{P_i(n)}.$$

Then the set of λ_τ-weights of the action on $\mathbb{P}\Big(\bigwedge^{P(m)} (V \otimes H)^* \Big)$ is

$$\{x\beta_1 + y\beta_2 + z\beta_3 \mid x + y + z = P(m)\},$$

and we may write all λ_τ-weights of points $y \in Y_\tau^{ss}$ uniquely in barycentric coordinates (x_{21}, x_{32}, x_{31}) as

$$(P_1(m)\beta_1 + P_2(m)\beta_2 + P_3(m)\beta_3) + x_{21}(\beta_2 - \beta_1) + x_{32}(\beta_3 - \beta_2) + x_{31}(\beta_3 - \beta_1),$$

with $x_{21}, x_{32}, x_{31} \geq 0$. With this convention, for a point $q \in Y_\tau^{ss}$ corresponding to a

sheaf $\mathcal{F}$, we have:

1) The short exact sequence

$$0 \to \mathcal{F}_1 \to \mathcal{F}^2 \to \mathcal{F}_2 \to 0$$

splits if and only if nonzero Plücker coordinates of q have weights with $x_{21} = 0$.

2) The short exact sequence

$$0 \to \mathcal{F}_2 \to \mathcal{F}/\mathcal{F}_1 \to \mathcal{F}_3 \to 0$$

splits if and only if nonzero Plücker coordinates of q have weights with $x_{32} = 0$.

Proof. To simplify notation, let us temporarily abuse notation and write P_i for $P_i(m)$. The first statement follows from the fact that each Plücker coordinate is formed by a choice of P elements of our chosen basis of $V \otimes H$, and each such basis element has weight β_i, for some $i = 1, 2, 3$. Next, recall that a quotient q has nonzero Plücker coordinate $v_1 \wedge \cdots \wedge v_P$ if and only if the images of the vectors $v_1, \ldots, v_P$ span $H^0(\mathcal{F}(m))$. Any such a spanning set must be obtainable in the following way: first choose a basis of the image of

$$H^0(q^1(m)) : V^1 \otimes H \to H^0(\mathcal{F}^1(m)),$$

which must consist of P_1 vectors of weight β_1, then extend to a basis for $H^0(q^2(m))$ by choosing P_2 vectors of weight either β_1 or β_2, and then finally extend to a basis of $H^0(q^3(m))$ by choosing a further P_3 vectors of weight either β_1, β_2, or β_3.

The two pairs of statements asserted to be equivalent are really just two ways of saying that the matrix $H^0(q(m))$ with respect to the weight filtration of λ_τ has a particular block form. Since it respects the filtration, we know $H^0(q(m))$ must be given by a $P \times \dim(V \otimes H)$ matrix of the form

$$\left(\begin{array}{c|c|c} A & D & F \\ \hline 0 & B & E \\ \hline 0 & 0 & C \end{array}\right)$$

A given Plücker coordinate then corresponds to a choice of P columns, and the coordinate will be nonzero iff this gives an invertible matrix, i.e. if that subset of the columns spans the image, with the weight of that Plücker coordinate determined by

which side of the two vertical lines those columns lie on. Consider the matrix given by the four blocks in the upper left,

$$\left(\begin{array}{c|c} A & D \\ \hline 0 & B \end{array}\right).$$

Then we see that we must have $D = 0$ if $\mathcal{F}^2 \cong \mathcal{F}_1 \oplus \mathcal{F}_2$, and so in finding a basis we must begin by choosing P_1 columns from before the first vertical line, and then choose P_2 columns from between the two vertical lines, and then the remainder might possibly lie anywhere. Consequently, all weights are of the form given above. Conversely, if $\mathcal{F}^2 \ncong \mathcal{F}_1 \oplus \mathcal{F}_2$ does not split, we must have $D \neq 0$, and then we can choose a nonzero column of $H^0(q(m))$ corresponding to which D has a nonzero column, and get a nonzero minor corresponding to a nonzero Plücker coordinate whose weight is not of the form stipulated. The statement in (2) is proved similarly. □

5.3.4 Remarks on $\ell > 3$

It should be clear from the above that, with the appropriate diligence, one could find results analogous to Proposition 5.3.3 in arbitrary length, and hence one could extend Theorem 5.3.2.5 to arbitrary ℓ, and finish off the moduli classification. The key difficulty is the need to write down an interpretation of the intermediate stabilisers to find out exactly which sheaves appear in which moduli spaces: that is to say, to have a full modular interpretation of the quotient of the complement of the exceptional locus. Another possibility is that one might be able to modify the proof of Theorem 3.3.1.12 slightly, to allow us to work with the full stabiliser in U from the beginning, rather than with the filtration by the U^i.

Chapter 6

Moduli Spaces of Singular Curves

Moduli spaces of curves are sufficiently important that the fact that they do not exist should not prevent us from studying them.
Tim Dokchitser

In this chapter we will consider our second application of non-reductive GIT methods to the construction moduli spaces of unstable objects: that of unstable, i.e. singular, curves.[1]

6.1 Constructions for Stable Curves

Before going on to think about singular curves, let us review some of the main ideas in the GIT constructions of the classical moduli spaces for stable curves. We will be very brief, as an excellent exposition already exists in [Mor09].

6.1.1 GIT construction of $\overline{M}_g$

6.1.1.1 Definition. We say that a curve is stable if its only singularities are nodes, and it has finite automorphism group.

The (coarse) moduli space of stable curves, $\overline{M}_g$, was first constructed by David Mumford using reductive GIT, as a quotient of a Chow variety [Mum77]. That said,

[1]By a curve we mean a variety over $\mathbb{C}$ of dimension one. In some situations we will apply the term more generally to allow non-reduced and reducible behaviour, e.g. when discussing Hilbert schemes of 'curves', but it will be clear what is meant from the context.

a construction via Hilbert schemes was later given by Gieseker in [Gie82], and this is the construction we will follow here. Fixing a degree d and genus g, or equivalently a Hilbert polynomial $P(m) = dm - g + 1$, Gieseker considers the Hilbert scheme of curves $C \subset \mathbb{P}^r$ of genus g and degree d, which we denote

$$\mathcal{H} \coloneqq \mathrm{Hilb}(\mathbb{P}^r, d, g).$$

6.1.1.2 Notation. To make notation easier to read, we denote the Hilbert point of a curve by $[C] \in \mathcal{H}$. We will denote the restriction to C of the line bundle $\mathscr{O}_{\mathbb{P}^r}(1)$ by $\mathscr{O}_C(1)$. That is, we define

$$\mathscr{O}_C(1) \coloneqq i^* \mathscr{O}_{\mathbb{P}^r}(1)$$

where $i : C \hookrightarrow \mathbb{P}^r$ is the inclusion.

The scheme $\mathcal{H}$ plainly characterises isomorphism classes of curves, together with some redundancy, and there is a locus on which that redudancy can be easily described. Consider the so-called n-canonical locus

$$\mathcal{K} = \{\text{smooth curves } C \text{ embedded via } \omega_C^{\otimes n}\}.$$

Since being smooth is an open condition, and the canonical bundle deserves its name, one can show that $\mathcal{K}$ is a locally closed subscheme. One then takes n sufficiently large that all curves of genus g will appear in this locus. This then carries an action of a suitable special linear group $\mathrm{SL}(r+1)$, induced by the natural action on $\mathbb{P}^r = \mathbb{P}(V)$. Isomorphism classes of smooth curves correspond to orbits of $\mathcal{K}$, since the action effectively just changes the basis of the global sections of $\omega_C^{\otimes n}$ that has been used to embed the curve. In [MFK94], Mumford shows that there is a geometric quotient

$$\mathcal{K}/\mathrm{SL}(r+1) = M_g,$$

the moduli space of smooth curves of genus g. In order to obtain a compactification of this moduli space via GIT, one must have an action on a projective variety. Hence, with respect to a parameter $m \gg 0$, one considers the Plücker embedding:

$$\mathcal{K} \subset \mathcal{H} \hookrightarrow \mathrm{Gr}(\mathrm{Sym}^m V^*, P(m)) \hookrightarrow \mathbb{P}\left(\bigwedge^{P(m)} \mathrm{Sym}^m V^*\right).$$

6.1.1.3 Notation. We denote the very ample line bundle on $\mathcal{H}$ corresponding to this embedding by Λ_m.

Under this embedding, a curve $[C] \in \mathcal{H}$, with ideal $I \trianglelefteq \mathbb{C}[X_0, \ldots, X_r]$, is defined as the point in the Grassmannian given by the surjection

$$[\mathrm{Sym}^m V^* \twoheadrightarrow H^0(\mathscr{O}_C(m))] \in \mathrm{Gr}(\mathrm{Sym}^m V^*, P(m)) \tag{6.1}$$

The Plücker coordinates on the Hilbert scheme thus correspond to choices of $P(m)$ monomials, each of degree m, in the variables $X_0, \ldots X_r$. The cooresponding Plücker coordinate will be nonzero for a curve C if and only if that set of monomials, upon restriction to C, gives a basis of $H^0(C, \mathscr{O}_C(m))$.

Hence, if we fix $\lambda : \mathbb{G}_m \to \mathrm{SL}(V)$ diagonalised by the coordinates X_i, acting with weights

$$\lambda(t)X_i = t^{w_i}X_i,$$

the Hilbert-Mumford function has the following form, with the minimum taken over all possible choices of sets $\{M_1, \ldots, M_{P(m)}\}$ of distinct monomials of degree m.

$$\mu([C], \lambda) = \min\left\{ \sum_{i=1}^{P(m)} \mathrm{wt}(M_i) \right\} \tag{6.2}$$

In other words, the value of the Hilbert-Mumford function is the weight of the minimally weighted monomial basis. Note that this expression, whilst essentially just a quadratic in m, is a minimum over a very large set, and hence extremely difficult to compute. This reflects the fact that, unlike for Simpson's GIT set-up for sheaves, it is not usually easy to say what the limit of the Hilbert point of a curve will be under an arbitrary one parameter subgroup. In addition to this, one must, at least in principle, check all one-parameter subgroups in order to prove that a given curve is stable.[2] Thus, instead of explicitly calculating this function in particular cases, Gieseker relies on an ingenious, if somewhat intractable, argument involving finding bounds for it. In this way, they prove:

6.1.1.4 Theorem. *[Gie82] If m and n are sufficiently large, the GIT quotient is*

$$\overline{M}_g = \overline{\mathcal{K}}^s/SL(r+1) = \{\textit{stable curves of genus } g\}/\textit{isom},$$

Hence, $\overline{M}_g$ is a projective variety.

[2]The level of difficulty of this is illustrated by the fact that a direct proof of the GIT stability of nodal curves was only found some forty years later, in [LW15]. In what follows, we will rely heavily on the fact that this function is computable in practice so long as the one parameter subgroup is an automorphism of the curve- that is, it belongs to the stabiliser in the Hilbert scheme.

6.1.2 Baldwin-Swinarski construction of $\overline{M}_{g,n}$

Next we consider a generalisation of Gieseker's result, due to Baldwin and Swinarski. A reasonable way to start thinking about singular curves is that such things are determined by the normalisation, as a curve with marked points, plus some extra data. Exactly what extra data, we will see in Section 6.2.1, and the rest of this chapter will be dedicated to the question of how best to capture that data. In any case, this motivates our interest in the moduli space of stable curves with marked points, $\overline{M}_{g,n}$.

6.1.2.1 Definition. A curve with marked points $(C, p_1, \ldots, p_\ell)$ is said to be stable if C is a stable curve, and the p_i are all distinct smooth points of C.

The GIT construction of $\overline{M}_{g,n}$ proceeds in fundamentally the same way as Gieseker's, although his estimates are not sharp enough, and must be replaced. One considers the parameter scheme

$$\hat{\mathcal{H}} \coloneqq \mathcal{H} \times (\mathbb{P}^r)^\ell$$

where $\mathbb{P}^\ell$ is thought of as its own Hilbert scheme of length 1 subschemes. One then considers a suitable closed incidence subscheme, which requires the points to lie on the curve, in addition to other conditions as in Gieseker's construction. Then, one linearises the diagonal $\mathrm{SL}(r+1)$ action on the $\hat{\mathcal{H}}$ using the very ample line bundle

$$L(m, m_1, \ldots, m_\ell) \coloneqq \Lambda_m \otimes \mathscr{O}_{\mathbb{P}^r}(m_1) \otimes \ldots \mathscr{O}_{\mathbb{P}^r}(m_\ell),$$

which has a natural $\mathrm{SL}(r+1)$ linearisation.

6.1.2.2 Theorem. *[Bal08] For suitable choices of the parameters $(d, m, m_1, \ldots, m_\ell)$, and a suitable locus in $\hat{\mathcal{H}}$, semistability coincides with stability and the GIT quotient realises the moduli space of stable marked curves, $\overline{M}_{g,n}$, as a projective variety.*

6.2 Preliminary Study of Singular Curves

6.2.1 Rosenlicht-Serre Construction

With the necessary background filled in, we can start to ask the question: *what data determines a singular curve?* For our purposes, the correct answer to this question is given by the Rosenlicht-Serre construction. This construction, given in Chapter 4 of

[Ser12], runs as follows. Let C be an algebraic curve and $\pi : \widetilde{C} \to C$ its normalisation. Define an equivalence relation $\sim$ on $S(\widetilde{C})$, the set of points of $\widetilde{C}$, by declaring two points equivalent if they map to the same point in C under π. Then one may recover C as a topological space by endowing $S(C)/\sim$ with the cofinite topology; to recover C as an algebraic curve, however, we also need to find a sheaf of rings.

6.2.1.1 Definition. Consider $\pi_*\mathscr{O}_{\widetilde{C}}/\mathscr{O}_C$ as a sheaf of $\mathscr{O}_C$-modules. Then define the *conductor sheaf*

$$\mathfrak{c} = Ann(\pi_*\mathscr{O}_{\widetilde{C}}/\mathscr{O}_C).$$

This sheaf of ideals on C is supported only at the singular points, and its stalk $\mathfrak{c}_p$ at a given point $p \in C$ is called the *conductor ideal* of the singularity at p.

6.2.1.2 Remark. An equivalent characterisation of the conductor ideal above is that it is the largest sheaf of functions on C that is also a sheaf of ideals on $\widetilde{C}$.

Denote by $\mathfrak{r}_p$ the Jacobson radical of the semi-local ring

$$\widetilde{\mathscr{O}}_p \coloneqq (\pi_*\mathscr{O}_{\widetilde{C}})_p = \bigcap_{q \in \pi^{-1}(p)} (\mathscr{O}_{\widetilde{C}})_q.$$

Then, for all singular points $p \in C$ we have inclusions

$$\mathfrak{c}_p \subset \mathscr{O}_{C,p} \subset \mathbb{C} + \mathfrak{r}_p \subset \widetilde{\mathscr{O}}_p.$$

Since the stalks of $\pi_*(\mathscr{O}_{\widetilde{C}})/\mathscr{O}_C$ are finite-dimensional $\mathbb{C}$-vector spaces, there must exist[3] some natural number c such that $\mathfrak{r}_p^c \subset \mathfrak{c}_p$, and hence

$$\mathbb{C} + \mathfrak{r}_p^c \subset \mathscr{O}_{C,p} \subset \mathbb{C} + \mathfrak{r}_p.$$

This indicates, roughly speaking, that all information about the singularity itself is captured by a choice of subspace of a finite-dimensional vector space

$$\frac{\mathscr{O}_{C,p}}{\mathfrak{r}_p^c} \subset \frac{\mathbb{C} + \mathfrak{r}_p}{\mathfrak{r}_p^c}. \tag{6.3}$$

The above conforms to our intuition that in order to know all data about a singularity, we need only know about those functions on it whose orders of vanishing are not too high, this is made precise by the following proposition.

[3] If we wish to make the dependence on the point clear, we will write c_p for the power corresponding to $p \in C$.

6.2.1.3 Proposition. *[Ser12] Let $\widetilde{C}$ be an irreducible smooth algebraic curve. Fix a finite subset of points on $\widetilde{C}$, and an equivalence relation on those points. Write these as $\{q_{ij}\}$, where $q_{1j}, \dots, q_{k_j j}$ are the points in the jth equivalence class and denote the set of equivalence classes by $\{p_1, \dots, p_\ell\}$. Define C to be the resultant quotient topological space, formed as the union*

$$(\widetilde{C} \setminus \{q_{ij}\}) \cup \{p_1, \dots, p_\ell\}$$

with the cofinite topology. For each $p \in C$, choose a subring $\mathscr{O}_p \subset \widetilde{\mathscr{O}}_p$ according to the rules:

i) For $i = 1, \dots, \ell$, $\mathscr{O}_{p_i}$ is a proper subring of $\widetilde{\mathscr{O}}_{p_i}$, and there exists $c_i \in \mathbb{N}$ such that

$$\mathbb{C} + \mathfrak{r}_{p_i}^{c_i} \subset \mathscr{O}_{p_i} \subset \mathbb{C} + \mathfrak{r}_{p_i}.$$

ii) For all $q \in C \setminus \{p_1, \dots, p_\ell\}$, set $\mathscr{O}_q = \widetilde{\mathscr{O}}_q$.

Then this data patches to give a sheaf of rings, endowing C uniquely with the structure of an algebraic curve. Conversely, any algebraic curve is determined up to isomorphism by this data.

The data described in the above proposition, which we will refer to as *Rosenlicht-Serre* data, is of two kinds: discrete and continuous. Of particular interest to us is choice of subrings that determine the singularities; by (6.3) we may think of this as a choice of points in an appropriate Grassmannian. In order to find a good notion of the 'Harder-Narasimhan type' of a curve, we will need a way to pull out some discrete invariant associated to that choice, which must then be combined in some way with the global data of the singular curve. A natural way to do this, which we pursue in Section 6.2.2, is to study a certain non-reductive instability stratification of these Grassmannians of local data. Thus, we proceed somewhat in the opposite direction than to the case of sheaves: rather than knowing the correct notion of Harder-Narasimhan type in advance, and showing that it can be instantiated in GIT[4], we will start by investigating a natural GIT set-up, and then interpreting it geometrically.

[4]To some extent it is purely a matter of historical accident that the notion of Harder-Narasimhan filtration did not originally come from GIT: as shown in [Sai13] it is possible to deduce its existence and basic properties purely from the GIT set-up, and an appropriate search for Kempf maximally destabilising one parameter subgroups, without knowing in advance what one is looking for.

6.2.2 Local study of curve singularities

We now turn to the local analogue of our moduli problem. This is a question that goes back to Zariski [ZT06], who investigated the so-called Zariski moduli space, of local rings of analytic curve singularities- that is, completed local rings of singular points on curves, up to isomorphism. As Zariski showed, this deceptively simple problem in fact leads to some non-trivial behaviour, and a lot of computational ingenuity was necessary to make any progress. Even before that, Ebey investigated the same question, in [Ebe65]. In our language, their work shows that in fact this is very naturally a problem in non-reductive GIT of the kind covered by the $\hat{U}$ theorems (Section 3.3), and applies certain elimination criteria to find explicitly, in some specific cases, a slice for the group action- the best that could be done without modern tools. In fact, even having non-reductive GIT does not immediately remove all of the difficulties associated to this problem: the condition $(\star)$ fails, and one simply translates the question into another, though perhaps more tractable, language. Fortunately for us, in order to answer our global moduli question we will not need to actually construct moduli spaces of curve singularities; we only need to understand the instability stratification as a discrete invariant.

6.2.2.1 Definition.

1) A curve singularity, or just singularity for short, is a ring $\mathscr{O}$ that appears as the local ring $\mathscr{O}_{C,p}$ for some curve C, singular at p. Denote the normalisation of the singularity- that is, the integral closure - by $\widetilde{\mathscr{O}}$.

2) We say two curve singularities are analytically isomorphic if they have isomorphic completions with respect to their maximal ideal.

3) Suppose that we have a curve C, which is singular at some $p \in C$. Then the normalisation $\pi : \widetilde{C} \to C$, plus completion with respect to the maximal ideal of the local ring at p, yields a map

$$\widehat{\mathscr{O}_{C,p}} \hookrightarrow (\pi_* \mathscr{O}_{\widetilde{C}})_q \simeq \bigoplus_{i=1}^{\ell} \mathbb{C}[[t_i]].$$

The number ℓ is the number of branches of the singularity. If $\ell = 1$ we say that the singularity is unibranch; otherwise, we say that it is multibranch. Equivalently, ℓ is the number of points in the fibre over p of the normalisation map.

4) We define the δ-invariant, or genus, of the singularity to be

$$\delta \coloneqq \dim_{\mathbb{C}} \widetilde{\mathscr{O}}/\mathscr{O} < \infty$$

This term agrees with the usual, global, notion of genus, as δ is the contribution of the singularity to the arithmetic genus of the curve. Note that, since it is finite, it is also invariant under passing to completions, so is an analytic isomorphism invariant.

The isomorphism on the right of (3) above is of course not canonical, but is defined up to the action of the automorphism group $\mathrm{Aut}(\bigoplus_{i=1}^{\ell} \mathbb{C}[[t_i]])$. Conversely, given two curves C_1, C_2 with singularities $p_i \in C_i$ having isomorphic completed local rings, we know by the universal property of normalisation[5] that there is an isomorphism $\varphi : \mathscr{O}_{\widetilde{C}_1,p_1} \to \mathscr{O}_{\widetilde{C}_2,p_2}$ making the below diagram commute.

$$\begin{array}{ccc} \mathscr{O}_{\widetilde{C}_1,p_1} & \overset{\varphi}{\dashrightarrow} & \mathscr{O}_{\widetilde{C},p_2} \\ {\scriptstyle \pi_1^*}\uparrow & & {\scriptstyle \pi_2^*}\uparrow \\ \mathscr{O}_{C_1,p_1} & \xrightarrow{\simeq} & \mathscr{O}_{C_2,p_2} \end{array} \tag{6.4}$$

We can then choose, for $j = 1, 2$, isomorphisms $\bigoplus_{i=1}^{\ell} \mathbb{C}[[t_i]] \simeq \mathscr{O}_{\widetilde{C}_j,p_j}$, to get a commutative diagram

$$\begin{array}{ccc} \bigoplus_{i=1}^{\ell} \mathbb{C}[[t_i]] & \xrightarrow{\simeq} & \bigoplus_{i=1}^{\ell} \mathbb{C}[[t_i]] \\ \uparrow & & \uparrow \\ \widehat{\mathscr{O}}_{C_1,p_1} & \xrightarrow{\simeq} & \widehat{\mathscr{O}}_{C_2,p_2} \end{array} \tag{6.5}$$

where the top horizontal arrow is induced by φ together with these choices. Thus, analytic isomorphism classes of ℓ-branched curve singularities correspond to certain orbits under this group action on the set of subrings of $\bigoplus_{i=1}^{\ell} \mathbb{C}[[t_i]]$.

In order to turn this into a GIT problem, we must fix some discrete invariants, instructed by the Rosenlicht-Serre construction. Let us suppose we are interested in singularities of genus δ. Since we only care about information that is invariant under analytic isomorphism, we may assume that in fact the semi-local ring of the normalisation is just $\bigoplus_{i=1}^{\ell} \mathbb{C}[t_i]_{(t_i)}$. The radical of this semi-local ring is the ideal $\oplus_i(t_i)$. For any fixed n, a subring $\mathscr{O} \hookrightarrow \bigoplus_{i=1}^{\ell} \mathbb{C}[t_i]$ which contains the nth power of

[5] Of the schemes given by the local rings, not the curves themselves.

the radical is then the same thing as a choice of subring

$$\overline{\mathscr{O}} \coloneqq {}^{\mathscr{O}}\!/_{\oplus_i(t_i^n)} \hookrightarrow \bigoplus_{i=1}^{\ell} {}^{\mathbb{C}[t_i]}\!/_{t_i^n}.$$

We can easily find a space that parameterises all such choices.

6.2.2.2 Definition. We define

$$\mathcal{A}^{\ell}_{n,\delta} \coloneqq \{\text{Genus } \delta \text{ subrings } \overline{\mathscr{O}} \subset \bigoplus_{i=1}^{\ell} {}^{\mathbb{C}[t_i]}\!/_{t_i^n}\} \subset \mathrm{Gr}(\bigoplus_{i=1}^{\ell} {}^{\mathbb{C}[t_i]}\!/_{t_i^n}, \delta) =: \mathrm{Gr}^{\ell}_{n,\delta}\,.$$

The Grassmannian on the right parameterises vector subspaces of codimension δ (or equivalently, δ-dimensional quotients).

6.2.2.3 Lemma. *[Ish80] The subset $\mathcal{A}^{\ell}_{n,\delta}$ is a closed subscheme $\mathrm{Gr}^{\ell}_{n,\delta}$, and is hence a projective variety.*

6.2.2.4 Remark. The above lemma is used by Ishii in [Ish82], who uses it to construct, using reductive GIT, quasi-projective moduli spaces of unibranch curves with fixed arithmetic genera at each singularity. Our construction may therefore be viewed as a refinement of theirs, with a natural generalisation to the multibranch case.

6.2.2.5 Definition. The *re-parameterisation group* $\mathbb{G}_n$, is the group that acts

$$\mathbb{G}_n \times {}^{\mathbb{C}[t]}\!/_{(t^n)} \to {}^{\mathbb{C}[t]}\!/_{t^n}$$

$$t \longmapsto \alpha_1 t + \alpha_2 t^2 + \cdots + \alpha_{n-1} t^{n-1}$$

where $\alpha_1 \in \mathbb{C}^*$ and $\alpha_i \in \mathbb{C}$ for $i = 2,..,n-1$. This group may be expressed as a semi-direct product $\mathbb{G}_n \simeq U_n \rtimes \mathbb{C}^*$, where U_n is its unipotent radical. In a representation, it looks like:

$$\mathbb{G}_{N,i} = \left\{ \begin{pmatrix} t_i^N & f_{N,(N-1)} & f_{N,(N-2)} & \dots & \dots & \alpha_N \\ \dots & \dots & \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & \dots & 0 & t_i^2 & f_{3,2} & a_2 \\ 0 & \dots & \dots & 0 & t_i & \alpha_1 \\ 0 & \dots & \dots & \dots & 0 & 1 \end{pmatrix} \mid \alpha_i \in \mathbb{C}\ \forall i, \alpha_1 \neq 0 \right\},$$

for suitable uniquely determined functions $f_{i,j}(\alpha_1, \ldots, \alpha_N)$.

6.2.2.6 Remark. Note that, from the point of view of the $\hat{U}$ theorems, we are in good shape so far, in that there is a natural a choice of distinguished $\mathbb{G}_m$: that is, there is a unique one parameter subgroup in the reparameterisation group that has positive weights on the Lie algebra of the unipotent radical. Hence below, where we consider products of such groups, the so-called 'positive cone' of Section 3.4 consists precisely of those one parameter subgroups that project to give the positive one parameter subgroup along each of the branches.

The action induced by the automorphism group $\mathrm{Aut}(\bigoplus_{i=1}^{\ell} \mathbb{C}[[t_i]])$ on $\mathrm{Gr}_{n,\delta}^{\ell}$ is, up to a finite group, nothing but the product of re-parameterisation groups. More precisely, the action of $\mathrm{Aut}(\bigoplus_{i=1}^{\ell} \mathbb{C}[[t_i]])$ factors through the group

$$\mathbb{G}_n^{\ell} \coloneqq (\prod_{i=1}^{\ell} \mathbb{G}_n) \rtimes S_{\ell},$$

where S_ℓ is the permutation group of the branches. Thus, the natural GIT problem for moduli of curve singularities is given by the action of $\mathbb{G}_n^{\ell}$ on $\mathrm{Gr}_{n,\delta}^{\ell}$. To finish off the set-up, we linearise the action in the natural way, by the usual Plücker embedding of the Grassmannian.

We can therefore study the instability stratification associated to this GIT picture, in the following way. Although the group is not reductive, the same principles apply as in Section 2.3. We first choose a maximal torus T, and an invariant inner product on its Lie algebra[6]. Then we analyse how the convex hull of the T-weights of a given $A \in \mathcal{A}_{n,\delta}^{\ell}$ changes as we sweep out by the rest of the group action. That is, for each $h \in \mathbb{G}_n^{\ell}$ we consider the convex hull of those T-weights that coorespond to coordinates which are non-zero for hA. In this way we associate to each point A in our parameter space its *instability type*. This is defined to be that weight of T that arises as the maximum (with respect to our chosen norm) of the closest points to zero of the convex hull of the T-weights of hA, as h runs through our group.

6.2.2.7 Remark. We emphasise that, though there is nothing to stop us defining it, unlike in the reductive case if there is no appropriate notion of stability in a particular non-reductive GIT set-up that can be defined by a Hilbert-Mumford criterion, there is no reason to expect the instability stratification as defined above to be particularly useful. (An extreme example would be when the group is unipotent!)

[6] As elsewhere, the choice of norm will not enter into our discussion, since there is a natural choice: the norm induced by the Killing form of the relevant special linear group, $\mathrm{SL}(\mathbb{C}[t]/(t^n))$

6.2.2.8 Definition. For a point $A \in \mathrm{Gr}^{\ell}_{n,\delta}$, we define the *Rosenlicht-Serre type* to be the associated instability type in the given GIT picture.

6.2.2.9 Remark. We note that the definition of the Rosenlicht-Serre type of a singularity, as given above, actually depends on the choice of n. In order for this to be an authentic invariant of curve singularities, or of the singular curves themselves, we must of course show that it is in some sense essentially independent of this choice (at least for n sufficiently large). We will do this, for the case of unibranch singularities, in Section 6.2.3.

6.2.3 Unibranch singularities

In our consideration of the global moduli problem, we will be principally be concerned with the case when all singularities of the curve are unibranch. In this case, the GIT set-up given above simplifies considerably, and an interpretation of the Rosenlicht-Serre type can readily be given.

Now let C be a curve with a singularity at $p \in C$, assumed to be unibranch, let $\pi : \widetilde{C} \to C$ be the normalisation as usual, and let $q = \pi^{-1}(p)$. Then in this case we have, for the semi-local ring,

$$(\pi_* \mathscr{O}_{\widetilde{C}})_p \cong \mathscr{O}_{\widetilde{C},q}$$

as rings, so that there is an inclusion of local rings $\mathscr{O}_{C,p} \subset \mathscr{O}_{\widetilde{C},q}$, and $\mathfrak{r}_p$ is just the maximal ideal of $\mathscr{O}_{\widetilde{C},q}$, denoted by $\mathfrak{m}_q$. Fixing $n = n_p$ as in the Rosenlicht-Serre construction, we see that the singularity is determined by a choice

$$\mathscr{O}_{C,p} \big/ \mathfrak{m}_q^n \subset \mathscr{O}_{\widetilde{C},q} \big/ \mathfrak{m}_q^n.$$

Taking completions of both sides has no effect, so we are really interested in choices of subring

$$R \subset \mathbb{C}[t] \big/ (t^n).$$

6.2.3.1 Definition.

1) Let $\nu_q : \mathscr{O}_{\widetilde{C},q} \to \mathbb{N} \cup \{\infty\}$ be the valuation of the local ring of the point q in the normalisation. Then we define the *semigroup of values* of the singularity at $p \in C$, to be

$$\Sigma_p(C) \coloneqq \{s \in \mathbb{N} \mid \exists f \in \mathscr{O}_{C,p} \text{ with } \nu_q(\pi^* f) = s\} \subset \mathbb{N}.$$

2) Since the δ-invariant is finite, there must exists $c = c_p \in \mathbb{N}$ such that $\mathbb{N}_{\geq c} \subset \Sigma_p(C)$. This number is called the *conductor*. The elements of $\mathbb{N} \setminus \Sigma_p(C)$, of which there are δ, are called the *gaps* in the semigroup.[7]

3) For a semigroup $\Sigma \subset \mathbb{N}$, we define the weight of Σ, to be the sum of the gaps.

4) Given a fixed semigroup $\Sigma \subset \mathbb{N}$, generated by $\alpha_1, \ldots, \alpha_k$, there is a there is a distinguished choice of curve having that semigroup: namely the curve $\operatorname{Spec} \mathbb{C}[t^{\alpha_1}, \ldots, t^{\alpha_k}]$. This canonical representative of the equivalence class will be referred to as the *monomial curve with semigroup* Σ, and has an important role to play in what follows. We will denote this curve by R_Σ.

6.2.3.2 Remark. We note that all of the above are analytic invariants. In addition, these definitions are consistent with the terms defined in Section 6.2.2, after passing to the completed local ring.

Our aim is to establish that, in the unibranch case, the data of the Rosenlicht-Serre type as we have defined it is independent of the choices involved, and in fact equivalent to the semigroup of values. This establishes, at least at the local level, a correspondence between GIT and a discrete invariant of curve singularities, and will be used extensively in what follows, where we try to lift this to the global moduli problem.

6.2.3.3 Remark. As one would expect, there is a natural definition of the semigroup of values of an ℓ-branch singularity, which uses the valuations at all of the points of the normalisation lying over the singular point, and gives a semigroup $\Sigma \subset \mathbb{N}^\ell$. This would be an important notion in extending the results of Section 6.3.3 and Section 6.3.4 to cover the multibranch case.

In the first place, we would like to be able to study all unibranch singularities of a given genus simultaneously, and to do that, we need to show that there is a single n that works for all of them- in effect proving a minature boundedness result. We can do that with the following easy observation.

6.2.3.4 Lemma. *If $n \geq 2\delta$, then we have $\mathscr{O} \supset \mathfrak{r}^n$ for all unibranch singularities of genus δ. Hence, all genus δ unibranch singularities correspond to points of* $\mathrm{Gr}_{2\delta}$.

[7]For us, unless stated otherwise, semigroup will always mean what is often called a numerical semigroup, that is a subset $\Sigma \subset \mathbb{N}$ containing $0 \in \mathbb{N}$ and closed under addition.

Proof. If $\mathscr{O}$ is a unibranch curve singularity of genus δ, then its semigroup of values has δ gaps. If the conductor is c, then for every $s \in \Sigma$, we must have $(c-1) - s \notin \Sigma$. Hence at most half of the integers in $\{1, \dots, (c-1)\}$ can be in Σ, and thus $\delta \geq c/2$. □

6.2.3.5 Notation. Let us simpify our notation in the unibranch case and write

$$\mathrm{Gr}_\delta \coloneqq \mathrm{Gr}^1_{2\delta,\delta},$$

setting $n = \delta$ light of the above proposition and dropping the notation for the number of branches. Similarly, we write $\mathcal{A}_\delta$ for $\mathcal{A}^\ell_{2\delta,\delta}$.

Now fix a genus of singularity we are interested in, $\delta > 0$. As we saw above, this action can be linearised in the standard way via the Plücker embedding

$$\mathcal{A}_\delta \subset \mathrm{Gr}_\delta \hookrightarrow \mathbb{P}\big((\textstyle\bigwedge^\delta \mathbb{C}[t]/(t^n))^*\big).$$

The following proposition relates our notion of Rosenlicht-Serre type, in the unibranch case, to many other well-known invariants.[8]

6.2.3.6 Proposition. *Two unibranch curve singularities have the same Rosenlicht-Serre type if and only if their semigroups have the same weight.*

Proof. We must show that the instability type, that is, the weight of the maximally destabilising one parameter subgroup, is equivalent data to the weight of the semigroup. But since there is a one-parameter subgroup $\lambda : \mathbb{G}_m \to \mathbb{G}$ having positive weights on the unipotent radical, and there is nothing else in the group, the search for a Kempf one parameter subgroup reduces to analysis of the Bialynicki-Birula stratification with respect to λ. If $A \in \mathcal{A}_\delta$, we claim that the point to which it flows under $\lambda(t)$ as $t \to 0$ is the monomial curve singularity with the same semigroup as A. Indeed, Plücker coordinates on $\mathcal{A}$ correspond to choices of δ elements of our chosen basis $\{1, t, \cdots, t^{n-1}\}$ of $\mathbb{C}[t]/(t^n)$. We can find which such coordinates are nonzero for A by choosing a basis of A, given by $f_1, \dots, f_\delta$, and expanding out the wedge $f_1 \wedge \cdots \wedge f_\delta$; moreover, this is independent of the choice of basis. So for each $\gamma_1, \dots, \gamma_\delta \in \Sigma(A)_{\leq 2\delta}$, take $f_i \in A$ having that order of vanishing, that is $f_i = t^{\gamma_i} + t^{\gamma_i+1}(\dots)$, where we may as well take the f_i to be monic. Then we have

$$f_1 \wedge \dots f_\delta = t^{\gamma_1} \wedge \cdots \wedge t^{\gamma_\delta} + \text{higher order terms.}$$

[8] For example, it was shown by Zariski in [ZT06] that two unibranch plane curve singularities have the same semigroup if and only if they are topologically equivalent.

This shows that the minimally weighted non-zero coordinate, i.e. the one to which A will flow as $t \to 0$, corresponds to the monomial curve singularity with the same semigroup.

□

6.2.3.7 Remark. This shows that the notion of Rosenlicht-Serre type is, at least in the unibranch case, entirely canonical. The proof shows that it is essentially independent of n, if n is taken sufficiently large depending on δ (cf. the Harder-Narasimhan stratification of the Quot scheme in Section 5.2). Thus it makes sense to talk about *the* Rosenlicht-Serre type of a unibranch singularity. In the multibranch case, it would be interesting to see whether there is a similar characterisation of Rosenlicht-Serre type at the local level, in intrinsic terms.

6.3 Moduli of Curves of fixed Rosenlicht-Serre type

We now turn our attention to the construction of moduli spaces of unstable, i.e. singular, curves. We will follow a method very like the one we used in Chapter 5 for sheaves; one important difference will be that we will not be studying the instability stratification associated to the classical GIT construction for moduli of stable curves. Rather, we will set things up slightly differently, working effectively with a smaller, and non-reductive, group, and a different parameter space. Nevertheless, this new set-up may be thought of as being in effect a small perturbation of the Baldwin-Swinarski construction, which we review below. The instability stratification associated to this group action will then be expressible in terms of the local data of the curve singularities, as studied in Section 6.2.2. We will give proofs only in the unibranch case, but it is likely that what we say can readily be generalised to include multibranch singularities as well, and we will give indications throughout of how this might be done.

6.3.1 From local to global

Motivated by Section 6.2.2, we define the following discrete invariant, which will be the analogue of instability type in our moduli problem.

6.3.1.1 Definition.

1) We say that a curve is unibranch if all of its singularities are unibranch; that is, if all of the completed local rings of its normalisation are isomorphic to $\mathbb{C}[[t]]$.

2) Let C be a unibranch curve with singularities at $p_1, \dots, p_\ell \in C$. We define the *Rosenlicht-Serre type* of C to be the vector $\sigma = (\Sigma_1, \dots, \Sigma_\ell)$ of numerical semigroups of the singularities of C.

Our ultimate goal is to construct moduli spaces for unibranch projective algebraic curves of fixed Rosenlicht-Serre type. To do this, we first gather some preliminary results that will be useful in lifting local, i.e. Rosenlicht-Serre, information about a curve up to the global level, as well as providing motivation and intuition for what comes later.

6.3.1.2 Notation. We adopt the following assumptions and notation from now until the start of Section 6.3.2. Let C be a curve of arbitrary genus, embedded in $\mathbb{P}^r$ via a very ample line bundle L. Denote by $\widetilde{C} \to C$ the normalisation map, and let $\widetilde{L} = \pi^* L$. For simplicity, we assume also that C has a unique singularity at $p \in C$, with $\pi^{-1}(p) = \{q_1, .., q_s\}$, so that C is not yet assumed to be unibranch. Fix $n \in \mathbb{N}$, and write $D = n \sum q_i$ for the corresponding divisor on $\widetilde{C}$. If C has more than one destabilising singularity, the analogues of these results still hold, although the notation is more cumbersome.

6.3.1.3 Lemma. *Let $\widetilde{L}$ be sufficiently ample that $\widetilde{L}_D \coloneqq \widetilde{L} \otimes \mathscr{O}_{\widetilde{C}}(-D)$ is still very ample on $\widetilde{C}$. Then, if $n \gg 0$ depending only on the singularity at $p \in C$, there is a sequence of vector spaces with inclusions*

$$H^0(\widetilde{C}, \widetilde{L}) \supset H^0(C, L) \supset H^0(\widetilde{C}, \widetilde{L}_D)$$

whose dual linear projections define rational maps

$$\mathbb{P}(H^0(\widetilde{C}, \widetilde{L})^*) \dashrightarrow \mathbb{P}(H^0(C, L)^*) \dashrightarrow \mathbb{P}(H^0(\widetilde{C}, \widetilde{L}_D)^*)$$

which restrict to (rational) maps

$$\widetilde{C} \to C \dashrightarrow \widetilde{C}$$

where the first map is π and the composition is an isomorphism.

Proof. It is clear that π^* gives the first inclusion. For the second, consider the short exact sequence on C,

$$0 \to \mathscr{O}_C \to \pi_*(\mathscr{O}_{\widetilde{C}}) \to \pi_*(\mathscr{O}_{\widetilde{C}})/\mathscr{O}_C \to 0.$$

Tensoring with L and taking global sections, we have

$$0 \to H^0(C, L) \to H^0(\widetilde{C}, \widetilde{L}) \to F_p \to 0$$

where F_p is a skyscraper sheaf supported only at p. If $s \in H^0(\widetilde{C}, \widetilde{L}_D)$, then the stalk of s at p must annihilate F_p. Hence it maps to zero under the second map in the above sequence, which means that it is contained in $\pi^* H^0(C, L)$. □

Lemma 6.3.1.3 provides a global version of the Rosenlicht-Serre construction, in a sense that we will make precise as we go along. As explained above, our task is to understand the instability stratification (or at least *an* instability stratification) for curves in a geometrically meaningful way: that is, roughly speaking, we are looking for maximally destabilising one parameter subgroups, in the sense of Section 2.3. Since it is evidently the singularities that make the curve moduli unstable it is natural to look for one parameter subgroups that somehow best capture the information about the singularities. First, though, we recall from [BFMV14] 4.4 how to find limits of Hilbert points of curves under one-parameter subgroups, using Gröbner bases.

6.3.1.4 Definition. Let $W = \langle x_0, \ldots, x_h \rangle$ be a vector space, $\lambda : \mathbb{G}_m \to \mathrm{SL}(W)$ a one parameter subgroup, acting diagonally with respect to these coordinates via

$$\lambda(t)(x_i) = t^{w_i} x_i.$$

Define a partial monomial ordering on $\mathbb{C}[x_0, \ldots x_h]$ in the following way. Given two monomials

$$x^\alpha = \prod_{i=1}^{h} x_i^{\alpha_i}$$

$$x^\beta = \prod_{i=1}^{h} x_i^{\beta_i}$$

we say $x^\alpha \prec_\lambda x^\beta$ if either

$$\deg x^\alpha < \deg x^\beta,$$

or

$$\deg x^\alpha = \deg x^\beta \text{ and } \mathrm{wt}_\lambda(x^\alpha) < \mathrm{wt}_\lambda(x^\beta),$$

where $\mathrm{wt}_\lambda(x^\alpha)$ denotes the λ-weight of the monomial. For $f \in \mathbb{C}[x_0, \dots x_h]$, denote by $\mathrm{lt}_\lambda f$ the leading term of f with respect to this ordering; that is, the sum of the monomials appearing in f of highest λ-weight, of those of highest degree.

6.3.1.5 Proposition. *([BFMV14] Corollary 4.16) Let $C \subset \mathbb{P}(W^*)$ be a curve of Hilbert polynomial $P(x) = ax + b$, so that*

$$[C] \in \mathrm{Hilb}(\mathbb{P}(W^*), a, b),$$

and denote by I the corresponding ideal. Then the limit point of $[C]$ in the Hilbert scheme is given by the ideal

$$\mathrm{in}_\lambda I \coloneqq \{\mathrm{lt}_\lambda f \mid f \in I\}.$$

6.3.1.6 Notation. From now until Section 6.3.2, write $W = H^0(C, L)$, and $V = H^0(\widetilde{C}, \widetilde{L}(-D))$, so that we have $W \supset V$.

6.3.1.7 Lemma. *Let $\lambda : \mathbb{G}_m \to \mathrm{SL}(W)$ be any 1PS such that the minimal weight space of λ is V. Then the flat limit, $\overline{C}$, of C under λ contains $\widetilde{C}$ as a component. More precisely, we have*

$$\overline{C} \cap \mathbb{P}(V^*) = \widetilde{C}$$

embedded as above, and the remainder of $\overline{C}$ is a configuration of curves $C_1, ..C_\ell$ meeting $\widetilde{C}$ at finitely many points.

Proof. Let $z_1, ..z_n \in V$ be coordinates on $\mathbb{P}(V^*)$, and extend to coordinates on $\mathbb{P}(W^*)$ by adding $x_1, .., x_m$. Let G be a Grobner basis for the ordering defined by λ. Then we may write

$$G = \{f_1, .., f_r, g_1, .., g_s\},$$

where $f_i = f_i(x_1, .., x_m, z_1, ..z_n)$ has at least one term containing some x_j, and the $g_i = g_i(z_1, .., z_n)$ have no dependence on any of the x_j's. Since the z_i are all in the minimal weight space, we have $\mathrm{lt}_\lambda(g_j) = g_j$ for all $j = 1, .., s$. If I is the ideal of C,

the ideal of the limit which must therefore be

$$in_\lambda I = (\mathrm{lt}(f_1), .., \mathrm{lt}(f_r), g_1, .., g_s).$$

Since each f_i contains a term depending on some x_j, so must its leading term. We have

$$\mathbb{P}(V^*) = \{x_i = 0 \ \forall \ i = 1, .., m\}$$

so the ideal of the intersection of the limit curve with this linear subspace is given by

$$J = (x_1, .., x_m, g_1, .., g_s).$$

On the other hand, by Lemma 6.3.1.3, the restriction to C of projection

$$\mathbb{P}(W^*) \dashrightarrow \mathbb{P}(V^*)$$

has image closure equal to $\widetilde{C}$. But the ideal of this image closure, considered as a subvariety of $\mathbb{P}(V^*)$, is nothing but $I \cap \mathbb{C}[z_0, .., z_n]$. By standard elimination ordering properties of Grobner bases, $(g_1, .., g_s)$ is a Grobner basis for this intersection.

□

Proposition 6.3.1.7 tells us that there is a certain class of one parameter subgroups that allow us to separate the information of the normalisation of a curve from the data of its singular points. Of course, unless we are very careful, in taking a limit under such a one parameter subgroup we will end up with something highly non-reduced that contains no information about C's singularities at all. The following proposition shows that this is not always the case, and will form the foundation for our investigation in Section 6.3.

6.3.1.8 Proposition. *Let $p \in C$ be the unique singularity, and suppose that it is unibranch, i.e. $q = \pi^{-1}(p) \in \widetilde{C}$. Let c be the conductor of the numerical semigroup Σ associated to the singularity at p. Choose coordinates $x_1, .., x_r, y, z_1, ..., z_s$ on $\mathbb{P}(W^*)$, labelled according to their valuations considered as functions on the normalisation, so that*

$$v_i \coloneqq \nu_q(\pi^* x_i) < c \text{ for all } i = 1, .., r$$

$$\nu_q(\pi^* y) = c$$

$$\nu_q(\pi^* z_i) > c$$

Let λ be the 1PS of $GL(W)$ defined by the weights

$$\lambda(t)x_i = t^{-v_i}x_i$$

$$\lambda(t)y = t^{-c}y$$

$$\lambda(t)z_j = t^{-c}z_i.$$

Then the limit of C under λ is the union of the two curves

$$R_\Sigma \subset \mathbb{P}(\langle x_1, ..., x_r, y\rangle^*)$$

$$\widetilde{C} \subset \mathbb{P}(\langle y, z_1, .., z_s\rangle^*),$$

where R_Σ is as defined at Definition 6.2.3.1, the monomial rational curve with semigroup Σ, parameterised in coordinates by minus the weights of λ, meeting $\widetilde{C}$ in a node at the point corresponding to q.

Proof. First choose as above a Gröbner basis

$$I_C = (f_1(x, y, z), ..., f_r(x, y, z), g_1(z), .., g_s(z)),$$

so that

$$I_{\overline{C}} = (\mathrm{lt}(f_1)(x, y, z), .., \mathrm{lt}(f_r)(x, y, z), g_1(z), .., g_s(z)).$$

Then we may choose generators for the remaining relevant ideals as follows:

$$I_{\widetilde{C}} = (g_1(z), ..., g_s(z), x_1, .., x_m)$$

$$I_{R_\Sigma} = (h_1(\underline{x}, y), .., h_n(\underline{x}, y), z_1, .., z_s).$$

Now choose a set of generators for the semigroup

$$\Sigma = \langle \gamma_1, .., \gamma_\ell \rangle.$$

Then we have

$$\hat{\mathcal{O}}_{R_\Sigma, p} \cong \mathbb{C}[[t^{\gamma_1}, .., t^{\gamma_\ell}]],$$

and there exist

$$g_i = s^{\gamma_i} + \sum_{\alpha_{i,j} > \gamma_i} a_{ij} s^{\alpha_{i,j}}$$

such that

$$\hat{\mathscr{O}}_{C,p} \cong \mathbb{C}[[g_1, .., g_\ell]].$$

Write ν_{R_Σ}, ν_C for the valuation in $\widehat{\mathscr{O}}_{R_\Sigma,p}$, $\widehat{\mathscr{O}}_{C,p}$ respectively. For convenience, let

$$w_\lambda : \mathbb{C}[x, y, x] \to \mathbb{Z}$$

be the function that outputs minus the λ-weight of a ring element. Abusing notation slightly, we will write $\hat{f}_{R_\Sigma}$, $\hat{f}_C$, for the images of an element $f \in W$ in the completed local rings $\widehat{\mathscr{O}}_{R,p}$ and $\widehat{\mathscr{O}}_{C,p}$, respectively.

Claim 1: $I_{\overline{C}} \leq I_R \cap I_{\widetilde{C}}$

First, Lemma 6.3.1.7 shows that $I_{\widetilde{C}} \geq I_{\overline{C}}$, so we will show $I_R \geq I_{\overline{C}} + (z_1, \ldots, z_s)$.

Let $\varphi_\Sigma : \mathbb{P}^1 \to \mathbb{P}(W^*)$ be the parameterisation map for R_Σ, so that $I_{R_\Sigma} = \ker \varphi_\Sigma^*$ Take a homogeneous $f \in in_\lambda I + (z_1, \ldots, z_s)$, and write $f = f_1 + f_2$ where $f_1 \in \in_\lambda I$ and $f_2 \in (z_1, \ldots, z_s)$. Then $\exists g \in I$ s.t. $\mathrm{lt}_\lambda\, g = f_1$. Write $f_1 = f_{11}(\underline{x}, \underline{y}) + f_{12}$ where $f_{12} \in (z_1, \ldots, z_s)$. Suppose $\hat{f}_C \neq 0$; then let $\nu_C(f) = \nu_1 < \infty$. So we can write

$$f_1 = a_{\nu_1} t^{\nu_1} + \sum_{\gamma > \nu_1} a_\gamma t^\gamma,$$

with $a_{\nu_1} =\neq 0$. Let $h = g - f_1$. We claim that $\nu_C(h) < \nu_1$. Indeed, if $h = \sum_i M_i$ is a sum of monomials, then $\mathrm{wt}_\lambda(M_i) > \mathrm{wt}_\lambda(f_1)$ for all i. Moreover, we may assume that $f_1 = f_1 1$ without loss of generality, so that $\mathrm{wt}_\lambda(f_1) = \nu_C(f_1) = \nu_1$. And for each monomial M_i, we have $\nu_C(M_i) \leq \mathrm{wt}_\lambda(M_i)$: this shows $\nu_C(h) < \nu_1$. Since $g = h - f_1 \in I$ we know that $\hat{g}_C = 0$, which contradicts $a_{\nu_1} \neq 0$. Hence $I_R \geq I_{\overline{C}} + (z_1, \ldots, z_s)$.

Claim 2: $\overline{C} \cap \mathbb{P}(x_i = 0\ \forall i) = \widetilde{C}$

Considering the corresponding ideals, it suffices to note that the leading terms $\mathrm{lt}(f_i(x, y))$ must all lie in $(x_1, .., x_m)$, since otherwise the point $[\underline{0} : 1 : \underline{0}]$, where y is the only non-zero coordinate, would not lie on $\overline{C}$.

Claim 3: $\overline{C} \cap \mathbb{P}(z_j = 0\ \forall j) = R$

By Claim 1, it suffices to show that $I_R \leq I_{\overline{C}} + (z_1, .., z_s)$. It is enough to show this for elements $f(x, y) \in I_R$ which are sums of monomials M_i such that there exists

d with $\nu_{R_\Sigma}(M_i) = d$ for all i. Considering the image, $\hat{f}_C$, of such an f in $\hat{\mathscr{O}}_{C,p}$, we see that its valuation is strictly higher than d, since the lowest powers of s cancel. Hence there is a correction term g such that $f + g \in I$, and the image $\hat{g}_{R_\Sigma}$ in $\widehat{\mathscr{O}}_{R,p}$ has $\nu_{R_\Sigma}(g) > d$. Since the generators of the semigroup have λ weight equal to their degrees (as polynomials in t) in the completed local ring, we may choose a g with $w_\lambda(g) > w_\lambda(f)$. Hence $\mathrm{lt}_\lambda(f+g) = f$, and $f \in I_{\overline{C}}$. Hilbert polynomial considerations rule out the possibility of other components. □

6.3.2 The GIT problem

Our first task is to find a parameter space for curves of fixed degree, genus, and Rosenlicht-Serre type up to the action of a linear algebraic group. Motivated by Section 6.3.1, we will attempt to do this by sandwiching each curve between two copies of its normalisation, so that we can make use of the invertibility of their dualising sheaves and hope to do something similar to standard GIT constructions of moduli of curves. First we will need some notation. Let C be a unibranch singular curve of genus g, and let $\{p_1, .., p_\ell\}$ be its singularities, with semigroups of values $(\Sigma_1, \ldots, \Sigma_\ell)$ having genera $(\delta_1, \ldots, \delta_\ell)$, and let $\delta = \sum_i \delta_i$, the contribution to the genus coming from the singularities. As above, we have $\pi : \widetilde{C} \to C$, the normalisation; this will be a smooth curve of genus $\widetilde{g} = g - \delta$. For $i = 1, \ldots, \ell$, denote by $q_i = \pi^{-1}(p_i)$ the point of $\widetilde{C}$ lying over p_i. For $n \gg 0$, define the divisor

$$D = n \sum_{i=1}^{\ell} q_i$$

on $\widetilde{C}$, and then fix $N \gg n \gg 0$, so that the invertible sheaves $\omega_{\widetilde{C}}^{\otimes N}$ and $\omega_{\widetilde{C}}^{\otimes N}(D)$ are both very ample, and we have the embeddings

$$\widetilde{C} \xhookrightarrow{\omega_{\widetilde{C}}^{\otimes N}} \mathbb{P}^a \tag{6.6}$$

$$\widetilde{C} \xhookrightarrow{\omega_{\widetilde{C}}^{\otimes N}(D)} \mathbb{P}^A \tag{6.7}$$

where $a = N(2\widetilde{g} - 2) - \widetilde{g} + 1$ and $A = a + n\ell$. These embed the curve $\widetilde{C}$ with respective degrees $d_1 \coloneqq N(2\widetilde{g} - 2)$ and $d_3 \coloneqq N(2\widetilde{g} - 2) + n\ell$. Our aim is to exploit these embeddings, which together make up for the lack of an invertible dualizing sheaf on C.

Now fix a vector space V_D, of dimension A and a linear subspace $V \subset V_D$ of codimension ℓ. Let

$$\rho : \mathbb{P}^A \dashrightarrow \mathbb{P}^a$$

be the associated linear projection. Consider the following Hilbert schemes

$$\mathcal{H}_1 = \mathrm{Hilb}(\mathbb{P}^A, d_1, \widetilde{g})$$

$$\mathcal{H}_2 = \mathrm{Hilb}(\mathbb{P}^A, d, g)$$

$$\mathcal{H}_3 = \mathrm{Hilb}(\mathbb{P}^a, d_3, \widetilde{g})$$

so that $\mathcal{H}_1$ is the Hilbert scheme of subschemes in $\mathbb{P}^A$ of Hilbert polynomial $P_1(t) = d_1 t - \widetilde{g} + 1$, and the others are defined similarly. Define also

$$\mathrm{Gr}_D \coloneqq \{\text{codimension } \delta \text{ subspaces of } {}^{V_D}\!/_V\}$$

and finally, assembling all of these together, we consider the projective scheme

$$\mathscr{H} \coloneqq \mathcal{H}_1 \times \mathcal{H}_2 \times \mathcal{H}_3 \times (\mathbb{P}^a)^\ell \times \mathrm{Gr}_D \,.$$

Now we can introduce our parameter space.

6.3.2.1 Proposition. *Fix points $q_1 \ldots, q_\ell \in \mathbb{P}^a$, and notation as above. Then there is a locally closed subscheme*

$$\mathcal{S} \subset \mathscr{H}$$

given by those points

$$([C_1], [C_2], [C_3], (q_1, \ldots, q_\ell), W) \in \mathcal{H}_1 \times \mathcal{H}_2 \times \mathcal{H}_3 \times (\mathbb{P}^a)^\ell \times \mathrm{Gr}_D$$

such that the following conditions hold.

i) q_i *lie on* C_3 *for* $i = 1, \ldots, \ell$.

ii) $[C_1] \in \mathcal{H}_1$ *and* $[C_3] \in \mathcal{H}_3$ *are isomorphic smooth curves, via a morphism*

$$\varphi : C_3 \to C_1$$

that extends the rational map ρ, and all their connected components have genus at least 2.

iii) $[C_3] \in \mathcal{H}_3$ *is embedded by* $\omega_{C_3}^{\otimes n}$*, and* $[C_1] \in \mathcal{H}_1$ *is embedded via* $\omega_{C_3}^{\otimes n}(\varphi^* D)$

iv) $[C_2]$ *has normalisation* $\pi : C_1 \to C_2$*, and the normalisation map extends the rational map* ρ*, and*

$$\varphi \circ \pi^{-1}(\ \mathrm{Sing}(C_2)) = \{q_1, \ldots, q_\ell\}.$$

v) We have

$$W = \frac{\pi^* H^0(C_2, \mathscr{O}_{\mathbb{P}^A}(1) \mid C_2)}{H^0(C_3, \omega_{C_3}^{\otimes n}(\varphi^* D))} \in \mathrm{Gr}_D \tag{6.8}$$

Proof. We follow standard proofs of similar statements, which rely on [Gro66] 12.2.1(x)- cf. for example [MFK94]. The first condition is evidently closed, as it can be obtained via an incidence correspondence. By imposing open conditions we may assume all curves are reduced, and that C_1 and C_3 are smooth. In addition, consider the closed condition given by requiring that we have $\rho(C_i) = C_3$, for $i = 1, 2$. More precisely, if $I_1, I_2 \trianglelefteq \mathbb{C}[X_0, \ldots X_A]$ are respectively the ideals of C_1 and C_2, and $I_3 \trianglelefteq \mathbb{C}[X_0, \ldots, X_a]$ is that of C_3, we require for $i = 1, 2$ that $I_i \cap \mathbb{C}[X_0, \ldots X_a] = I_2$. Now let $\mathbb{P}^{A-a-1} \subset \mathbb{P}^A$ be the indeterminacy locus of the linear projection ρ, given by the kernel of the projection $V_D \to {}^{V_D}\!/_{V}$. Since C_1 is smooth, every connected component must either be contained the linear space $\mathbb{P}^{A-a-1}$, or meet it in finitely many points. Denote by C_1' the union of the connected components of C_1 not contained entirely in this subspace. Then ρ defines a rational map between the smooth curves C_1' and C_3, which must therefore be isomorphic via a unique map $\varphi : C_3 \to C_1$ extending ρ. But if we restrict to the locus where C_1 and C_3 are such that every irreducible component has genus at least 2, this ensures that C_1' is in fact isomorphic to C_3, since they both have the same genus as C_2. So C_1 and C_3 are isomorphic via φ. Writing $Q_i = \varphi^{-1}(q_i)$, we may denote by $D_Q = \varphi^* D = N \sum_i Q_i$ the divisor obtained by pulling back $D = \sum_i q_i$ to C_1. By the usual methods, the conditions that C_3 is n-canonically embedded, and C_1 is embedded via $\omega_{C_1}^n(D_Q)$, are both closed. Since $\rho(C_1) = C_2$, we know that there is a rational map from C_1 to C_2, and since C_1 is smooth, it must in fact be a morphism. Adding in the condition that C_2 be integral (which is open since we are working over an algebraically closed field), we see that this map must be birational, and hence the morphism $\pi : C_1 \to C_2$ induced by ρ must be the normalisation. Finally, the condition on the Grassmannian is determined by the other data. □

From now on we will simplify notation as best we can, and write

$$[\underline{C}] = ([\widetilde{C}_Q], [C], [\widetilde{C}_q], \underline{q}, W) \in \mathcal{S}$$

for points of $\mathcal{S}$. Similarly to the sheaves case, with the Quot scheme, we will say that such a point corresponds to a curve C, since it is the singular curve we are really interested in. As in the proof above we take $Q_i = \varphi^{-1}(q_i)$, and write

$$D_Q = \varphi^* D_q = N \sum_i Q_i$$

for the divisor obtained by pulling back $D_q = N \sum_i q_i$. We also identify $\widetilde{C}_Q$ and $\widetilde{C}_q$ via the isomorphism given above. Note that, although the q_i and Q_i are the same considered as points on the abstract curves $\widetilde{C}_q$ and $\widetilde{C}_Q$ via this identification, they represent different points in their respective projective spaces. We will also abuse notation slightly and write

$$\omega_Q = \omega_{\widetilde{C}}^{\otimes n}(D_Q)$$

$$\omega_q = \omega_{\widetilde{C}}^{\otimes n}$$

for the line bundles on $\widetilde{C}$ obtained by restricting, respectively, the sheaves $\mathscr{O}_{\mathbb{P}^A}(1)$ and $\mathscr{O}_{\mathbb{P}^a}(1)$.

The scheme $\mathcal{S}$ evidently parameterises singular curves of a certain sort, and it is clear that there is some redundant information which we must remove to obtain a moduli space. Let us investigate exactly how much freedom there is, given an isomorphism class of singular curve C with the right Hilbert polynomial.

6.3.2.2 Notation. We will consider coordinate systems on $\mathbb{P}^A = \mathbb{P}(V_D^*)$ of the form

$$[X_0 : \cdots : X_r : Y_1 : \cdots : Y_\ell : Z_0 : \ldots Z_s],$$

$$\langle X_0, \ldots, X_r, \ldots, Y_1, \ldots, Y_\ell \rangle = V.$$

We fix, for each $i = 1, \ldots, \ell$, the points

$$q_i = [0 : \cdots : 0 : 1 : 0 : \cdots : 0] \in \mathbb{P}^a = \mathbb{P}(V^*)$$

to be where all coordinates vanish except for Y_i. In representing matrices in $SL(V_D)$ in block form we will separate X's, Y's and Z's by solid lines, for ease of parsing.

6.3.2.3 Definition. We consider the non-reductive group $H \triangleleft \mathrm{SL}(V_D)$, given by

$$H \coloneqq \left\{ \left(\begin{array}{c|c|c} A & 0 & 0 \\ \hline U_1 & B & U_2 \\ \hline 0 & 0 & C \end{array} \right) \in H_1 \rtimes \mathbb{G} \right\} \cap SL(V_D)$$

where

i) H_1 is the group of matrices of the form

$$\left(\begin{array}{c|c|c} A_1 & 0 & 0 \\ \hline U_1 & B_1 & 0 \\ \hline 0 & 0 & I \end{array} \right)$$

whose top left 2×2 block is the subgroup of $SL(V)$ that fixes the points

$$q_i \in \mathbb{P}^a = \mathbb{P}(V^*),$$

so that B_1 is diagonal.

ii) We define

$$\mathbb{G} \coloneqq \prod_{i=1}^{\ell} \mathbb{G}_{N,i},$$

where each $\mathbb{G}_{N,i}$ is a re-parameterisation group as defined at Definition 6.2.2.5.

iii) The isomorphism $H/H_1 \cong \mathbb{G}$, is via the identification (6.9) below.

6.3.2.4 Proposition. *The scheme $\mathcal{S}$ carries an action of H, and isomorphism classes of connected unibranch curves C of genus $g = \tilde{g} + \delta$, with a choice of ordering on their singular points, are in one-to-one correspondence with orbits of closed points under this group action.*

Proof. Since elements of H act naturally on V_D and fix V, they induce a diagonal action on $\mathcal{H}_1 \times \mathcal{H}_2 \times \mathcal{H}_3 \times (\mathbb{P}^a)^\ell$. The action on the Grassmannian factor comes about in the following way. For any point of $\mathcal{S}$ we can canonically identify

$$V_D/V \cong \frac{H^0(\widetilde{C}, \omega_{\widetilde{C}}^{\otimes N}(D))}{H^0(\widetilde{C}, \omega_{\widetilde{C}}^{\otimes N})} \cong \mathscr{O}_{\widetilde{C},q_i}/\mathfrak{m}_{q_i}^N$$

and hence we may also make a choice of identification

$$V_D/V \cong \frac{H^0(\widetilde{C}, \omega_{\widetilde{C}}^{\otimes N}(D))}{H^0(\widetilde{C}, \omega_{\widetilde{C}}^{\otimes N})} \simeq \bigoplus_{i=1}^{\ell} \mathbb{C}[t_i]/(t_i^N). \tag{6.9}$$

Any W coming from a point of $\mathcal{S}$ as described at (6.8) gives a direct sum of subrings, and hence we can think of this choice up to the action of the product of re-parameterisation groups above. Hence, fix such an identification and thus the Grassmannian acquires a well-defined $\mathbb{G}$ action, which induces a H action. The full action of H on $\mathcal{S}$ is given by taking the diagonal action on the full product.

Suppose we have points $[\underline{C}_1], [\underline{C}_2] \in \mathcal{S}$, where

$$[\underline{C}_j] = ([\widetilde{C}_Q^j], [C^j], [\widetilde{C}_q^j], \underline{q}, W^j),$$

and suppose $C^1 \cong C^2$, as pointed curves. Then it follows that $\widetilde{C}_q^1 \cong \widetilde{C}_q^1$ and $\widetilde{C}_Q^1 \cong \widetilde{C}_Q^1$ are isomorphic abstract curves, embedded in each of $\mathbb{P}^a$ and $\mathbb{P}^A$, respectively, by the same line bundle. So there is a matrix $M_1 \in \mathrm{SL}(V)$, which respects the inclusion $W \subset V$, taking one to the other. Moreover this matrix must fix the points q_i on the C_q^j's, since these points are distinguished.
Next, choose local parameters on $\widetilde{C}_q^1$ at each q_i, and pull them back to local parameters on $\widetilde{C}_Q^1$. Doing the same for $\widetilde{C}_q^2$ and $\widetilde{C}_Q^2$, we see that the two choices must be related, up to order N by the action of an element of $\mathbb{G}$; the remainder can be dealt with by the rest of H. Equivalently, we can think of these local parameters as giving us a choice of complement to V_D in V, which by construction is well-defined up to the action of $\mathbb{G}$. □

By the above proposition, the naive answer to our moduli problem should be the quotient of $\mathcal{S}$ by H, and as usual we want to consider the unstable strata picture associated to this GIT problem. Though H is not reductive, this still makes perfectly good sense. We will see that we are in a setting in which Hilbert-Mumford can be shown to apply, and so it is still reasonable to look for maximally destabilising one-parameter subgroups, as in Section 2.3. We do this, as in the reductive case, by choosing a maximal torus, and then analysing how the convex hull of the torus weights changes as we sweep out by the action of H. Hence, the procedure outlined in Chapter 4 can be followed: if we can interpret the instability strata in terms of the geometry of our curves, and find categorical quotients of those instability strata, then the above proposition tells us that we will have coarse moduli spaces for whichever curves those strata parameterise.

To complete this set-up as a non-reductive GIT problem, we need our group to act linearly on a projective variety. We recall that, for $m \gg 0$, a Hilbert scheme of curves has an embedding into projective space, as in Section 6.1. Hence, fix m sufficiently large, depending on the parameters chosen so far, so as to work uniformly for $\mathcal{H}_1$, for $i = 1, 2$, and denote the associated very ample line bundle on $\mathcal{H}_i$ by $\Lambda_{m,i}$. We choose $m_1, \ldots, m_\ell$ as in Section 6.1.2, so that $L = L(m, m_1, \ldots, m_\ell)$ is the Baldwin-Swinarski linearisation on $\mathcal{H}_3 \times (\mathbb{P}^a)^\ell$. We recall that there is also a Plücker embedding of our Grassmannian Gr_D, and denote the corresponding very ample line bundle by L_{Gr}. Putting all of this together, we define a line bundle on $\mathscr{H}$ as

$$\mathscr{L} = \mathscr{L}(M, m, m_1, \ldots, m_\ell) \coloneqq \Lambda_{m,1} \otimes \Lambda_{m,2} \otimes (L(m, m_1, \ldots, m_\ell))^{\otimes M} \otimes L_{\mathrm{Gr}}.$$

This line bundle defines an embedding of the product $\mathscr{H}$ into a product of projective spaces, and is hence very ample. Since each of the factors in the product has a natural linearisation of the H action, so does the product itself.

6.3.2.5 Remark. We will be thinking of M as very large compared to all of the other parameters we have defined so far- the idea behind this, roughly speaking, is that in order for a curve C to be stable with respect to this action, its normalisation, considered as a curve with marked points, must be Baldwin-Swinarski stable. Despite the different parameter space and non-reductive group, we are working, in effect, with a small perturbation of the Baldwin-Swinarski GIT set-up.

6.3.2.6 Notation. To denote a particular choice of our parameters, we will write

$$\underline{m} \coloneqq (M, m, m_1, \ldots, m_\ell).$$

6.3.3 The associated graded curve

We can now begin to study the instability stratification with respect to the GIT problem set up above. Motivated by the case of sheaves, we are looking for something like a Harder-Narasimhan filtration for curves. This of course will correspond to a weighted filtration of V_D - that is, a one-parameter subgroup - but there is no meaningful sense in which we can be said to have an induced filtration by subcurves. However, we can still make a good guess at what the 'associated graded curve' should be.

6.3.3.1 Definition.

1) Let C be a unibranch curve having singularities at $\{p_1, \ldots, p_\ell\}$, of genera δ_i. At each $p_i \in C$, let Γ_i be the semigroup, let c_i be the conductor of Γ_i, and choose generators $\{\gamma_i^j\} \subset \Gamma_i$. Define

$$N_i = \max(\{c_i\} \cup \{\gamma_i^j\})$$

and take $N = \max\{N_i\}$. For each i, define $R_i(C)$ to be the rational curve of degree N parameterised by

$$\mathbb{P}^1 \to \mathbb{P}^{N-\delta_i}$$

$$[s:t] \longmapsto [s^N : s^{N-v_1}t^{v_1} : s^{N-v_2}t^{v_2} \ldots : t^N]$$

where the v_i are the elements of $\Gamma_i \cap \mathbb{N}_{\leq N}$.

2) We define the *associated graded curve* of C to be the curve obtained from the normalisation, $\widetilde{C}$, by attaching the rational curves R_i to the points p_i such that the R_i are all disjoint and $\widetilde{C}$ meets each $R_i(C)$ at p_i, at a node. We denote this curve by $\overline{C}$.

6.3.3.2 Remark.

i) The rationale behind this definition, is that it separates the information about C that is irrelevant to its stability, i.e. the isomorphism class of the normalisation, from the relevant information: the singularity data. These are the curves which we will want to lie in the analogue of Z_β^{ss} for our instability stratification. Hence, in some sense, we want to think of the graded curve $\overline{C}$ as the curve that loses the least amount of information about C, whilst still having a $\mathbb{G}_m$ in its automorphism group that 'knows' at least something about the singularities.[9]

ii) A key feature of the unibranch case is that, associated to each unibranch semigroup there is a distinguished projective rational curve having that semigroup-namely the one above, the monomial curve. This reflects the fact that, at the local level, the unibranch re-parameterisation group only has a rank one torus. If the singularities of C are multibranch, we can still define the associated graded, but the difference is that we will have to make a choice of $\mathbb{G}_m$ inside the multibranch re-parameterisation group of the singularity. The natural thing to do

[9] Just as in the sheaves case we lose information by passing to the associated graded.

in that case would be to take the $\mathbb{G}_m$ giving the maximally destabilising one parameter subgroup in the Rosenlicht-Serre Grassmannian.[10]

Shortly we will exhibit a one-parameter subgroup that degenerates points in $\mathcal{S}$ corresponding to a singular curve C into its associated graded. We can now define the one-parameter subgroup that will play the role that $\beta_{m,n}$ did in Chapter 5.

6.3.3.3 Definition. Fix a given choice of coordinates as in Notation 6.3.2.2. Define a one parameter subgroup diagonalised by these coordinates $\theta(t) : \mathbb{G}_m \to H$ via

$$\theta(t) = \ \mathrm{diag}(1, 1, \ldots, 1, 1, t^{-1}, t^{-2}, \ldots, t^{2-N}, t^{1-N}, 1, t^{-1}, \ldots, t^{1-N}, 1, \ldots, \ldots, t, 1)$$

That is to say, all coordinates belonging to V have weight 0, and the remainder is given by taking the diagonal in a chosen maximal torus of the product of re-parameterisation groups.

6.3.3.4 Proposition. *Consider the point in the parameter scheme*

$$[\underline{C}] = ([\widetilde{C}_Q], [C], [\widetilde{C}_q], \underline{q}, W) \in \mathcal{S}.$$

Then we have

$$\lim_{t\to 0} \theta(t)[\underline{C}] = ([\widetilde{C}^\circ], [\overline{C}], \widetilde{C}_q, \underline{q}, \overline{W}) \subset \overline{\mathcal{S}} \subset \mathscr{H}$$

where

i) $\widetilde{C}^\circ$ is a curve obtained from $C_q \subset \mathbb{P}^a \subset \mathbb{P}^A$ by attaching ℓ copies of $\mathbb{P}^1$ of degree N, passing through Q_i and meeting $\widetilde{C}_q$ in a node at q_i. Each of these curves is embedded via a complete linear system $\mathscr{O}_{\mathbb{P}^A}(N)$.

ii) The curve $\overline{C}$ is the associated graded of C, embedded via taking the linear projection of $\widetilde{C}$ onto the span of C. That is, $\overline{C} \subset \mathbb{P}^A$ is obtained from $\widetilde{C}_Q \subset \mathbb{P}^A$ by attaching the curves $R_i(C)$, each of degree N, with the singularity at Q_i, and meeting $\widetilde{C}_Q$ at a node at q_i

iii) The linear space $\overline{W}$ in Gr_D is obtained as in the proof of Proposition 6.2.3.6.

Proof. The proof of both (*i*) and (*ii*) in the case $\ell = 1$ is the same, mutatis mutandis, as that of Proposition 6.3.1.8, while (*iii*) is proved at Proposition 6.2.3.6. The generalisation to more than one singularity, except for the addition of notational clutter, is no harder. □

[10] cf. Chapter 3 of [BRV17] and in particular theorem 3.5.

6.3.3.5 Remark. If all singularities of C are local complete intersections, we can prove this proposition in a more straightforward way, which also generalises to the multibranch case without increased difficulty. Following [HH13], we can consider the weights of λ on the deformation spaces of the singularities of $\overline{C}$. These singularities are of two kinds: those coming from C, and those not coming from C, of which the latter are all nodes. If the curve C is a local complete intersection then the Hilbert scheme H_2 is smooth at C, so we have smooth point of $\mathcal{S}$. Hence, by analysing the weights of λ we can show that the basin of attraction of $\underline{\overline{C}} \in \mathcal{S}$ contains no smoothings of the singularities coming from C, and all smoothings of those singularities not coming from C.

6.3.3.6 Remark. It is also worth noting a dissimilarity with the sheaves case: we have, for all unibranch curves with these discrete invariants, the same Kempf one parameter subgroup (although we will have a different weight). This is an artefact of the way we have set things up, which is not likely to persist in the multibranch case.

6.3.4 Construction of $M_{g,\sigma}$

We now turn our attention to the construction of the moduli space of unibranch curves of fixed Rosenlicht-Serre type.

6.3.4.1 Theorem. *Fix a unibranch Rosenlicht-Serre type $\sigma = (\Sigma_1, \ldots, \Sigma_\ell)$. Then there is a quasi-projective coarse moduli space $M_{g,\sigma}$ parameterising irreducible projective algebraic curves of Rosenlicht-Serre type σ of arithmetic genus g up to isomorphism, which admits a projective completion $\overline{M}_{g,\sigma}$ constructed via non-reductive GIT.*

We will prove this proposition in stages, and first we need some preliminaries. Proposition 6.3.3.4 allows us to formulate the following definition, motivated by Section 4.2. This is the linearisation that will play the role played by $\mathscr{L}^{per}_{\beta(\tau)}$ in the sheaves case. Note, though, that unlike the sheaves case, the one-parameter subgroup $\theta : \mathbb{G}_m \to H$ does not depend on the Rosenlicht-Serre type, only on the discrete invariants necessary to set up the parameter space. The weight of its action, however, does depend on the Rosenlicht-Serre type, and not on any other data of the curve.[11] that is to say: if C is a curve of Rosenlicht-Serre type σ, and $[\underline{C}] \in \mathcal{S}$ a corresponding point, then the value of the Hilbert-Mumford function $\mu^{\mathscr{L}}([\underline{C}], \theta)$ depends only on the Rosenlicht-Serre type of C, by the above.

[11] In fact, one can easily calculate it explicitly, using (6.2).

6.3.4.2 Definition.

i) Fix a Rosenlicht-Serre type σ. Define the character $\chi_\sigma = \chi_\sigma(\underline{M}) : \mathbb{G}_m \to H$, to be that given by the weight of the action of θ on the associated graded curves of any curve of Rosenlicht-Serre type σ.

ii) Denote by $\mathscr{L}_\sigma = \mathscr{L}_\sigma(\underline{M})$ the linearisation given by twisting the linearisation $\mathscr{L}$ by the character $-\chi_\sigma$. Similarly, denote by $\mathscr{L}_\sigma^{per}$ the linearisation given by twisting $\mathscr{L}$ by $-\chi_s(1+\varepsilon)$, for some $\varepsilon > 0$ chosen sufficiently small.

Let us now take the closure $\overline{\mathcal{S}} \subset \mathcal{H}$, which is then a projective variety. We will quotient by the action of H in three stages, using techniques from both Sections 3.1 and 3.3. First, we will quotient by the action of the group

$$H_1' \coloneqq \left\{ \left(\begin{array}{c|c|c} A_1' & 0 & 0 \\ \hline U_1' & B_1' & 0 \\ \hline 0 & 0 & I \end{array} \right) \in H_1 \trianglelefteq H \mid \det A_1' \det B_1' = 1 \right\}$$

Recalling H_1 from Definition 6.3.2.3, this means that B_1' is diagonal, and the whole 2×2 block lies in $\mathrm{SL}(V)$. Let us write $G_1 \coloneqq \mathrm{SL}(V)$, and consider the locus

$$\mathcal{T} \coloneqq G_1 \mathcal{S}.$$

Then taking the closure in the ambient projective space yields $\overline{\mathcal{T}}$, a projective variety. Further G_1 is a reductive group, and our linearisation $\mathcal{L}_\sigma^{per}$ extends to a linearisation of the G_1 action on $\mathcal{T}$: this is because the line bundle itself is defined and very ample on the whole of $\mathcal{H}$, and the character χ_σ extends to a character of G_1 (the trivial one).

6.3.4.3 Lemma. *We have a G_1-equivariant isomorphism $\mathcal{T} \cong G_1 \times^{H'1} \mathcal{S}$, where the latter is as defined at Definition 3.1.3.1.*

Proof. The scheme $\mathcal{S}$ is H_1-invariant, so there is a morphism $G_1 \times^{H_1} \mathcal{S} \to \mathcal{T}$ with image T. Now let $[\underline{C}] = ([C_1], [C_2], [C_3], \underline{r}, W) \in \mathcal{T}$. There is a projection

$$f : \mathcal{T} \to (\mathbb{P}^a)^\ell,$$

under which $\mathcal{S} = f^{-1}(\{q_1, \ldots, q_\ell\})$. And

$$G_1 \cdot \{q_1, \ldots, q_\ell\} \cong G \times^{H_1} \{q_1, \ldots, q_\ell\},$$

so we have

$$\mathcal{T} = G_1\mathcal{S} = G \cdot f^{-1}(\{q_1, \ldots, q_\ell\}) \cong G \times^{H_1} f^{-1}(\{q_1, \ldots, q_\ell\}) = G_1 \times^{H_1'} \mathcal{S}.$$

□

It follows from the above that a quotient of $\mathcal{T}$ by G_1 is equivalent to a quotient of $\mathcal{S}$ by H_1'. Using classical GIT there is a categorical quotient

$$\varphi_{\overline{\mathcal{T}}} : \overline{\mathcal{T}}^{ss} \to \overline{\mathcal{T}} /\!\!/ G_1 =: \mathcal{X}$$

which is a projective variety, and the semistable locus is determined by the Hilbert-Mumford criterion. Since we have chosen $M \gg 0$ and very large compared to all the other parameters, and since semistability and stability coincide (in the usual reductive GIT sense) for the Baldwin-Swinarski quotient, together with the Hilbert-Mumford criterion this implies that we have $\mathcal{T} \subset \overline{\mathcal{T}}^s$. Hence, we have a quasi-projective geometric quotient

$$\mathcal{S} \to \mathcal{S}/H_1,$$

with a natural projective completion provided by $\mathcal{X}$. This scheme carries an induced action of the residual group, which we can write as

$$H_2' \coloneqq H/H_1' \cong U_1 \rtimes L_1,$$

with U_1 being the unipotent radical of the re-parameterisation group, and L_1 being reductive. This action has a linearisation induced from $\mathscr{L}_\sigma^{per}$ on $\mathcal{H}$, which we will denote by $\overline{\mathscr{L}_\sigma^{per}}$. Moreover, we are now in the situation where the $\hat{U}$ Theorems can be applied, since we have a one parameter subgroup $\theta : \mathbb{G}_m \to H \to H/H_1$, which has positive weights on the unipotent radical U_1. In fact we can use Theorem 3.3.1.8.

6.3.4.4 Lemma. *The subset* $\{x \in \mathcal{X} \mid \mathrm{Stab}_{U_1}(x) = \{e\}\}$ *is dense.*

Proof. This essentially follows because of the finiteness of automorphism groups of irreducible curves of genus at least 2. A full analysis of when curves may have infinite automorphism groups, and which automorphism groups they may have, may be found in [Pop78]. We first argue that it suffices to prove the statement for $\mathcal{S}$. Note that the image

$$\varphi : \mathcal{S} \to \mathcal{S}/H_1 \subset \mathcal{X}$$

is dense in $\mathcal{X}$. Take $[\underline{C}] \in \mathcal{S}$. Since the quotient is geometric and H-equivariant, if

$u \in U_1$ fixes $\varphi([\underline{C})]$ then there is $h \in H_1'$ such that $h[\underline{C}] = u[\underline{C}]$. Since the intersection $H_1' \cap U_1$ is trivial, it follows that if the stabiliser of a point in U_1 on one side is infinite, it must be infinite on the other side also. So we can prove the statement by proving it for $\mathcal{S}$.

Now, for $[\underline{C}] \in \mathcal{S}$, any element of $\mathrm{Stab}_{U_1}([\underline{C}])$ must in particular lie in the stabiliser of the Hilbert point $[\widetilde{C}_Q] \in \mathcal{H}_1$. By [BFMV14] Lemma 6.1, recalling that $[\widetilde{C}_Q]$ is n-canonical, the stabiliser in U_1 of this Hilbert point is a unipotent subgroup of the automorphism group of the normalisation $\widetilde{C}$ of C, considered as an abstract curve. A non-trivial unipotent group must have a subgroup isomorphic to $\mathbb{G}_a$, but since $\widetilde{C}$ is smooth with all connected components having genus at least 2, it follows that we have

$$\mathrm{Stab}_{U_1}([\underline{C}]) = \{e\} \ \forall [\underline{C}] \in \mathcal{S}.$$

This proves the lemma. □

We now restrict our attention to the Rosenlicht-Serre type we are interested in.

6.3.4.5 Definition. We define the σ-level

$$Y_\sigma^0 = \{\underline{[C]} = ([\widetilde{C}_Q], [C], [\widetilde{C}_q], \underline{q}, W) \in \mathcal{S} \mid C \text{ has Rosenlicht-Serre type } \sigma\}.$$

Similarly, define

$$Z_\sigma = \{\lim_{t\to 0} \theta(t)\underline{[C]} \mid \ \underline{[C]} \in Y_\sigma\},$$

the associated graded curves.

6.3.4.6 Lemma. *The level $Y_\sigma^0 \subset \mathcal{S}$ is an H-invariant locally closed subscheme, as is $Z_\sigma \subset \mathcal{S}$.*

Proof. First observe that the H-invariance is clear. Next, we can endow Y_σ^0 with the structure of a locally closed subscheme by defining it as a scheme to be the fibre product

$$\begin{array}{ccc} Y_\sigma^0 & \longrightarrow & \mathrm{Gr}_D^\sigma \\ \downarrow & & \downarrow \\ \mathcal{S} & \longrightarrow & \mathrm{Gr}_D \end{array} \tag{6.10}$$

Here the bottom arrow is the projection onto the last factor. The vertical arrow on the right is the inclusion, of those subrings of Rosenlicht-Serre type σ, which by [Ish80], gives Gr_D^σ the structure of a locally closed subscheme of Gr_D. And since locally closed immersions are stable under arbitrary base change, we are done. A similar proof works for Z_σ. □

Hence, taking projective closure, we have $\overline{Y^0_\sigma} \subset \mathscr{H}$, a projective variety.

6.3.4.7 Definition. Define $\mathcal{Y}^0_\sigma = \varphi(Y^0_\sigma)$.[12] There is an induced map

$$p_\sigma : \mathcal{Y}^0_\sigma \to \mathcal{Z}_\sigma =: \varphi(Z_\sigma)$$

which takes a point to its limit in the minimal weight space; i.e. to the point representing its associated graded curve.

6.3.4.8 Proposition.

1) The scheme $\mathcal{Y}^0_\sigma$ has a geometric quotient by the action of $\hat{U}_1 \coloneqq U_1 \rtimes_\theta \mathbb{G}_m$.

2) Following this up with the induced reductive quotient by $L_1 \cong H'_2/U_1$, there is a quasi-projective geometric quotient

$$(\mathcal{Y}^0_\sigma)^s \to \mathcal{W}_\sigma.$$

Proof. For this we use the $\hat{U}$ Theorems. Embed $\mathcal{Y}^0_\sigma \subset \overline{\mathcal{X}}_\sigma$ in its projective closure, and consider the action of H'_2 on it. With respect to the restriction of the H'_2 linearisation $\overline{\mathscr{L}^{per}_\sigma}$, by construction the distinguished $\theta : \mathbb{G}_m \to H \to H'_2$ has precisely one strictly negative weight, that corresponding to the image of Z_σ, the associated graded curves, and the rest of its weights are all strictly positive. By Lemma 6.3.4.4 and its proof, which applies equally to $\mathcal{Y}^0_\sigma$, the generic stabiliser in U_1 of points in $\overline{\mathcal{Y}^0_\sigma}$ are trivial. So we can perform the blow-up process on $\overline{\mathcal{Y}^0_\sigma}$, terminating in a projective variety

$$\psi : \widetilde{\mathcal{Y}^0_\sigma} \to \overline{\mathcal{Y}^0_\sigma}$$

such that all U_1-stabilisers are trivial. Moreover, the U_1-stabilisers of points in $\mathcal{Y}^0_\sigma$, and their images under the map p_σ, are all trivial in the first place. So these points all survive the process without being blown up, and we can think of

$$\mathcal{Y}^0_\sigma \subset \widetilde{\mathcal{Y}^0_\sigma}$$

via the identification provided by ψ. This proves the first statement- the second then follows from Theorem 3.3.1.8, as the residual group L_1 is reductive.

□

[12] This has a canonical scheme structure by standard properties of reductive GIT quotients, as a geometric quotient of Y^0_σ by the H'_1 action: see [New78] 3.4.3 and the universality statement in Theorem 1.10 of [MFK94].

Conclusion of the proof of Theorem 6.3.4.1. Let us now put all of this together. By composing our two quotient maps

$$Y_\sigma^0 \to \mathcal{Y}_\sigma^0 \dashrightarrow \mathcal{W}_\sigma,$$

we have established that some open subset of Y_σ^0 must have a geometric quotient by the action of H. In fact, this subset is all of Y_σ^0: we can understand this in terms very similar to Remark 2.3.0.12. Taking the limit of a point of Y_σ^0 under the one parameter subgroup θ takes it to the corresponding point of Z_σ^0, its associated graded curve. The group that plays the role of $\operatorname{Stab}\beta$ is then H_1', and the statement then reduces to Theorem 6.1.2.2. □

6.3.5 Remarks on generalisations

It will not have escaped the reader's notice that no attempt was made to keep track of what appears in the boundary of this construction. To that extend, the proof resembles Chapter 5 of [MFK94]. A priori, they might be points coming from blow ups of Hilbert schemes, or at least curves in the original product of Hilbert schemes that are non-reduced. However, it is likely that a thorough analysis of the boundary could be performed without a great deal of work. To do this, rather than working with a parameter space that places a singular curve between two copies of its normalisation, it would be preferable in some respects to make use of the full force of the Baldwin-Swinarski construction, essentially working with a parameter space of curves together with their *nodalisations*. This done, one could ask the question of whether the resultant moduli space is in fact projective, by analysing the fibre of the natural map to $\overline{M}_{g,n}$.

Another natural area for generalisation is the extension to non-nodal multibranch singularities, and most of what we have said carries over to that case with very minimal alteration. As indicated in Remark 6.3.3.2, a definition of the Rosenlicht-Serre type, and associated graded curve, is possible, and, as stated in Remark 6.3.3.5, an analogue of Proposition 6.3.3.4 can presumably be shown to hold.

Appendix

Diagrams and Algorithm for Projective Completions

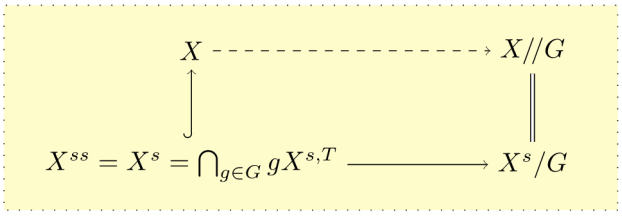
by Eloise Hamilton

1 Diagrams for GIT, reductive and non-reductive

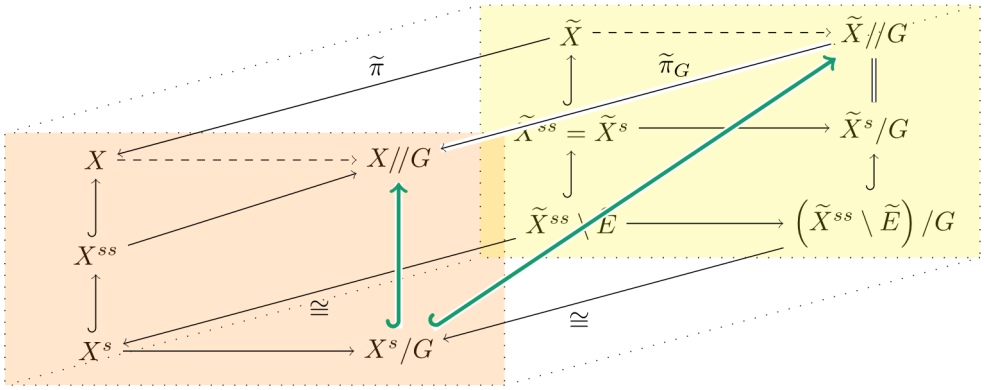
In this section we summarise the main results from Geometric Invariant Theory (GIT), both in the classical setting of [MFK94] (Section 1.1) and in the non-reductive setting of [BDHK16] (Sections 1.2 and 1.3), providing diagrams to illustrate the constructions involved.

1.1 Classical GIT for reductive groups

Let X be an irreducible projective variety over an algebraically closed field of characteristic 0 and let G be a reductive group acting linearly on X with respect to an ample line bundle $L \to X$. Then according to classical GIT, there exist open subsets $X^s \subseteq X^{ss} \subseteq X$ and a projective variety $X/\!/G = \operatorname{Proj} \bigoplus_{k\geq 0} H^0(X, L^{\otimes k})^G$ with a rational map $q_G : X \dashrightarrow X/\!/G$ such that q_G restricts to a good quotient $X^{ss} \to X/\!/G$ and induces a geometric quotient $X^s \to q_G(X^s) = X^s/G$. Note that in the terminology of [MFK94], X^s is the properly stable locus; we will use the notation X^{Ms} to denote the stable locus in Mumford's sense. When X^s is dense in X, the GIT quotient $X/\!/G$ can be viewed as a projective completion of the geometric quotient X^s/G. Moreover,



(a) When $X^{ss} = X^s \neq \emptyset$



(b) When $\emptyset \neq X^s \subsetneq X^{ss}$. The green arrows denote projective completions.

Figure .1: Classical GIT for reductive groups

the semistable and stable loci are characterised by Hilbert-Mumford criteria so that

$$X^{(s)s} = \bigcap_{g\in G} gX^{(s)s,T}$$

for a fixed maximal torus $T \subseteq G$. In what follows, if X^s is empty but X^{Ms} is not, then stability should be replaced by Mumford-stability. Note however that the Mumford-stable locus, unlike the stable locus, is not described by a Hilbert-Mumford-type criterion.

If $\emptyset \neq X^s = X^{ss}$, then the GIT quotient $X//G$ is a projective geometric quotient for the stable locus of X, and hence can have at worst only finite quotient singularities. Figure .1a illustrates this situation.

If $\emptyset \neq X^s \subsetneq X^{ss}$, then by [Kir85, Rei89] there exists a sequence of blow-ups starting from X and resulting in a projective variety $\widetilde{X}$ with a blow-down map $\widetilde{\pi} : \widetilde{X} \to X$ such that $\widetilde{X}^{ss} = \widetilde{X}^s \neq \emptyset$ for a suitably linearised action of G on $\widetilde{X}$ (see [BDHK16, §3] for more detail). By construction of the blow-ups, $X^s = \widetilde{\pi}(\widetilde{X}^s \setminus \widetilde{E})$ where $\widetilde{E}$ denotes the exceptional divisor of the blow-down map $\widetilde{\pi}$. Figure .1b illustrates this blowing up construction. The two green arrows denote the two projective completions of X^s/G, namely $X//G$ and $\widetilde{X}//G$, where the latter can be viewed as a partial desingularisation

of the former.

If $\emptyset = X^s = X^{Ms} \neq X^{ss}$, then although the blow-ups can again be performed, for the resulting variety $\widetilde{X}$ the semistable locus is empty (see [BDHK16, Remark 3.3]). In general, if the semistable locus for the linearised action of a reductive group G on a projective variety X is empty, then we are not in a situation where classical GIT can be usefully applied. This is one of the motivations for considering non-reductive GIT. Indeed, given an invariant inner product on the Lie algebra of G, by the instability-stratification for X we have:

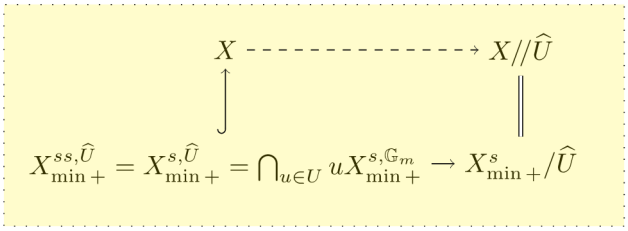
$$X = \bigsqcup_{\beta \in \mathcal{B}} S_\beta$$

where $S_0 = X^{ss}$ and $S_\beta = G \times_{P_\beta} Y_\beta^{ss}$ for $\beta \neq 0$ (see [BDHK16, §8.2] and [BHK17, §1] for more details). Note that when G is connected, S_β is irreducible if and only if Y_β^{ss} is irreducible. If $X^{ss} = \emptyset$, then there must exist an unstable stratum S_β with $\beta \neq 0$ such that $S_\beta \cong GY_\beta^{ss} = G \times_{P_\beta} Y_\beta^{ss}$ is open in X. Taking a quotient of an invariant open subset of S_β by G is thus equivalent to taking a quotient of an invariant open subset of Y_β^{ss} by the parabolic subgroup P_β, which is non-reductive in general. However, the subgroup P_β satisfies the property that its unipotent radical is internally graded. That is, $P_\beta = U_\beta \rtimes L_\beta$, where U_β is its unipotent radical and L_β a Levi subgroup, and there exists a central one-parameter subgroup $\lambda_\beta : \mathbb{G}_m \to L_\beta$ such that $\lambda_\beta(\mathbb{G}_m)$ acts on $\operatorname{Lie} U_\beta$ with strictly positive weights. In the following two subsections we summarise results from [BDHK16] for graded extensions of unipotent groups (Section 1.2) and for non-reductive groups with internally graded unipotent radicals (Section 1.3), providing diagrams to illustrate them.

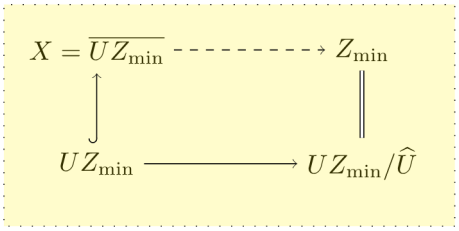
1.2 GIT for graded extensions of unipotent groups

Let $\widehat{U} = U \rtimes \mathbb{G}_m$ be the semi-direct product of a non-trivial unipotent group with the multiplicative group $\mathbb{G}_m$ such that $\mathbb{G}_m$ acts on $\operatorname{Lie} U$ with strictly positive weights. Suppose that $\widehat{U}$ acts linearly on a projective variety X with ample line bundle L. The subvarieties $Z_{\min}$ and $X^0_{\min}$ are defined as in [BDHK16, Definition 1.1].

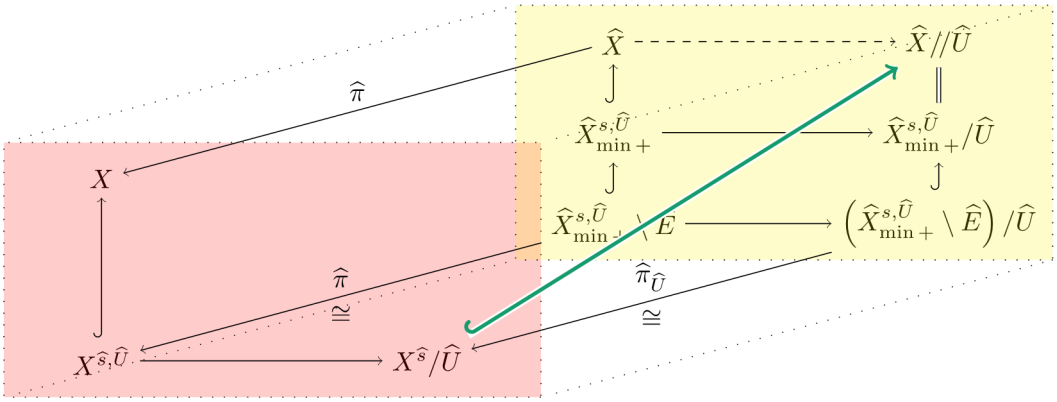
Since $\widehat{U}$ is non-reductive, the ring of invariants $\bigoplus_{k \geq 0} H^0(X, L^{\otimes k})^{\widehat{U}}$ may not be finitely generated and thus there is no obvious analogue of the GIT quotient. Instead, an 'enveloping quotient' $X \mathbin{/\!\!/} \widehat{U}$, can be defined, admitting a rational map $q_{\widehat{U}} : X \dashrightarrow X \mathbin{/\!\!/} \widehat{U}$ which restricts to a morphism on a suitably defined semistable locus $X^{ss,\widehat{U}}$. Unlike in the classical case, this morphism may not be surjective, and its image may



(a) When (ss=s($\widehat{U}$)) holds and $X \neq \overline{UZ_{\min}}$



(b) When (ss=s($\widehat{U}$)) holds and $X = \overline{UZ_{\min}}$



(c) When (ss=s($\widehat{U}$)) does not hold. The green arrow denotes a projective completion.

Figure .2: GIT for graded extensions of unipotent groups

not even be a subvariety of $X \wr \widehat{U}$ but only a constructible subset. Nevertheless, a stable locus $X^{s,\widehat{U}}$ can be defined so that $q_{\widehat{U}}(X^{s,\widehat{U}})$ is a subvariety of $X \wr \widehat{U}$ and a geometric quotient for the action of $\widehat{U}$ on X (see [BDHK16, §2] for more detail). By contrast with classical GIT, the semistable and stable loci for the $\widehat{U}$ action cannot in general be described by Hilbert-Mumford-type criteria.

Nevertheless, if the action satisfies an additional condition denoted (ss=s($\widehat{U}$)), namely that $\mathrm{Stab}_U(z) = \{e\}$ for all $z \in Z_{\min}$, then by [BDHK16, Theorem 1.2], after taking a tensor power of the linearisation and twisting it by an appropriate character, all of the properties of classical GIT in the case when semistability coincides with stability can be recovered. That is, the invariants form a finitely generated graded ring and $X \wr \widehat{U}$ is the associated projective variety. Moreover, a stable locus

$$X^{s,\widehat{U}}_{\min+} := X^{0}_{\min} \setminus UZ_{\min}$$

can be defined such that, if non-empty, $q_{\widehat{U}} : X^{s,\widehat{U}}_{\min+} \to X \wr \widehat{U}$ is surjective and a geometric quotient for the action of $\widehat{U}$ on $X^{s,\widehat{U}}_{\min+}$. Finally, the stable locus $X^{s,\widehat{U}}_{\min+}$ is

characterised by a Hilbert-Mumford-type criterion:

$$X^{s,\widehat{U}}_{\min+} = \bigcap_{u\in U} uX^{s,\mathbb{G}_m}_{\min+}$$

where $X^{s,\mathbb{G}_m}_{\min+}$ is the stable locus in the reductive sense for the action of $\lambda(\mathbb{G}_m)$ on X with respect to the tensored and twisted linearisation. Since the enveloping quotient $X \wr\!\wr \widehat{U}$ satisfies all of the properties of classical GIT in the case where semistability coincides with stability, we denote it by $X//\widehat{U}$ instead. Figure .2a illustrates this situation. Note that if $X^{s,\widehat{U}}_{\min+} = \emptyset$, then $UZ_{\min}$ must be open in X. In this case, $Z_{\min}$ is a geometric quotient for the action of $\widehat{U}$ on $UZ_{\min}$ (see [BDHK16, Remark 1.3]). Figure .2b illustrates this situation.

If $\mathrm{Stab}_U(x)$ is not trivial but of constant and strictly positive dimension for all $x \in X^0_{\min}$, and the same is true for subgroups $U^{(j)}$ appearing in a series $U \supseteq U^{(1)} \supseteq \cdots \supseteq U^{(s)} = \{e\}$ normalised by H with $U^{(j)}/U^{(j+1)}$ abelian, then the same conclusions as stated in the above paragraph hold (see [BDHK16, Remark 1.6]). This situation is analogous to the reductive case when $X^s = \emptyset$ but $\emptyset \neq X^{Ms} = X^{ss}$. For this reason we denote the above weaker condition for the $\widehat{U}$-stabiliser groups in $X^0_{\min}$ by (ss=Ms($\widehat{U}$)).

If the condition (ss=s($\widehat{U}$)) is not satisfied for the action of $\widehat{U}$ on X, then provided there is some $x \in X$ such that $\mathrm{Stab}_U(x) = \{e\}$ (this condition is analogous to the condition that $X^s \neq \emptyset$ in the reductive case), by [BDHK16, Theorem 1.8] there is a sequence of $\widehat{U}$-equivariant blow-ups of X resulting in a projective variety $\widehat{X}$ with a blow-down map $\widehat{\pi} : \widehat{X} \to X$ such that for a suitably defined linearised action of $\widehat{U}$ on $\widehat{X}$ the condition (ss=s($\widehat{U}$)) is satisfied. Figure .2c illustrates the blowing up procedure. If $\mathrm{Stab}_U(x)$ is strictly positive-dimensional for all $x \in X$, then a sequence of blow-ups can still be performed to obtain a variety $\widehat{X}$ satisfying (ss=Ms($\widehat{U}$)) instead (see [BDHK16, Remark 1.10]). This is analogous to the case in classical GIT when $\emptyset = X^s \neq X^{Ms} \subsetneq X^{ss}$. Note that while in the reductive case the semistable locus must be non-empty in order to be able to perform the sequence of blow-ups, the $\widehat{U}$-blow-ups can always be performed. Unlike in Figure .1b, we obtain only one projective completion for the geometric quotient $X^{\widehat{s},\widehat{U}}/\widehat{U}$, coming from $\widehat{X}$, and it cannot in general be viewed as a partial desingularisation of another projective completion.

By analogy with classical GIT, a version of the stable locus inside X, denoted $X^{\widehat{s},\widehat{U}}$, can be defined as

$$X^{\widehat{s},\widehat{U}} := \widehat{\pi}(\widehat{X}^{s,\widehat{U}}_{\min+} \setminus \widehat{E})$$

where $\widehat{E}$ denotes the exceptional divisor for the blow-down map $\widehat{\pi} : \widehat{X} \to X$. If

$\widehat{X}^{s,\widehat{U}}_{\min+} = \emptyset$ then we define $X^{\widehat{s},\widehat{U}} := \hat{\pi}(U\widehat{Z}_{\min} \setminus \widehat{E})$ instead, where $\widehat{Z}_{\min}$ is the analogue of $Z_{\min}$ for $\widehat{X}$.

1.3 GIT for internally graded unipotent radical

More generally, let $H = U \rtimes R$ be a linear algebraic group with internally graded unipotent radical acting linearly on X. We let $\lambda : \mathbb{G}_m \to R$ denote a central one-parameter subgroup such that $\lambda(\mathbb{G}_m)$ acts on $\operatorname{Lie} U$ with strictly positive weights and we let R_λ denote the reductive quotient group $R/\lambda(\mathbb{G}_m)$.

When the condition (ss=s($\widehat{U}$)) is satisfied for the action of $\widehat{U} = U \rtimes \lambda(\mathbb{G}_m)$ on X, then by Section 1.2 there exists a projective geometric quotient $X//\widehat{U}$ for the action of $\widehat{U}$ on $X^{s,\widehat{U}}_{\min+}$ (if $X^{s,\widehat{U}}_{\min+} = \emptyset$ so that $X = \overline{UZ_{\min}}$ then in what follows we simply replace all occurrences of $X^{s,\widehat{U}}_{\min+}$ and $X//\widehat{U}$ with $UZ_{\min}$ and $Z_{\min}$ respectively). The projective variety $X//\widehat{U}$ has an induced action of the reductive group R_λ which is linear with respect to an ample line bundle $L_{\widehat{U}}$ on $X//\widehat{U}$ such that $q^*_{\widehat{U}} L_{\widehat{U}}$ is a tensor power of L where $q_{\widehat{U}} : X^{s,\widehat{U}}_{\min+} \to X//\widehat{U}$ denotes the quotient map. If $(X//\widehat{U})^{ss,R_\lambda} \neq \emptyset$ then we can further quotient by R_λ to obtain a projective variety $X//H := \left(X//\widehat{U}\right)//R_\lambda$ which is a geometric quotient for the action of R_λ on $(X//\widehat{U})^{s,R_\lambda}$, and for the action of H on

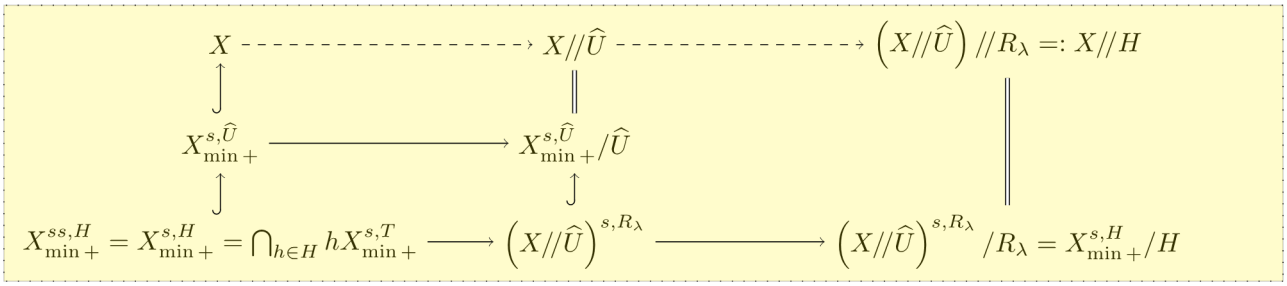
$$X^{s,H}_{\min+} := q^{-1}_{\widehat{U}}((X//\widehat{U})^{s,R_\lambda}) \cap X^{s,\widehat{U}}_{\min+}.$$

The semistable locus is analogously defined as $X^{ss,H}_{\min+} := q^{-1}_{\widehat{U}}((X//\widehat{U})^{ss,R_\lambda}) \cap X^{s,\widehat{U}}_{\min+}$ and both the stable and semistable loci satisfy Hilbert-Mumford-type criteria:

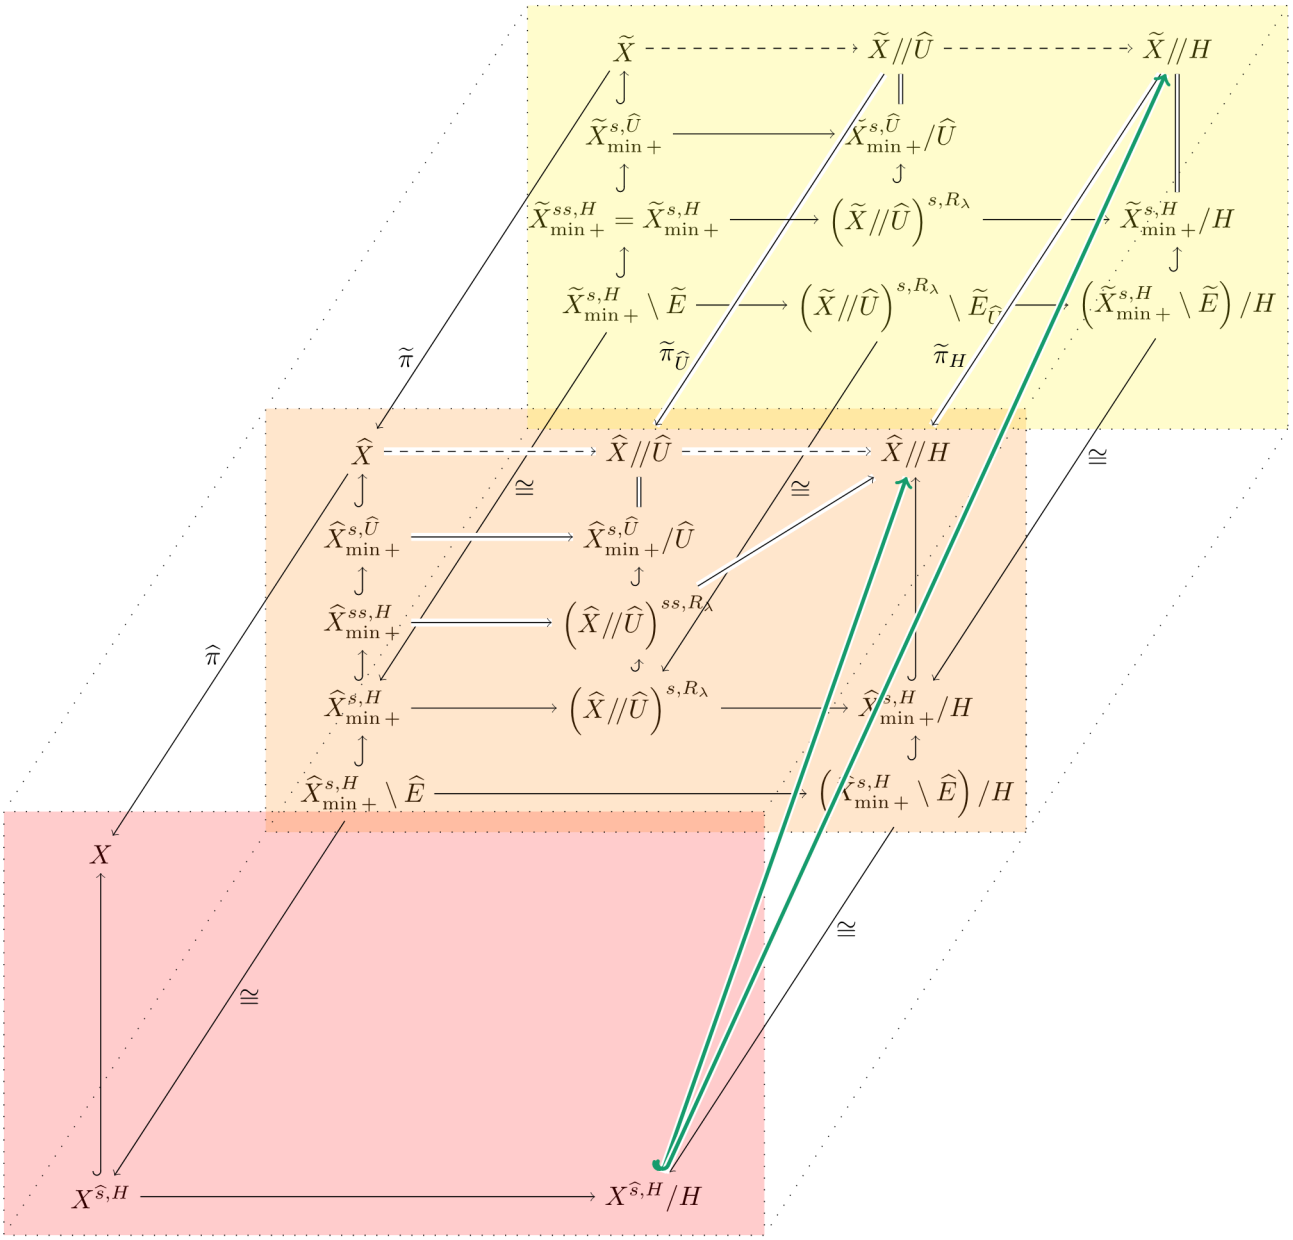
$$X^{(s)s,H}_{\min+} = \bigcap_{h\in H} hX^{(s)s,T}_{\min+}$$

for a fixed maximal torus $T \subseteq R$ containing the central one-parameter subgroup which grades U. If semistability coincides with stability for the action of R_λ on $X//\widehat{U}$, then $X//H$ is a geometric quotient for the action of H on $X^{s,H}_{\min+}$ since in this case $X^{ss,H}_{\min+} = X^{s,H}_{\min+}$. Figure .3a illustrates this situation.

If the condition (ss=s($\widehat{U}$)) is not satisfied, then as described in Section 1.2 a sequence of blow-ups can be performed to obtain a projective variety $\widehat{X}$ for which either (ss=s($\widehat{U}$)) or (ss=Ms($\widehat{U}$)) holds. For the projective variety $\widehat{X}//\widehat{U}$, if $\emptyset \neq (\widehat{X}//\widehat{U})^{s,R_\lambda} \subsetneq (\widehat{X}//\widehat{U})^{ss,R_\lambda}$, then as described in Section 1.1 an additional sequence of blow-ups can be performed to produce a projective variety $\widetilde{\widehat{X}//\widehat{U}}$ such that semistability coincides with stability (or Mumford-stability) for a suitably lifted action of R_λ. Equivalently



(a) When (ss=s($\widehat{U}$)) for the action of $\widehat{U}$, $X \neq \overline{UZ_{\min}}$ and $\left(X//\widehat{U}\right)^{ss,R_\lambda} = \left(X//\widehat{U}\right)^{s,R_\lambda} \neq \emptyset$



(b) When $\left(\widehat{X}//\widehat{U}\right)^{s,R/\lambda(\mathbb{G}_m)} \neq \emptyset$, $\widehat{X} \neq \overline{U\widehat{Z}_{\min}}$ and $\widetilde{X} \neq \overline{U\widetilde{Z}_{\min}}$. The green arrows denote projective completions.

Figure .3: GIT for linear algebraic groups with internally graded unipotent radical (in the setting of [BDHK16])

we can construct a variety $\widetilde{X}$ with a blow-down map to $\widehat{X}$ and a linearised action of H such that $\widetilde{X}//\widehat{U} = \widetilde{\widehat{X}//\widehat{U}}$ is a projective geometric quotient for the action of H on $\widetilde{X}^{s,H}_{\min+}$. An alternate stable locus inside X, denoted $X^{\widehat{s},H}$, is then defined as

$$X^{\widehat{s},H} := \widehat{\pi}(\widehat{X}^{s,H}_{\min+} \setminus \widehat{E}) = \widehat{\pi}(\widetilde{\pi}(\widetilde{X}^{s,H}_{\min+} \setminus \widetilde{E})) \setminus \widehat{E})$$

where $\widehat{E}$ and $\widetilde{E}$ denote the exceptional divisors for the blow-down maps $\widehat{\pi} : \widehat{X} \to X$ and $\widetilde{\pi} : \widetilde{X} \to \widehat{X}$ respectively. Figure .3b illustrates this sequence of blow-ups. By construction, $X^{\widehat{s},H}/H$ is a geometric quotient for the action of X on $X^{\widehat{s},H}$ and $\widetilde{X}//\widehat{U}$ is a projective completion of this quotient.

2 Algorithm for internally graded quotients

In the set-up of Section 1.3, if H acts linearly on X, then provided the stable or Mumford-stable locus for the induced action of $R_\lambda = R/\lambda(\mathbb{G}_m)$ on $\widehat{X}//\widehat{U}$ (or on $Z_{\min}$ if $\overline{UZ_{\min}} = X$) is non-empty, we obtain a non-empty open subset $X^{\widehat{s},H} \subseteq X$ such that $X^{\widehat{s},H}/H$ is a geometric quotient for the action of H, as well as a projective completion $\widetilde{X}//H$ which is itself a projective geometric quotient for the action of H on an open subset of $\widetilde{X}$. This is the main result of [BDHK16]. The purpose of this section is to describe a generalisation of this result to cases where the above condition on the induced reductive action is not satisfied, based on the inductive construction from [BHK17]. The proof is in the form of an algorithm, called the 'blow-up algorithm'.

Let $H = U \rtimes R$ be a linear algebraic group with unipotent radical U graded by a central one-parameter subgroup $\lambda : \mathbb{G}_m \to R$ of R, acting linearly on an irreducible projective variety X. We assume moreover that X is non-empty and that H is connected; otherwise we replace H with its identity component. We denote the trivial one-parameter subgroup by λ_0. As before, we let R_λ denote the reductive quotient group $R/\lambda(\mathbb{G}_m)$. Let $\mathcal{L}$ denote the linearisation for the action of H on X. Given this data, together with an invariant inner product on $\operatorname{Lie} R$, the blow-up algorithm produces a non-empty open subset $X^{\widehat{s},H}$ of X which admits a quasi-projective geometric quotient by the action of H (cf. [BHK17]), together with a projective completion $\mathrm{Bl}(X)//H$ of the geometric quotient $X^{\widehat{s},H}/H$, where $\mathrm{Bl}(X)$ is a projective variety with a suitably linearised action of H obtained from a sequence of H-equivariant blow-ups starting from X. Moreover, $\mathrm{Bl}(X)//H$ is itself a projective geometric quotient for the action of H on an open subset of $\mathrm{Bl}(X)$.

This blow-up algorithm is based on [BHK17]. It will be described in Section 2.2

and is illustrated in Figure .5. It relies on another self-standing algorithm called the 'replacement algorithm', which we describe in Section 2.1 and which is illustrated in Figure .4. The replacement algorithm also appears on the left-hand side of Figure .5.

2.1 Replacement algorithm

This algorithm is illustrated in Figure .4 and is described in words below.

Replacement algorithm. *Suppose that* $(X, H, \lambda, \mathcal{L})$ *encodes the data consisting of the linearised action of a non-reductive group* H *with internally graded unipotent radical* U *and Levi subgroup* R *on an irreducible projective variety* X *as above, and let* $R_\lambda = R/\lambda(\mathbb{G}_m)$*. Here we fix an adjoint-invariant innver product on* $\operatorname{Lie} R$*; this restricts to an adjoint-invariant inner product on* $\operatorname{Lie} R'$ *for any* $R' \leq R$*. One of three conditions must be satisfied for the action of* H *on* X*:*

$(\star)$ $\lambda(\mathbb{G}_m)$ *acts non-trivially on* X*;*

$(\star\star)$ $\lambda(\mathbb{G}_m)$ *acts trivially on* X *and* $X^{ss,R_\lambda} \neq \emptyset$*; or*

$(\star\star\star)$ $\lambda(\mathbb{G}_m)$ *acts trivially on* X *and* $X^{ss,R_\lambda} = \emptyset$*.*

Note that if $\lambda(\mathbb{G}_m)$ *acts trivially on* X*, then* U *must also act trivially on* X*, and so* X *is already a geometric quotient for the action of* $\widehat{U}$ *on* X*. Thus it suffices to consider the action of* $H/\lambda(\mathbb{G}_m) = R/\lambda(\mathbb{G}_m) = R_\lambda$ *on* X*, and in particular we can consider the semistable locus for this reductive action which determines which of* $(\star\star)$ *and* $(\star\star\star)$ *is satisfied.*

The input of the replacement algorithm is a 4*-tuple* $(X, H, \lambda, \mathcal{L})$ *satisfying* $(\star\star\star)$*. The algorithm then produces another such* 4*-tuple* $(X', H', \lambda', \mathcal{L}')$ *satisfying either* $(\star)$ *or* $(\star\star)$*. The construction of* $(X', H', \lambda', \mathcal{L}')$ *from* $(X, H, \lambda, \mathcal{L})$ *is depicted in Figure .4, and can be described in words as follows:*

1 Step. *Since by assumption* $(X, H, \lambda, \mathcal{L})$ *satisfies* $(\star\star\star)$*, the semistable locus for the induced action of* $R_\lambda = R/\lambda(\mathbb{G}_m)$ *on* X *is empty. Thus there exists an unstable stratum* S_β *with* $\beta \neq \emptyset$ *which is open in* X*, such that* $X = \overline{S_\beta} = \overline{R_\lambda Y_\beta^{ss}}$*. As observed in Section 1.1, this means that* $\overline{Y_\beta^{ss}}$ *is irreducible and non-empty. Let* P_β *denote the parabolic subgroup of* R_λ *determined by* β *and let* $\lambda_\beta : \mathbb{G}_m \to R_\lambda$ *denote the corresponding one-parameter subgroup. We then replace* $(X, H, \lambda, \mathcal{L})$ *by* $(\overline{Y_\beta^{ss}}, P_\beta, \lambda_\beta, \mathcal{L}|_{\overline{Y_\beta^{ss}}})$ *and proceed to Step* 2*.*

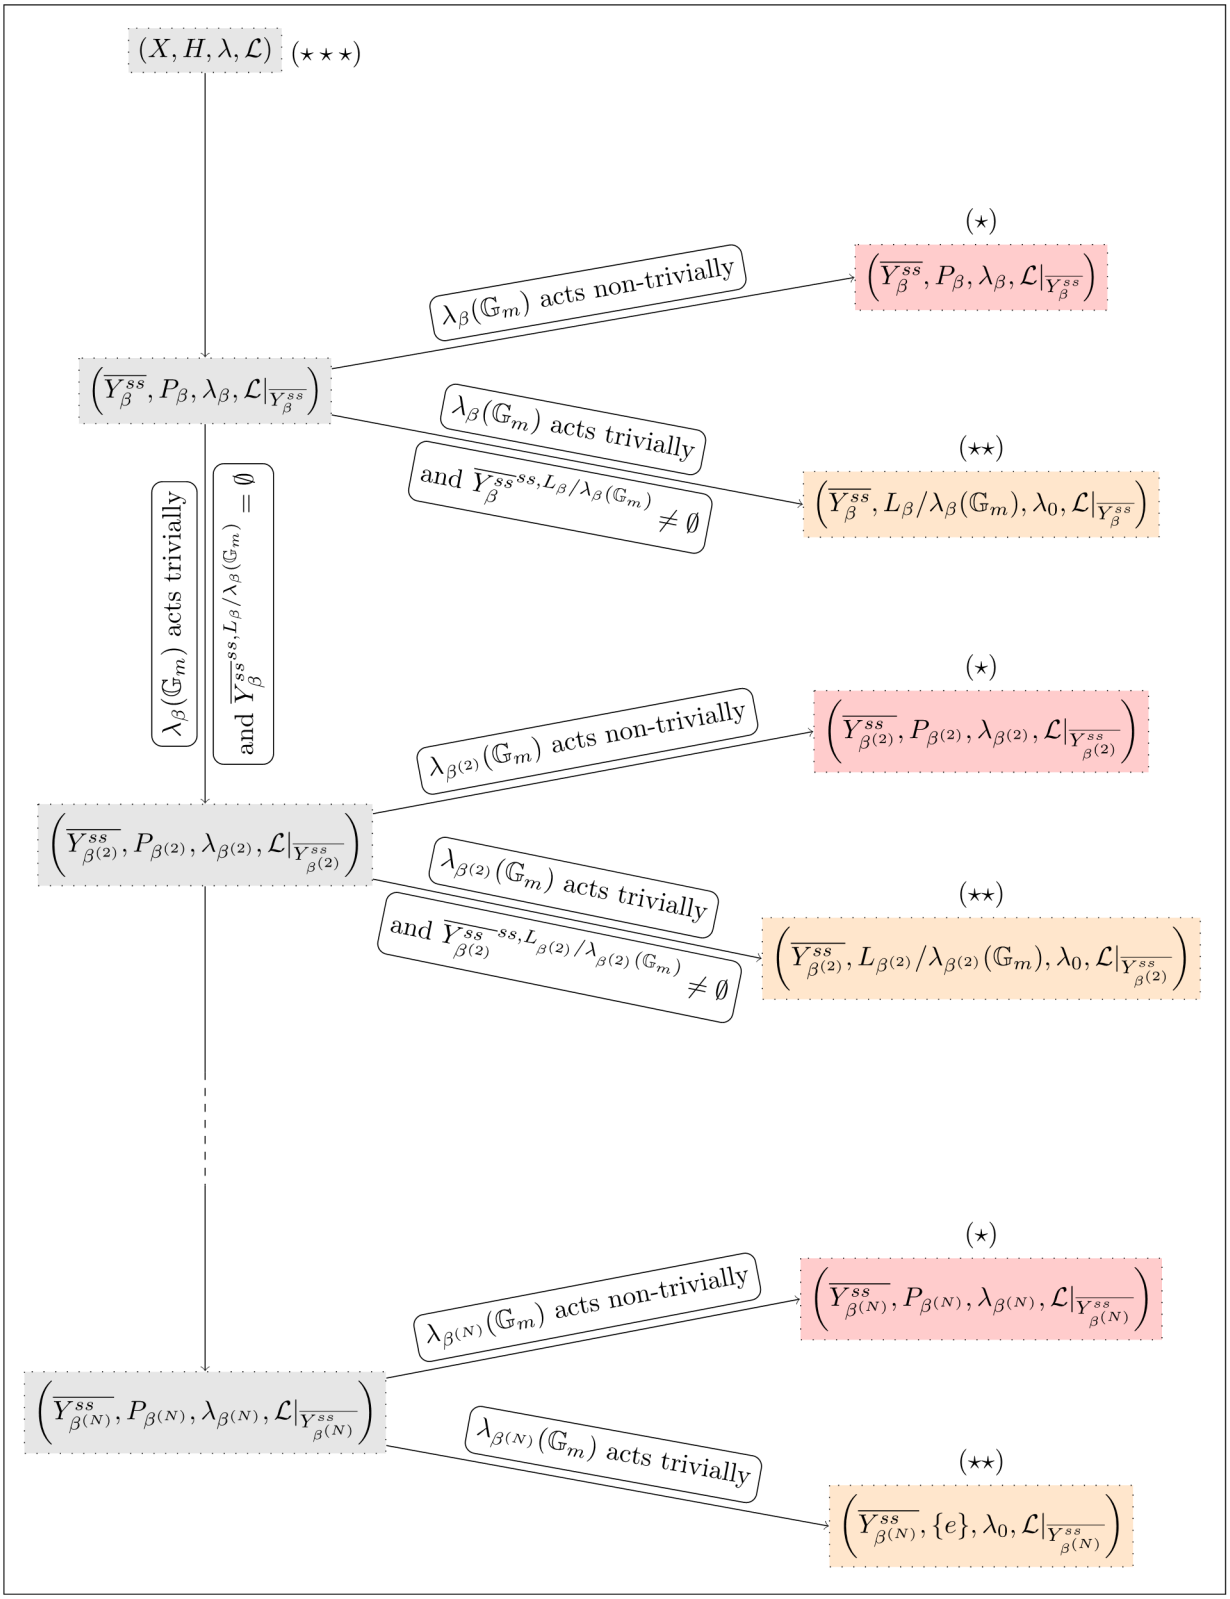


Figure .4: Replacement algorithm. The algorithm starts in the top left gray rectangle, and proceeds by following the arrows according to the conditions specified in the framed labels. The algorithm terminates when a red or orange rectangle is reached.

2 Step. *If* $(\overline{Y^{ss}_\beta}, P_\beta, \lambda_\beta, \mathcal{L}|_{\overline{Y^{ss}_\beta}})$ *satisfies* $(\star)$*, then we set* $(X', H', \lambda', \mathcal{L}') = (\overline{Y^{ss}_\beta}, P_\beta, \lambda_\beta, \mathcal{L}|_{\overline{Y^{ss}_\beta}})$ *and the algorithm terminates here.*

If $(\overline{Y^{ss}_\beta}, P_\beta, \lambda_\beta, \mathcal{L}|_{\overline{Y^{ss}_\beta}})$ *satisfies* $(\star\star)$*, then we set* $(X', H', \lambda', \mathcal{L}') = (\overline{Y^{ss}_\beta}, P_\beta/\lambda_\beta(\mathbb{G}_m), \lambda_0, \mathcal{L}|_{\overline{Y^{ss}_\beta}})$ *and the algorithm also terminates here.*

If $(\overline{Y^{ss}_\beta}, P_\beta, \lambda_\beta, \mathcal{L}|_{\overline{Y^{ss}_\beta}})$ *satisfies* $(\star\star\star)$*, then for the induced action of* R_λ *on* X*, there exists an unstable stratum* $S_{\beta^{(2)}}$ *with* $\beta^{(2)} \neq \emptyset$ *which is open in* $\overline{Y^{ss}_\beta}$*, such that* $\overline{Y^{ss}_\beta} = \overline{S_{\beta^{(2)}}} = \overline{R_\lambda Y^{ss}_{\beta^{(2)}}}$ *and* $\overline{R_\lambda Y^{ss}_{\beta^{(2)}}}$ *is a non-empty and irreducible projective variety (see Section 1.1). Let* L_β *denote a Levi subgroup of* P_β *and let* $P_{\beta^{(2)}}$ *denote the parabolic subgroup of* $L_\beta/\lambda_\beta(\mathbb{G}_m)$ *determined by* $\beta^{(2)}$*. Moreover, let* $\lambda_{\beta^{(2)}} : \mathbb{G}_m \to L_\beta/\lambda_\beta(\mathbb{G}_m)$ *denote the corresponding one-parameter subgroup. We then replace* $(\overline{Y^{ss}_\beta}, P_\beta, \lambda_\beta, \mathcal{L}|_{\overline{Y^{ss}_\beta}})$ *by* $(\overline{Y^{ss}_{\beta^{(2)}}}, P_{\beta^{(2)}}, \lambda_{\beta^{(2)}}, \mathcal{L}|_{\overline{Y^{ss}_{\beta^{(2)}}}})$ *and proceed to Step* 3*.*

3 Step. *Repeat Step* 2*, except with the output* $(\overline{Y^{ss}_{\beta^{(2)}}}, P_{\beta^{(2)}}, \lambda_{\beta^{(2)}}, \mathcal{L}|_{\overline{Y^{ss}_{\beta^{(2)}}}})$ *of the previous step as the starting* 4*-tuple.*

More generally, assuming that for each Step $i \leq k$ *the starting* 4*-tuple satisfies* $(\star\star\star)$*, then Step* $k+1$ *can be defined inductively. Let* $(\overline{Y^{ss}_{\beta^{(i)}}}, P_{\beta^{(i)}}, \lambda_{\beta^{(i)}}, \mathcal{L}|_{\overline{Y^{ss}_{\beta^{(i)}}}})$ *be the output of Step* i *for* $1 \leq i \leq k$ *(for* $i = 1$ *we let* $\beta^{(1)} = \beta$*).*

Step k + 1. *Repeat Step k, replacing* $(\overline{Y^{ss}_{\beta^{(k-1)}}}, P_{\beta^{(k-1)}}, \lambda_{\beta^{(k-1)}}, \mathcal{L}|_{\overline{Y^{ss}_{\beta^{(k-1)}}}})$ *with* $(\overline{Y^{ss}_{\beta^{(k)}}}, P_{\beta^{(k)}}, \lambda_{\beta^{(k)}}, \mathcal{L}|_{\overline{Y^{ss}_{\beta^{(k)}}}})$ *as the starting* 4*-tuple. If* $(\star\star\star)$ *holds, then* $\lambda_{\beta^{(k)}}(\mathbb{G}_m)$ *acts trivially on* $\overline{Y^{ss}_{\beta^{(k)}}}$ *and the semistable locus for the induced action of* $L_{\beta^{(k)}}/\lambda_{\beta^{(k)}}(\mathbb{G}_m)$ *on* $\overline{Y^{ss}_{\beta^{(k)}}}$ *is empty, where* $L_{\beta^{(k)}}$ *denotes a Levi subgroup of* $P_{\beta^{(k)}}$*. In this case, we let* $\beta^{(k+1)}$ *denote the non-trivial character determining the unstable stratum* $S_{\beta^{(k+1)}}$ *for the action of* $L_{\beta^{(k)}}/\lambda_{\beta^{(k)}}(\mathbb{G}_m)$ *on* $\overline{Y^{ss}_{\beta^{(k)}}}$ *such that* $\overline{Y^{ss}_{\beta^{(k)}}} = \overline{S_{\beta^{(k+1)}}}$*. Moreover, let* $Y^{ss}_{\beta^{(k+1)}}$ *denote the locally closed subset of* $\overline{Y^{ss}_{\beta^{(k)}}}$ *determined by this action such that* $S_{\beta^{(k+1)}} = L_{\beta^{(k)}}/\lambda_{\beta^{(k)}}(\mathbb{G}_m)Y^{ss}_{\beta^{(k+1)}}$*. Finally, let* $\lambda_{\beta^{(k+1)}} : \mathbb{G}_m \to L_{\beta^{(k)}}/\lambda_{\beta^{(k)}}(\mathbb{G}_m)$ *and* $P_{\beta^{(k+1)}}$ *denote the non-trivial one-parameter subgroup and parabolic subgroup respectively of* $L_{\beta^{(k)}}/\lambda_{\beta^{(k)}}(\mathbb{G}_m)$ *determined by* $\beta^{(k+1)}$*. With this notation, repeating Step k then produces the* 4*-tuple* $(\overline{Y^{ss}_{\beta^{(k+1)}}}, P_{\beta^{(k+1)}}, \lambda_{\beta^{(k+1)}}, \mathcal{L}|_{\overline{Y^{ss}_{\beta^{(k+1)}}}})$*.*

2.1.1 Proposition. *The replacement algorithm terminates after a finite number of steps.*

Proof. If we arrive at Step i for $i \geq 1$, then the only way for the algorithm to continue on to Step $i+1$ is if the starting 4-tuple $(\overline{Y^{ss}_{\beta^{(i-1)}}}, P_{\beta^{(i-1)}}, \lambda_{\beta^{(i-1)}}, \mathcal{L}|_{\overline{Y^{ss}_{\beta^{(i-1)}}}})$ satisfies $(\star\star\star)$. If so, then the 4-tuple must be replaced with $(\overline{Y^{ss}_{\beta^{(i)}}}, P_{\beta^{(i)}}, \lambda_{\beta^{(i)}}, \mathcal{L}|_{\overline{Y^{ss}_{\beta^{(i)}}}})$. But we have that $\dim L_{\beta^{(i)}} \leq \dim P_{\beta^{(i)}} \leq \dim L_{\beta^{(i-1)}}/\lambda_{\beta^{(i-1)}}(\mathbb{G}_m) < \dim L_{\beta^{(i-1)}}$, since $\lambda_{\beta^{(i-1)}}(\mathbb{G}_m)$ is a non-trivial one-parameter subgroup of $L_{\beta^{(i-1)}}$. Hence at each step the dimension

of the Levi subgroup appearing in the 4-tuple strictly decreases. This implies that if at each step the starting 4-tuple satisfies $(\star\star\star)$, we must eventually reach a Step N with starting 4-tuple $(\overline{Y^{ss}_{\beta(N-1)}}, P_{\beta(N-1)}, \lambda_{\beta(N-1)}, \mathcal{L}|_{\overline{Y^{ss}_{\beta(N)}}})$ such that $L_{\beta(N-1)}/\lambda_{\beta(N-1)}(\mathbb{G}_m)$ is trivial. Hence the semistable locus for the induced action of $L_{\beta(N-1)}/\lambda_{\beta(N-1)}(\mathbb{G}_m) = \{e\}$ on $\overline{Y^{ss}_{\beta(N-1)}}$ is all of $\overline{Y^{ss}_{\beta(N-1)}}$ and so $(\star\star)$ is automatically satisfied. Thus $(\star\star\star)$ cannot be satisfied by the starting 4-tuple at Step N and so the algorithm terminates. □

2.2 Blow-up algorithm

This algorithm is illustrated in Figure .5 and described in words below.

Blow-up algorithm. *Given a 4-tuple $(X, H, \lambda, \mathcal{L})$ encoding the data consisting of the linearised action of a non-reductive group H with internally graded unipotent radical U and Levi subgroup R on an irreducible projective variety X, this algorithm produces:*

i) a non-empty open subset $X^{\widehat{s},H} \subseteq X$ such that $X^{\widehat{s},H}/H$ is a geometric H-quotient; and

ii) a projective completion $\mathrm{Bl}(X)//H$ of $X^{\widehat{s},H}/H$ for a projective variety $\mathrm{Bl}(X)$ with a blow-down map $\pi : \mathrm{Bl}(X) \to X$ resulting from a sequence of H-equivariant blow-ups starting from X, such that the following condition, denoted $(\bullet)$, is satisfied: $\mathrm{Bl}(X)$ admits a linearised action of H with respect to a line bundle L_H which is a suitably perturbed tensor power of π^*L such that the action of $\widehat{U}$ on $\mathrm{Bl}(X)$ satisfies the condition (ss=s($\widehat{U}$)) and $\mathrm{Bl}(X)//X$ is a geometric quotient by H of an H-invariant open subset of $\mathrm{Bl}(X)$.

1 Case. *Suppose that $(X, H, \lambda, \mathcal{L})$ satisfies $(\star)$. We start in the red square of the right-hand rectangle in Figure .5. According to the diagram two paths can be taken,* 1 *or* 2*. Both correspond to performing the sequence of $\widehat{U}$-blow-ups described in Section 1.2. If $\widehat{X}^{s,\widehat{U}}_{\min+} \neq \emptyset$, then we obtain a non-empty projective variety $\widehat{X}//\widehat{U}$ which is a geometric quotient for the action of $\widehat{U}$ on $\widehat{X}^{s,\widehat{U}}_{\min+} \subseteq \widehat{X}$. If $\widehat{X}^{s,\widehat{U}}_{\min+} = \emptyset$, then we consider instead the projective variety $\widehat{Z}_{\min}$ which is a geometric quotient for the action of $\widehat{U}$ on $U\widehat{Z}_{\min}$. In either case, the non-empty projective variety admits an induced action of $R_\lambda = R/\lambda(\mathbb{G}_m)$, and the semistable locus for this linearised action determines the path to be taken, as described by the following subcases. For simplicity, we describe these subcases under the assumption that $\widehat{X}^{s,\widehat{U}}_{\min+} \neq \emptyset$, but if this is not the case then it suffices to replace all occurrences of $\widehat{X}^{s,\widehat{U}}_{\min+}$ and $\widehat{X}//\widehat{U}$ with $U\widehat{Z}_{\min}$ and $\widehat{Z}_{\min}$ respectively.*

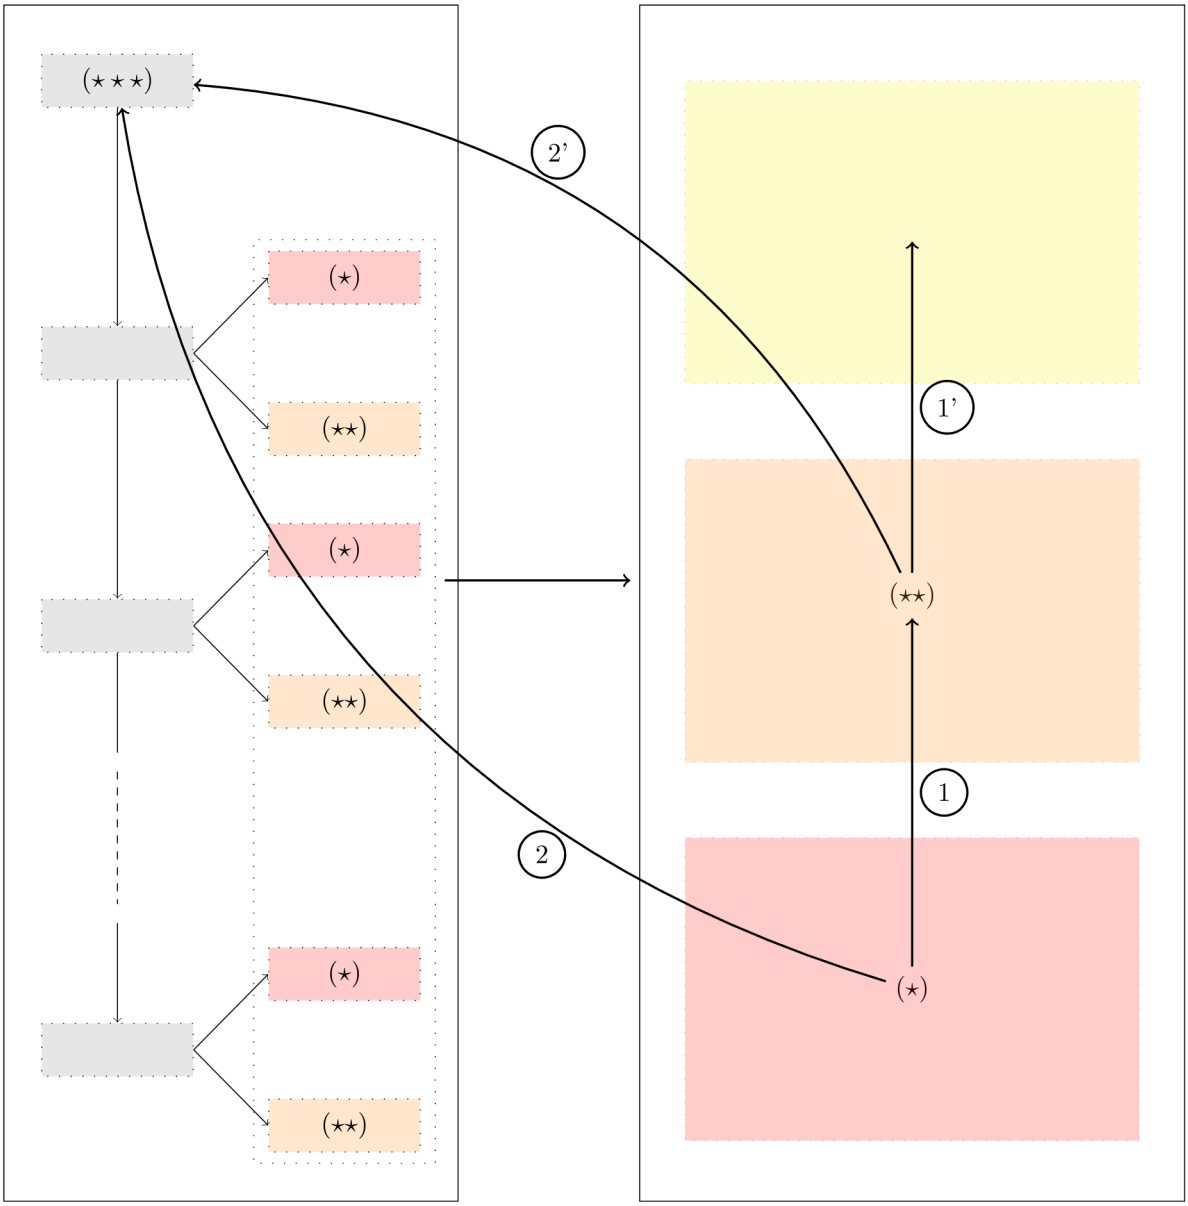
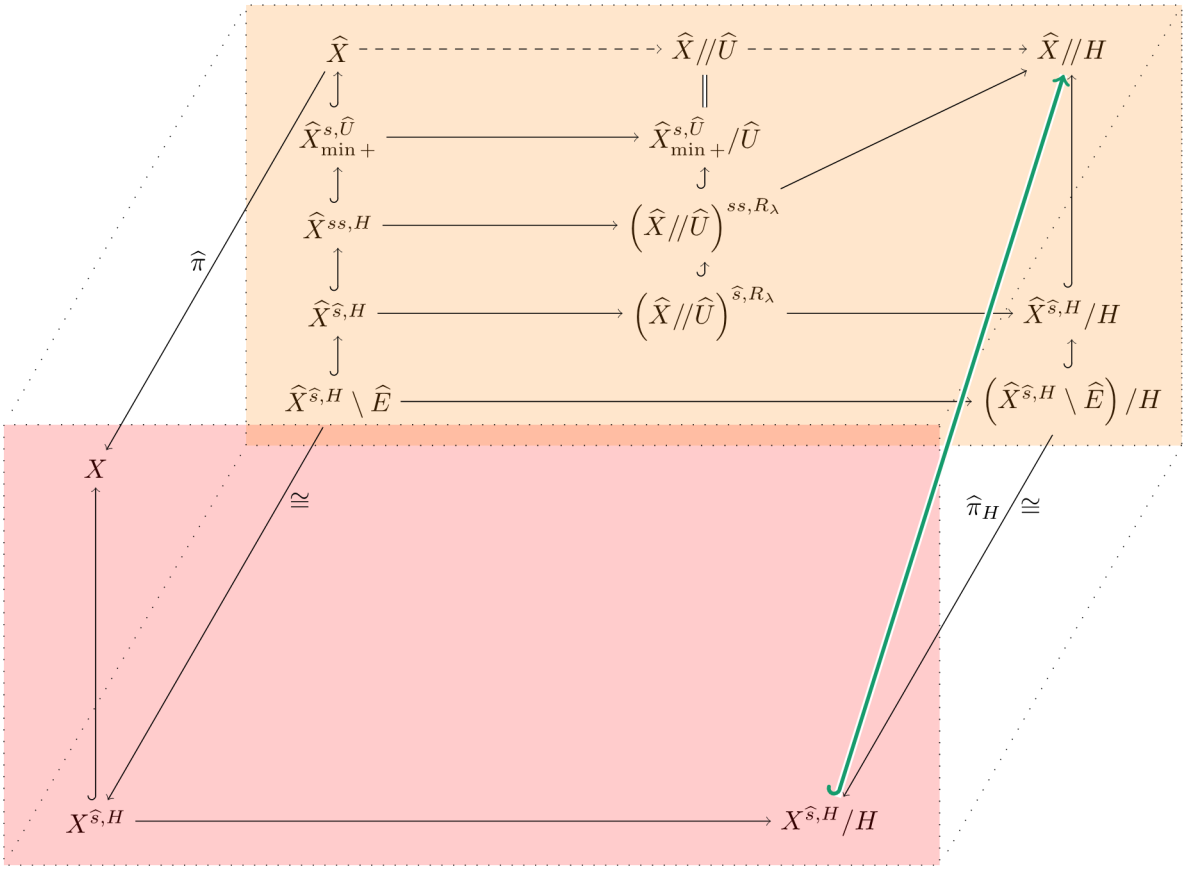
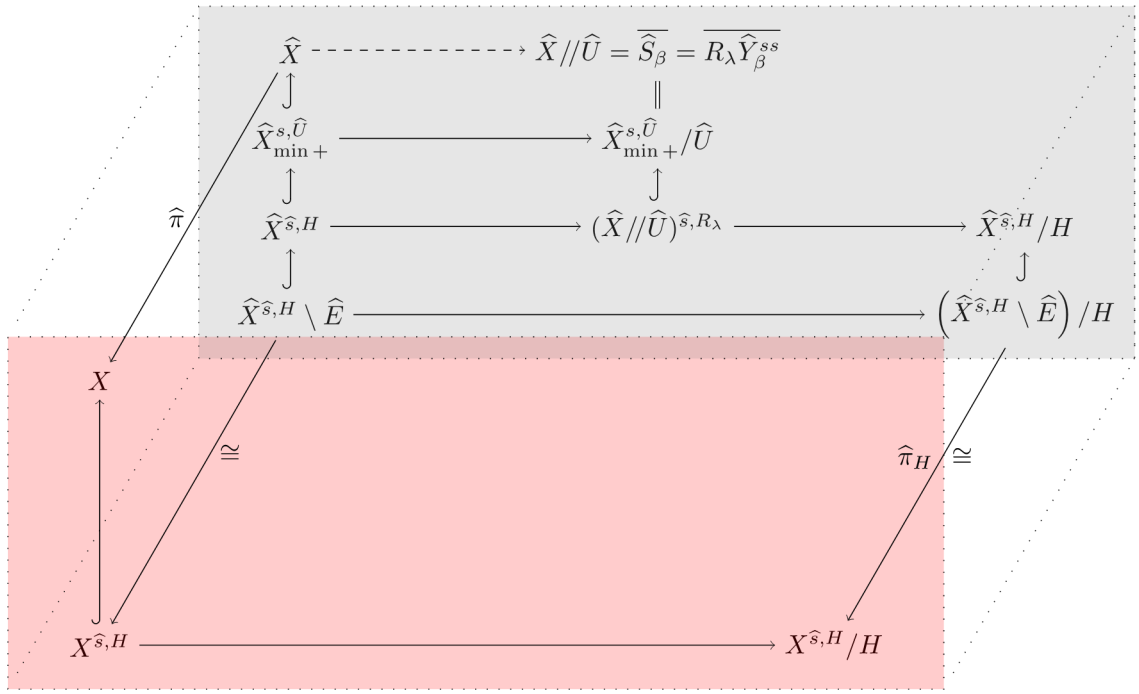


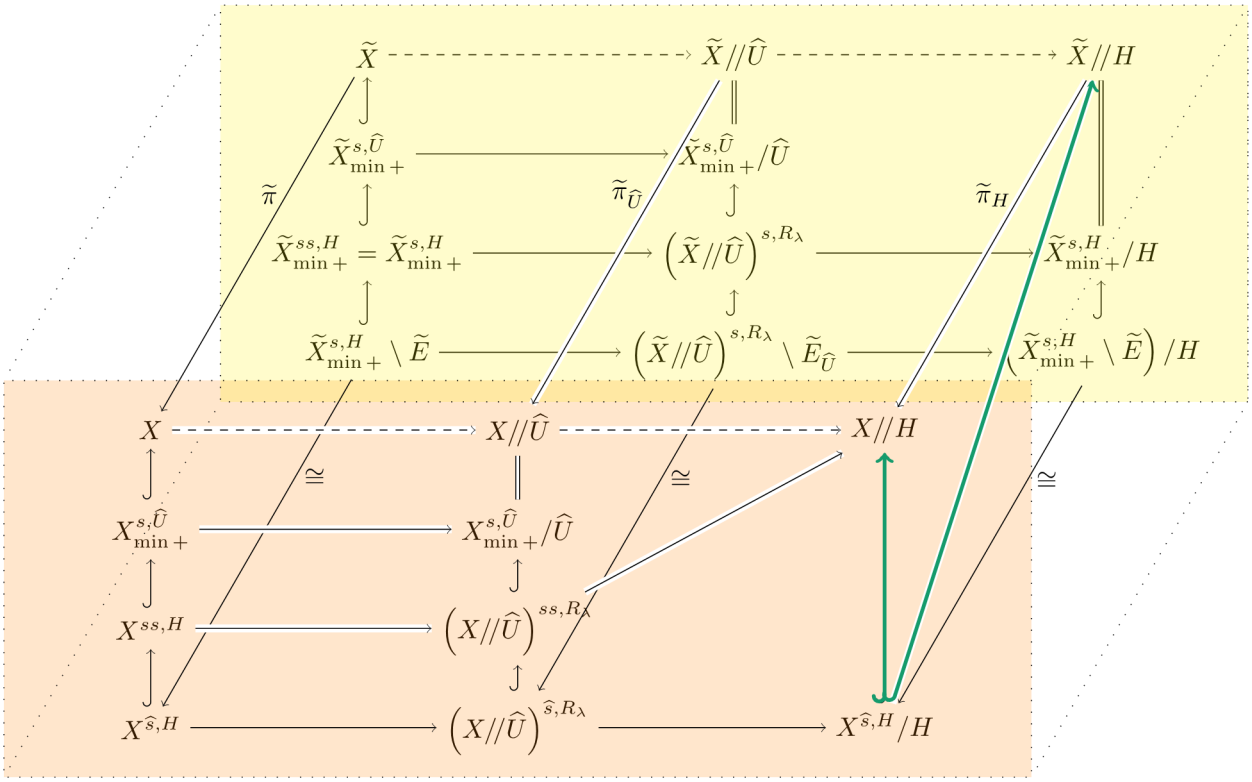
Figure .5: Blow-up algorithm for the action of an internally graded linear algebraic group on a projective variety. The left-hand side corresponds to the replacement algorithm, described in Figure .4. The sequence of blow-ups corresponding to each of the four paths 1, 2, $1'$ and $2'$ are described in Figure .6. Given a 4-tuple $(X, H, \lambda, \mathcal{L})$, the algorithm starts in the gray square in the top left or the red or orange squares on the right-hand side, depending on which of the three conditions $(\star)$, $(\star\star)$ and $(\star\star\star)$ is satisfied for the 4-tuple. At any point in the algorithm if we find ourselves on the gray square in the top left (either at the start or after taking paths 2 or $2'$), then we run the replacement algorithm. The output determines whether we proceed to the red or to the orange rectangle on the right-hand side: if the output is red we move to the red rectangle, if it is orange we move to the orange rectangle. Figure .7 depicts the different paths which can be taken within the algorithm and illustrates why the algorithm must terminate in the yellow rectangle after a finite number of steps (Proposition 2.2.1).



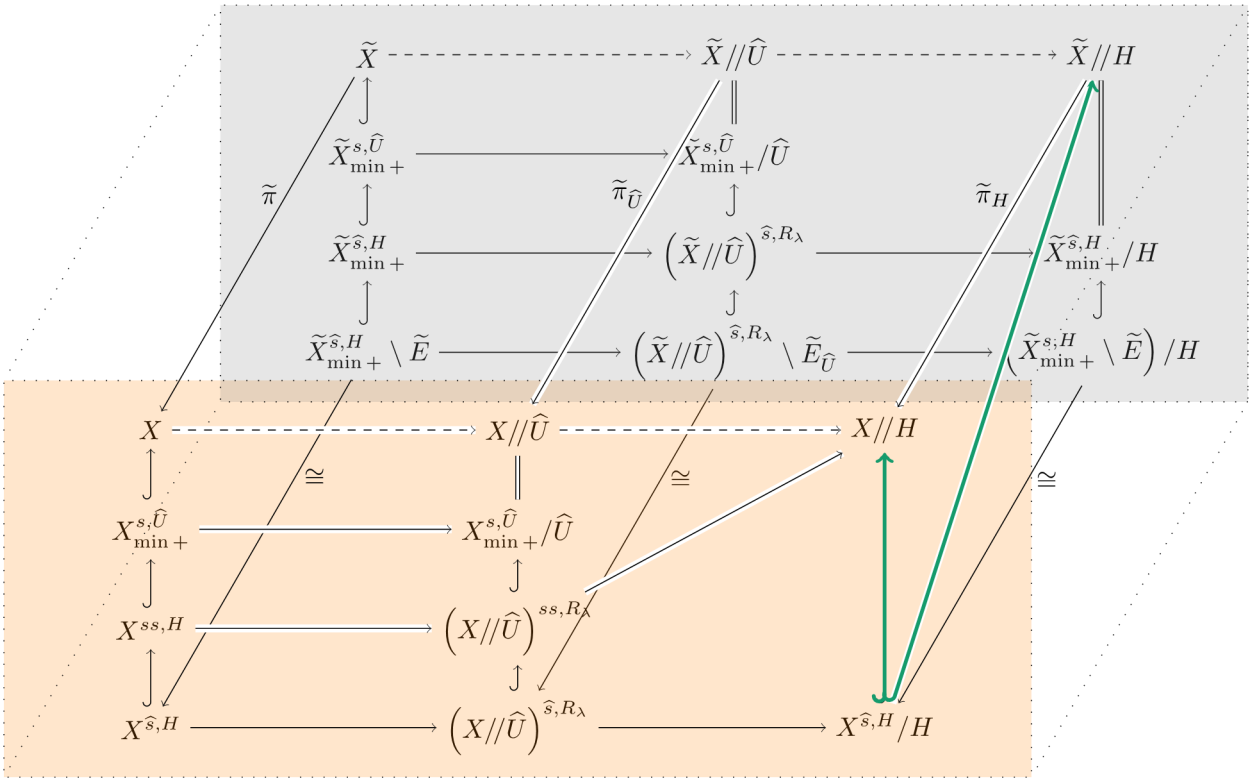
(a) Blow-up corresponding to Path 1 in Figure .5 (when $\widehat{X} \neq \overline{U\widehat{Z}_{\min}}$)



(b) Blow-up corresponding to Path 2 in Figure .5



(c) Blow-up corresponding to Path 1′ in Figure .5 (when $\widetilde{X} \neq \overline{U\widetilde{Z}_{\min}}$)



(d) Blow-up corresponding to Path 2′ in Figure .5 (when $\widetilde{X} \neq \overline{U\widetilde{Z}_{\min}}$)

Figure .6: Blow-ups from the blow-up algorithm

1.1 Case. *If* $(\widehat{X}//\widehat{U})^{ss,R_\lambda} \neq \emptyset$*, then path* 1 *must be taken, as illustrated in Figure .6a. By assumption, the* 4*-tuple* $(\widehat{X}//\widehat{U}, R_\lambda, \lambda_0, \widehat{\mathcal{L}})$ *then satisfies* $(\star\star)$*, where* λ_0 *denotes the trivial one-parameter subgroup of* R_λ*. We then proceed to Case* 2 *applied to the* 4*-tuple* $(\widehat{X}//\widehat{U}, R_\lambda, \lambda_0, \widehat{\mathcal{L}})$ *in order to construct the open subset* $(\widehat{X}//\widehat{U})^{\widehat{s},R_\lambda} \subseteq \widehat{X}//\widehat{U}$ *and the variety* $\mathrm{Bl}(\widehat{X}//\widehat{U})$*. Having done so, we set*

$$X^{\widehat{s},H} := \widehat{\pi}\left(\left(q_{\widehat{U}}^{-1}\left(\left(\widehat{X}//\widehat{U}\right)^{\widehat{s},R_\lambda}\right) \cap \widehat{X}^{s,\widehat{U}}_{\min+}\right) \setminus \widehat{E}\right) \tag{.1}$$

where $\widehat{E}$ *is the exceptional divisor corresponding to the* $\widehat{U}$*-blow-up* $\widehat{\pi} : \widehat{X} \to X$*. The variety* $\mathrm{Bl}(X)$ *is defined as follows. Let* $(\mathrm{Bl}_{(i)}(\widehat{X}//\widehat{U}))_{i=0}^n$ *be the sequence of blow-ups determining* $\mathrm{Bl}(\widehat{X}//\widehat{U})$*, so that* $\mathrm{Bl}_{(0)}(\widehat{X}//\widehat{U}) = \widehat{X}//\widehat{U}$ *and* $\mathrm{Bl}_{(n)}(\widehat{X}//\widehat{U}) = \mathrm{Bl}(\widehat{X}//\widehat{U})$*. For* $i \in \{1, \ldots, n\}$*,* $\mathrm{Bl}_{(i)}(\widehat{X}//\widehat{U})$ *is the blow-up of* $\mathrm{Bl}_{(i-1)}(\widehat{X}//\widehat{U})$ *along a closed subvariety, say* $W_{(i-1)}$*. We then set* $\mathrm{Bl}_{(0)}(\widehat{X}) = \widehat{X}$ *and inductively define* $\mathrm{Bl}_{(i)}(\widehat{X})$ *for* $i \geq 1$ *to be the blow-up of* $\mathrm{Bl}_{(i-1)}(\widehat{X})$ *along the closure of the subvariety* $q_{\widehat{U}}^{-1}(W)$ *in* $\mathrm{Bl}_{(i-1)}(\widehat{X})$ *and set* $\mathrm{Bl}(\widehat{X}) = \mathrm{Bl}_{(n)}(\widehat{X})$*. By post-composing with the blow-down map* $\widehat{\pi} : \widehat{X} \to X$*, we obtain a blow-down map* $\mathrm{Bl}(\widehat{X}) \to X$ *and we set* $\mathrm{Bl}(X) = \mathrm{Bl}(\widehat{X})$*.*

1.2 Case. *If* $(\widehat{X}//\widehat{U})^{ss,R_\lambda} = \emptyset$*, then path* 2 *must be taken, as illustrated in Figure .6b. From the assumption determining this subcase, the semistable locus for the action of* R_λ *on* $\widehat{X}//\widehat{U}$ *is empty. Thus the* 4*-tuple* $(\widehat{X}//\widehat{U}, R_\lambda, \lambda_0, \widehat{\mathcal{L}})$ *satisfies* $(\star\star\star)$*. We then proceed to Case* 3 *applied to the* 4*-tuple* $(\widehat{X}//\widehat{U}, R_\lambda, \lambda_0, \widehat{\mathcal{L}})$ *in order to construct the open set* $(\widehat{X}//\widehat{U})^{\widehat{s},R_\lambda} \subseteq \widehat{X}//\widehat{U}$ *and the variety* $\mathrm{Bl}(\widehat{X}//\widehat{U})$*. Having done so, we set*

$$X^{\widehat{s},H} := \widehat{\pi}\left(\left(q_{\widehat{U}}^{-1}\left(\left(\widehat{X}//\widehat{U}\right)^{\widehat{s},R_\lambda}\right) \cap \widehat{X}^{s,\widehat{U}}_{\min+}\right) \setminus \widehat{E}\right) \tag{.2}$$

and define $\mathrm{Bl}(X)$ *in terms of* $\mathrm{Bl}(\widehat{X}//\widehat{U})$ *as in Case* 1.1*.*

2 Case. *Suppose that* $(X, H, \lambda, \mathcal{L})$ *satisfies* $(\star\star)$*. We start in the orange square of the right-hand rectangle in Figure .5. According to the diagram two paths can be taken,* $1'$ *or* $2'$*. Both correspond to performing the sequence of reductive blow-ups described in Section 1.1. Since* X^{ss,R_λ} *is non-empty, we can perform a reductive blow-up of* X *to obtain a projective variety* $\widetilde{X}$ *acted upon linearly by* R_λ *such that semistability coincides with stability (or Mumford-stability if there are no properly stable points) for the action of* R_λ *on* $\widetilde{X}$*. The semistable locus for this linearised action determines the path to be taken, as described in the following subcases.*

2.1 Case. *If* $\widetilde{X}^{ss,R_\lambda} \neq \emptyset$*, then path* $1'$ *must be taken, as illustrated in Figure .6c. Taking this path corresponds to performing the* R_λ*-blow-up of* X *to obtain a projective variety* $\widetilde{X}$ *for which* $\widetilde{X}^{ss,R_\lambda} = \widetilde{X}^{(M)s,R_\lambda} \neq \emptyset$*. We then define*

$$X^{\widehat{s},H} = \widetilde{\pi}\left(\widetilde{X}^{ss,R_\lambda} \setminus \widetilde{E}\right), \tag{.3}$$

where $\widetilde{E}$ *is the exceptional divisor corresponding to the blow-up* $\widetilde{\pi} : \widetilde{X} \to X$*, and*

$$\mathrm{Bl}(X) = \widetilde{X}.$$

The algorithm terminates here.

2.2 Case. *If* $\widetilde{X}^{ss,R_\lambda} = \emptyset$*, then path* $2'$ *must be taken, as illustrated in Figure .6d. Taking this path corresponds to performing the* R_λ*-blow-up of* X *to obtain a projective variety* $\widetilde{X}$ *for which* $\widetilde{X}^{ss,R_\lambda} = \emptyset$*. Thus the* 4*-tuple* $(\widetilde{X}, R_\lambda, \lambda_0, \widetilde{\mathcal{L}})$ *satisfies* $(\star\star\star)$*. We then proceed to Case* 3 *applied to the* 4*-tuple* $(\widehat{X}//\widehat{U}, R_\lambda, \lambda_0, \widehat{\mathcal{L}})$ *in order to construct the open set* $\widetilde{X}^{\widehat{s},R_\lambda} \subseteq \widetilde{X}$ *and the variety* $\mathrm{Bl}(\widetilde{X})$*. Having done so, we set*

$$X^{\widehat{s},H} = \widetilde{\pi}\left(\widetilde{X}^{\widehat{s},R_\lambda} \setminus \widetilde{E}\right)$$

where $\widetilde{E}$ *is the exceptional divisor corresponding to the blow-up* $\widetilde{\pi} : \widetilde{X} \to X$*, and*

$$\mathrm{Bl}(X) = \mathrm{Bl}(\widetilde{X}).$$

3 Case. *Suppose that* $(X, H, \lambda, \mathcal{L})$ *satisfies* $(\star\star\star)$*. Then we can apply the replacement algorithm to this* 4*-tuple to obtain another* 4*-tuple* $(X', H', \lambda', \mathcal{L}')$ *satisfying* $(\star)$ *or* $(\star\star)$*.*

3.1 Case. *If* $(\star)$ *is satisfied, then* $(X', H', \lambda', \mathcal{L}')$ *must be of the form* $(\overline{Y^{ss}_{\beta(k)}}, P_{\beta(k)}, \lambda_{\beta(k)}, \mathcal{L}|_{\overline{Y^{ss}_{\beta(k)}}})$ *for some* $k \geq 1$ *(see Figure .4). We then proceed to Case* 1 *applied the* 4*-tuple* $(X', H', \lambda', \mathcal{L}')$ *in order to construct the open set* $X'^{\widehat{s},H'} = \overline{Y^{ss}_{\beta(k)}}^{\widehat{s},P_{\beta(k)}} \subseteq \overline{Y^{ss}_{\beta(k)}}$ *and the variety* $\mathrm{Bl}(X')$*. Having done so, we set*

$$X^{\widehat{s},H} = R\left(L_\beta\left(\cdots\left(L_{\beta(k-1)}\left(\overline{Y^{ss}_{\beta(k)}}^{\widehat{s},P_{\beta(k)}} \cap Y^{ss}_{\beta(k)}\right)\cdots\right)Y^{ss}_{\beta(2)}\right) \cap Y_{\beta^{ss}}\right).$$

Moreover, we define $\mathrm{Bl}(X)$ *as follows. Let* $(\mathrm{Bl}_{(i)}(X'))_{i=0}^n$ *be the sequence of blow-ups determining* $\mathrm{Bl}(X')$ *so that* $\mathrm{Bl}_{(0)}(X') = X'$ *and* $\mathrm{Bl}_{(n)}(X') = \mathrm{Bl}(X')$*. For* $i \in \{1, \ldots, n\}$*,* $\mathrm{Bl}_{(i)}(X')$ *is the blow-up of* $\mathrm{Bl}_{(i-1)}(X')$ *along a closed subvariety of* $\mathrm{Bl}_{(i-1)}(X')$*, say* $W_{(i-1)}$*. We then set* $\mathrm{Bl}_{(0)}(X) = X$ *and inductively define* $\mathrm{Bl}_{(i)}(X)$ *for* $i \geq 1$ *to be the blow-up of* $\mathrm{Bl}_{(i-1)}(X)$ *along the closure of* $R(L_\beta(\cdots(L_{\beta(k-1)}W_{(i-1)})\cdots))$ *inside*

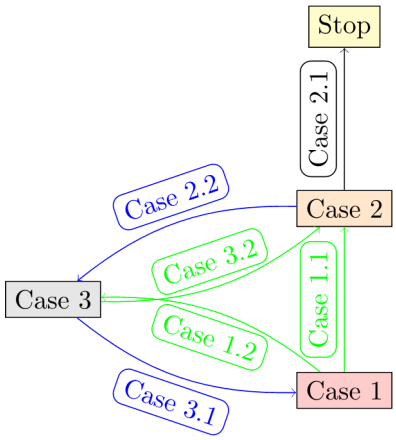


Figure .7: Proof that the blow-up algorithm must terminate. A green (respectively blue) arrow indicates that the dimension of the group decreases (respectively decreases or remains the same) upon passing from one case to the other. Since each loop contains at least one green arrow, when going through a loop the dimension of the group must always decrease.

$\mathrm{Bl}_{(i-1)}(X)$.

3.2 Case. *If* $(X', H', \lambda', \mathcal{L}')$ *satisfies* $(\star\star)$ *instead, then it is of the form* $(\overline{Y^{ss}_{\beta^{(k)}}}, P_{\beta^{(k)}}/\lambda_{\beta^{(k)}}(\mathbb{G}_m), \lambda_{\beta^{(k)}},$ *for some* $k \geq 1$. *We then proceed to Case* 1 *applied to the* 4*-tuple* $(X', H', \lambda', \mathcal{L}')$ *in order to construct the open set* $X'^{\widehat{s},H'} = \overline{Y^{ss}_{\beta^{(k)}}}^{\widehat{s},P_{\beta^{(k)}}} \subseteq \overline{Y^{ss}_{\beta^{(k)}}}$ *and the variety* $\mathrm{Bl}(X')$. *Having done so, we set*

$$X^{\widehat{s},H} = R\left(L_\beta\left(\cdots\left(L_{\beta^{(k-1)}}\left(\overline{Y^{ss}_{\beta^{(k)}}}^{\widehat{s},P_{\beta^{(k)}}/\lambda_{\beta^{(k)}}(\mathbb{G}_m)} \cap Y^{ss}_{\beta^{(k)}}\right)\cdots\right)Y^{ss}_{\beta^{(2)}}\right)\cap Y_{\beta^{ss}}\right)$$

and define $\mathrm{Bl}(X)$ *in terms of* $\mathrm{Bl}(X')$ *as in Case* 3.1.

2.2.1 Proposition. *The blow-up algorithm terminates after a finite number of steps.*

Proof. Suppose that the initial 4-tuple is $(X, H, \lambda, \mathcal{L})$. By construction, the blow-up algorithm can only terminate by reaching Case 2.1.

If $(\star)$ is satisfied for $(X, H, \lambda, \mathcal{L})$, then we start in Case 1, according to which $(X, H, \lambda, \mathcal{L})$ is replaced by $(\widehat{X}//\widehat{U}, R_\lambda, \lambda_0, \widehat{\mathcal{L}})$. Since $(\star)$ holds for $(X, H, \lambda, \mathcal{L})$, the one-parameter subgroup $\lambda(\mathbb{G}_m)$ acts non-trivially on X and hence is itself non-trivial. Thus we have that $\dim R_\lambda = \dim R/\lambda(\mathbb{G}_m) < \dim R \leq \dim H$, and so whether we move to Case 2 (if Case 1.1 holds) or Case 3 (if Case 1.2 holds), the new starting 4-tuple involves a group of dimension strictly smaller than the dimension of H.

If $(\star\star)$ is satisfied for $(X, H, \lambda, \mathcal{L})$, then we start in Case 2. If Case 2.1 holds, then the algorithm terminates at this step. If Case 2.2 holds, then according to

the algorithm we must proceed to Case 3, using the 4-tuple $(\widetilde{X}, R_\lambda, \lambda_0, \widetilde{\mathcal{L}})$. Since $\dim R_\lambda = \dim R/\lambda(\mathbb{G}_m) \leq \dim R \leq \dim H$ where equality may hold as $\lambda(\mathbb{G}_m)$ acts trivially on X, we are applying Case 3 to a 4-tuple for which the dimension of the group is smaller than or equal to the dimension of H.

If $(\star\star\star)$ is satisfied for $(X, H, \lambda, \mathcal{L})$, then we start in Case 3, according to which the 4-tuple must be replaced either by $(\overline{Y^{ss}_{\beta^{(k)}}}, P_{\beta^{(k)}}, \lambda_{\beta^{(k)}}, \mathcal{L}|_{\overline{Y^{ss}_{\beta^{(k)}}}})$ (Case 3.1) or by $(\overline{Y^{ss}_{\beta^{(k)}}}, P_{\beta^{(k)}}/\lambda_{\beta^{(k)}}(\mathbb{G}_m), \lambda_0, \mathcal{L}|_{\overline{Y^{ss}_{\beta^{(k)}}}})$ (Case 3.2) for some $k \geq 1$. Note that by construction, we have that $\dim P_{\beta^{(k)}} \leq \dim R/\lambda(\mathbb{G}_m) \leq \dim R \leq \dim H$, where equality may hold throughout (since $\lambda(\mathbb{G}_m)$ could be trivial). However, since $\lambda_{\beta^{(k)}}(\mathbb{G}_m)$ is a non-trivial one-parameter subgroup, we have that $\dim P_{\beta^{(k)}}/\lambda_{\beta^{(k)}}(\mathbb{G}_m) < \dim P_{\beta^{(k)}}$. Thus in Case 3.1 we obtain a new 4-tuple for which the dimension of the group is smaller than or equal to the dimension of H, while in Case 3.2 the dimension of the group becomes strictly smaller.

The only way in which the algorithm might not terminate is if we could travel between Cases 1, 2, and 3 infinitely many times. In Figure .7, this corresponds to circling through Cases 1, 2, and 3 via loops formed by the arrows coloured green (indicating that the dimension of the group decreases) and blue (indicating that the dimension of the group decreases or stays the same). However, as the diagram shows, every such closed loop contains at least one green arrow. Thus each time we travel through a loop, the dimension of the group must decrease. If we did this infinitely often, we would eventually reach a case for which the starting 4-tuple $(X', H', \lambda', \mathcal{L}')$ satisfies $H' = \{e\}$. Since $\lambda'(\mathbb{G}_m)$ is then trivial, and since $X'^{ss, R'/\lambda'(\mathbb{G}_m)} = X'^{ss, \{e\}} = X' \neq \emptyset$, Case 2 applies. Note that we also then have $X'^{s, R'/\lambda'(\mathbb{G}_m)} = X'^{ss, R'/\lambda'(\mathbb{G}_m)} \neq \emptyset$, so the reductive blow-up is X' itself, and hence $\widetilde{X'}^{ss, R'/\lambda'(\mathbb{G}_m)} \neq \emptyset$. Thus Case 2.1 would be automatically satisfied, and so the algorithm would terminate.

□

Bibliography

[Bal08] E. Baldwin. A GIT construction of moduli spaces of stable maps in positive characteristic. *Journal of the London Mathematical Society*, 78(1):107–124, 2008.

[BB73] A. Bialynicki-Birula. Some theorems on actions of algebraic groups. *Annals of Mathematics*, 98(3):480–497, 1973.

[BDHK] G. Bérczi, B. Doran, T. Hawes, and F. Kirwan. Geometric invariant theory for graded unipotent groups and applications. *Journal of Topology*, 11(3):826–855.

[BDHK16] G. Bérczi, B. Doran, T. Hawes, and F. Kirwan. Projective completions of graded unipotent quotients. *ArXiv e-prints*, July 2016.

[BFMV14] G. Bini, F. Felici, M. Melo, and F. Viviani. *Geometric Invariant Theory for Polarized Curves*. Lecture Notes in Mathematics. Springer International Publishing, 2014.

[BHJK] G. Bérczi, V. Hoskins, J Jackson, and F. Kirwan. Quotients of unstable subvarieties and moduli spaces of sheaves of fixed Harder-Narasimhan type, 2. *In Preparation*.

[BHK17] G. Bérczi, V. Hoskins, and F. Kirwan. Stratifying quotient stacks and moduli stacks. *To appear in the Proceedings of the 2017 Abel Symposium*, 2017.

[BJK] G. Bérczi, J. Jackson, and F. Kirwan. Variation of non-reductive geometric invariant theory. *To appear in Surveys in Differential Geometry*. To appear.

[BPMN] L. Brambila-Paz, O. Mata, and N. Nitsure. Moduli stacks and moduli schemes for rank 2 unstable bundles. arXiv preprint 0911.2301.

[Bri17a] M. Brion. *Autmorphisms and Endomorphisms of Projective Varieties*. Proceedings of Symposia in Pure Mathematics. American Mathematical Society, 2017.

[Bri17b] M. Brion. Some structure theorems for algebraic groups. *Proc. Symp. Pure Math. Vol. 94.*, September 2017.

[BRV17] D. Bejleri, D. Ranganathan, and R. Vakil. Motivic Hilbert zeta functions of curves are rational. *arXiv preprint arXiv:1710.04198*, 2017.

[DH98] I. V. Dolgachev and Y. Hu. Variation of geometric invariant theory quotients. *Publications Mathématiques de L'IHÉS*, 87(1):5–51, 1998.

[DK07] B. Doran and F. Kirwan. Towards non-reductive geometric invariant theory. *Pure and Applied Mathematics Quarterly*, 3, 04 2007.

[Ebe65] S. Ebey. The classification of singular points of algebraic curves. *Transactions of the American Mathematical Society*, 118:454–471, 1965.

[FGI+05] B. Fantechi, L. Göttsche, L. Illusie, S. Kleiman, N. Nitsure, and A. Vistoli. *Fundamental Algebraic Geometry: Grothendieck's FGA Explained*. Mathematical surveys and monographs. American Mathematical Society, 2005.

[Gie82] D. Gieseker. *Lectures on moduli of curves*, volume 69 of *Tata Institute of Fundamental*

Research Lectures on Mathematics and Physics. Published for the Tata Institute of Fundamental Research, Bombay; Springer-Verlag, Berlin-New York, 1982.

[Gro66] A. Grothendieck. Éléments de géométrie algébrique. IV. Étude locale des schémas et des morphismes de schémas. III. *Inst. Hautes Études Sci. Publ. Math.*, (28):255, 1966.

[Haw15] T. Hawes. *Non-reductive geometric invariant theory and compactifications of enveloped quotients*. PhD thesis, University of Oxford, 2015.

[HH13] B. Hassett and D. Hyeon. Log minimal model program for the moduli space of stable curves: the first flip. *Annals of Mathematics*, 177:1–58, 2013.

[HK12] V. Hoskins and F. Kirwan. Quotients of unstable subvarieties and moduli spaces of sheaves of fixed Harder–Narasimhan type. *Proceedings of the London Mathematical Society*, 105(4):852–890, 2012.

[HL10] D. Huybrechts and M. Lehn. *The geometry of moduli spaces of sheaves*. Cambridge University Press, 2010.

[HN75] G. Harder and M. S. Narasimhan. On the cohomology groups of moduli spaces of vector bundles on curves. *Math. Ann.*, 212:215–248, 1975.

[Hos] V. Hoskins. Stratifications for moduli of sheaves and moduli of quiver representations. arXiv preprint 1407.4057.

[Hos11] V. Hoskins. *Moduli spaces of complexes of sheaves*. PhD thesis, Oxford University, UK, 2011.

[Ish80] S. Ishii. Moduli of subrings of a local ring. *Journal of Algebra*, 67(2):504–516, 1980.

[Ish82] S. Ishii. Global moduli of equisingular complete curves. *Journal of Algebra*, 78(2):255–264, 1982.

[Kem78] G.R. Kempf. Instability in invariant theory. *Annals of Mathematics*, 108(2):299–316, 1978.

[Kir84] F. Kirwan. *Cohomology of Quotients in Symplectic and Algebraic Geometry*. Mathematical Notes - Princeton University Press. Princeton University Press, 1984.

[Kir85] F. Kirwan. Partial desingularisations of quotients of nonsingular varieties and their Betti numbers. *Annals of Mathematics*, 122(1):41–85, 1985.

[LW15] J. Li and X. Wang. Hilbert–Mumford criterion for nodal curves. *Compositio Mathematica*, 151(11):2076–2130, 2015.

[Mar77] M. Maruyama. Moduli of stable sheaves, i. *J. Math. Kyoto Univ.*, 17(1):91–126, 1977.

[Mar78] M. Maruyama. Moduli of stable sheaves, ii. *J. Math. Kyoto Univ.*, 18(3):557–614, 1978.

[MFK94] D. Mumford, J. Fogarty, and F.. Kirwan. *Geometric Invariant Theory*. Springer-Verlag, Berlin, 1994.

[Mor09] I. Morrison. GIT constructions of moduli spaces of stable curves and maps. *Surveys in Differential Geometry*, 14(1):315–370, 2009.

[Mum66] David Mumford. *Lectures on curves on an algebraic surface*. With a section by G. M. Bergman. Annals of Mathematics Studies, No. 59. Princeton University Press, Princeton, N.J., 1966.

[Mum77] D. Mumford. *Stability of projective varieties*. L'Enseignement Mathématique, Geneva, 1977. Lectures given at the "Institut des Hautes Études Scientifiques", Bures-sur-Yvette, March-April 1976, Monographie de l'Enseignement Mathématique, No. 24.

[Nag60] M. Nagata. On the fourteenth problem of Hilbert. In *Proc. Internat. Congress Math.*

1958, pages 459–462. Cambridge Univ. Press, New York, 1960.

[Nag63] M. Nagata. Invariants of group in an affine ring. *J. Math. Kyoto Univ.*, 3(3):369–378, 1963.

[New78] P. E. Newstead. *Introduction to Moduli Problems and Orbit Spaces*. T.I.F.R. Lecture Notes. Springer-Verlag, 1978.

[Nit11] N. Nitsure. Schematic Harder–Narasimhan filtration. *Internat. J. Math*, 22(10):1365–1373, 2011.

[Pop78] V. L. Popov. Algebraic curves with an infinite automorphism group. *Mathematical notes of the Academy of Sciences of the USSR*, 23(2):102–108, 1978.

[PPS+94] V.L. Popov, A.N. Parshin, T.A. Springer, I.R. Shafarevich, and E.B. Vinberg. *Algebraic Geometry IV: Linear Algebraic Groups Invariant Theory*. Encyclopaedia of Mathematical Sciences. Springer Berlin Heidelberg, 1994.

[PV94] V. L. Popov and E. B. Vinberg. *Invariant Theory*, pages 123–278. Springer Berlin Heidelberg, Berlin, Heidelberg, 1994.

[Rei89] Z. Reichstein. Stability and equivariant maps. *Inventiones mathematicae*, 96(2):349–383, Jun 1989.

[Rob90] P. Roberts. An infinitely generated symbolic blow-up in a power series ring and a new counterexample to Hilbert's fourteenth problem. *Journal of Algebra*, 132(2):461 – 473, 1990.

[Rud97] A. Rudakov. Stability for an abelian category. *Journal of Algebra*, 197(1):231 – 245, 1997.

[Sai13] A.Z. Saiz. GIT characterizations of harder-narasimhan filtrations. caracterizaciones git de filtraciones de harder-narasimhan. Tesis inédita de la Universidad Complutense de Madrid, Facultad de Ciencias Matemáticas, Departamento de Álgebra, leída el 07/06/2013, Julio 2013.

[Ser12] J.P. Serre. *Algebraic Groups and Class Fields*. Graduate Texts in Mathematics. Springer New York, 2012.

[Sha77] S. Shatz. The decomposition and specialization of algebraic families of vector bundles. *Compositio Math.*, 35:163–187, 1977.

[Sim94] C. T. Simpson. Moduli of representations of the fundamental group of a smooth projective variety. ii. *Publications mathématiques de l'IHÉS*, 80(1):5–79, 1994.

[Tha96] M. Thaddeus. Geometric invariant theory and flips. *J. Amer. Math. Soc.*, 9(3):691–723, 1996.

[Wei32] R. Weitzenböck. Über die invarianten von linearen gruppen. *Acta Math.*, 58:231–293, 1932.

[ZT06] O. Zariski and B. Teissier. *Problème des modules pour les branches planes*. University lecture series. American Mathematical Society, 2006.