Non-Reductive Geometric Invariant Theory and its applications to Higgs bundles

Eloise Hamilton
Balliol College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity Term 2020

Acknowledgements

Above all I would like to thank my first supervisor Frances Kirwan. It has been a tremendous honour and privilege to study under her supervision. I am immensely grateful for her generosity with her time, for her patience, for her guidance, and for making my research these last four years so enjoyable and exciting. I cannot thank her enough. I would also like to thank my second supervisor Paul Ziegler for his helpful advice and for always making himself available to answer my questions.

My research has greatly benefitted from discussions with Gergely Bérczi, Joshua Jackson, Victoria Hoskins and Yuuji Tanaka, and I would like to thank them for their ongoing help and support.

Amongst the many other mathematicians who have helped me during my DPhil, I am particularly grateful to Steven Rayan for hosting me for a research visit at the University of Saskatchewan, and to Laura Schaposnik for inviting me to a number of conferences.

The Oxford Mathematical Institute has provided a wonderful working environment. My time in the Mathematical Institute would not have been the same without all the friends I made there. Thank you also to the Graduate Studies Administrator Sandhya Patel for her efficiency and kindness.

Balliol College has played an important role in my life outside the Mathematical Institute: I am grateful to the Dining Hall staff, to the Balliol College Rowing Club, to the Chapel Choir and to the Balliol MCR. They have made my time in Oxford thoroughly enjoyable and I am thankful for the wonderful friends I made there. Thank you also to The Oxford University Lawn Tennis Club and to my teammates for the good times we shared on and off the courts.

Thank you to David Smyth and Timothy Trudgian who, after encouraging me to apply to Oxford when I was at the Australian National University, have continued to provide support throughout my time here.

A special thank you to the John Monash Foundation for believing in me and for providing the scholarship which allowed me to study in Oxford. I also gratefully acknowledge the Centre for Quantum Geometry of Moduli Spaces at Aarhus University and the Danish National Research Foundation for funding me for the last part of my DPhil.

Finally, I am indebted to my family for their constant support throughout my studies.

Abstract

Geometric Invariant Theory (GIT) is a powerful theory for constructing and studying the geometry of moduli spaces in algebraic geometry, as quotients of a parameter space by a linear group action. However GIT only applies to classification problems involving reductive group actions, and only produces a moduli space for the subclass of 'stable' objects. Non-Reductive GIT was developed over the last ten years to overcome these limitations.

The aim of this thesis is firstly to describe a systematic way of constructing and studying the geometry of quotients using both classical and Non-Reductive GIT, and secondly to show how these methods can be applied to the classification problem for Higgs bundles.

In the first part we describe an algorithm which, given the linear action of a linear algebraic group with so-called 'internally graded unipotent radical' on a projective variety, uses classical and Non-Reductive GIT to produce a non-empty open subset of the variety admitting a quasi-projective geometric quotient, together with an explicit projective completion of this quotient. Assuming the initial variety is smooth, we describe a procedure for calculating the Poincaré series of the resulting projective completion, which first requires showing that it can have at worst finite quotient singularities. We then obtain as a result a formula which can be viewed as the non-reductive analogue of the formula for the Poincaré series of the partial desingularisation of classical GIT quotients.

In the second part we show how Non-Reductive GIT can be used to refine two instability stratifications of the stack of Higgs bundles: the Higgs Harder-Narasimhan and the Harder-Narasimhan stratifications. These refined stratifications satisfy the property that each stratum admits a quasi-projective coarse moduli space with an explicit projective completion. In the rank 2 case we provide a complete moduli-theoretic interpretation of these refined stratifications and study the geometry of the corresponding moduli spaces.

Contents

Introduction

Geometric Invariant Theory, classical and Non-Reductive. Classification problems in algebraic geometry typically exhibit 'moduli'[1], in the sense that discrete invariants are not enough to distinguish between all the isomorphism classes of the objects to be classified. In other words, isomorphism classes can vary in a continuous way. Moduli spaces can be thought of as solutions to such classification problems: a moduli space is a variety (or a scheme, or even an algebraic space) whose points correspond to isomorphism classes of objects, and whose geometry captures the way in which the objects can vary continuously. As a result, information about the objects can be extracted more efficiently by studying the geometry of the moduli space instead, and a wide range of algebro-geometric tools is available for this purpose.

While moduli spaces are the key to solving classification problems in algebraic geometry, constructing them is generally a difficult task. Significant progress was made in the 1960s by Mumford who developed Geometric Invariant Theory (GIT) primarily as a tool for constructing moduli spaces in algebraic geometry [76]. Indeed, given a collection of algebro-geometric objects to classify, it is often possible to construct a variety which parametrises these objects by equipping them with some extra structure. Although the redundancy introduced by this extra structure prevents such a variety from being a moduli space, this redundancy is typically captured by the action of a group, in the sense that two points in the variety correspond to equivalent objects if and only if they lie in the same orbit under the group action. As a result, the problem of constructing a moduli space becomes one of constructing a quotient of the parameter space by the given group action. Moreover, the geometry of moduli spaces constructed in this way can be described in terms of the geometry of their parameter spaces, which is often easier to describe. By providing a framework and method for constructing quotients of varieties

[1] The concept of moduli can be traced back to Riemann's work from the 1850s on the classification of smooth projective curves over $\mathbb{C}$, in which he uses heuristic arguments to conclude that any such curve of genus g is determined up to isomorphism by $3g-3$ complex parameters.

by certain linear group actions, GIT has proven thus far to be one of the most successful theories for constructing and studying moduli spaces in algebraic geometry.

Given a classification problem which can be reduced to the problem of constructing a quotient for the action of a linear algebraic group G on a projective variety X, then in order for classical GIT to apply, the group G must be reductive. Under this assumption, and provided the action can be linearised[2], GIT produces open subsets $X^s \subseteq X^{ss} \subseteq X$ of X, called the stable and semistable locus respectively, together with a projective variety $X/\!/G$ which is a good quotient[3] for the action of G on X^{ss}, and which restricts to a geometric quotient[4] $X^s/G \subseteq X/\!/G$ for the action of G on X. The geometric quotient X^s/G then represents a quasi-projective moduli space for objects parametrised by stable points in X, with a projective completion given by $X/\!/G$. It is important to note that objects which are parametrised by unstable points of X (i.e. points which are not semistable) cannot be classified using classical GIT. Moreover, if the group is not reductive, then classical GIT simply does not apply. This limits the scope of classification problems to which classical GIT can be applied, a limitation which has led to the development in recent years of a generalisation of GIT, called Non-Reductive GIT [23, 6, 7, 8].

As its name suggests, Non-Reductive GIT applies to the actions of groups which are not necessarily reductive. More precisely, it applies to the actions of linear algebraic groups H containing a semi-direct product $\widehat{U} := U \rtimes \mathbb{G}_m$ where $\mathbb{G}_m$ is a one-parameter subgroup of H acting with positive weights on the Lie algebra of the unipotent radical U of H. Given the linear action of such a group H on a projective variety X, provided an assumption analogous to the assumption that 'semistability coincides with stability' is satisfied for the action of $\widehat{U}$ on X, an effective analogue of classical GIT can be obtained (see [8]). That is, Non-Reductive GIT similarly produces open subsets[5] $X^{\text{`}s\text{'}} \subseteq X^{\text{`}ss\text{'}} \subseteq X$ of X such that $X^{\text{`}ss\text{'}}$ admits a projective good quotient $X/\!/H$, and this quotient restricts to a geometric quotient $X^{\text{`}s\text{'}}/H$ for the action of

[2] This means that the action of G lifts to an action on an ample line bundle on X.

[3] A *good quotient* of a variety Y by a group H is a variety Z and a morphism $\pi : Y \to Z$ such that:

(i) π is H-invariant, surjective and affine;
(ii) the pull-back map $\phi^* : \mathcal{O}_Z \to \phi_*\mathcal{O}_Y$ induces an isomorphism of sheaves $\mathcal{O}_Z \cong (\phi_*\mathcal{O}_Y)^H$ where $(\phi_*\mathcal{O}_Y)^H(U) = \mathcal{O}_Y(\phi^{-1}(U))^H$ for any open subset $U \subseteq Z$;
(iii) if W_1 and W_2 are closed H-invariant subsets of Y, then $\pi(W_1)$ and $\pi(W_2)$ are closed and disjoint subsets of Z.

[4] A good quotient is a *geometric quotient* if it is also an orbit space.

[5] We use single quotation marks around the superscripts ^{s} and ^{ss} to indicate stability and semistability in the sense of Non-Reductive GIT, which is subtly different from the corresponding notions in classical GIT, refer to Sections 1.1.2 and 1.1.3.

H on X. The key difference from classical GIT is that this Non-Reductive GIT construction can be reapplied to the complement of the stable locus in X, so that X can be stratified inductively in such a way that each stratum admits a quasi-projective geometric quotient with an explicit projective completion. Thus Non-Reductive GIT can be used to obtain complete solutions to classification problems involving reductive groups (by enabling the classification of unstable objects), and also opens the door to solving new classification problems involving non-reductive groups.

The concept of 'stability'[6] lies at the heart of both classical and Non-Reductive GIT, since it is by restricting to the 'stable' locus (in either the classical or Non-Reductive sense) that geometric quotients can be obtained. But therein also lies the difficulty in applying GIT to classification problems in algebraic geometry: to obtain meaningful results, it is necessary for each application to interpret 'stability' in terms of properties of the objects themselves. This can be very difficult to do in practice [41, Chap 1].

The classification problem for Higgs bundles presents a particularly interesting example from this perspective: the condition of stability arising from classical GIT corresponds, somewhat surprisingly, to the condition of being a solution to the self-duality equations on a Riemann surface (these are equations arising from theoretical physics) [51]. Thus if Non-Reductive GIT can be applied to classify unstable Higgs bundles, just as classical GIT can be used to classify stable Higgs bundles, we can ask whether the resulting notion of 'stability' (in the non-reductive sense, and hence applicable to unstable Higgs bundles) might have a corresponding interpretation in terms of solutions to equations from theoretical physics. Although the question is not addressed in this thesis, it can be viewed as a motivation for the application of Non-Reductive GIT to the classification problem for Higgs bundles.

Non-Reductive GIT for Higgs bundles. The origin of Higgs bundles can be traced back to elementary particle physics, and more specifically to its mathematical formulation as a Yang-Mills theory. Indeed, Higgs bundles were first introduced in 1987 by Hitchin as solutions to the so-called self-duality equations on a Riemann surface (these equations correspond to the dimensional reduction of a special class of the Yang-Mills equations on Euclidean 4-space) [51].

[6] Again we use the single quotation marks to cover stability in both the classical and Non-Reductive sense.

The physical importance of the Yang-Mills equations is that they describe the various forces in the standard model of physics, thanks in part to the incorporation of a 'Higgs field' into Yang-Mills theory[7] [60, §1]. The term Higgs bundle was coined in reference to this Higgs field: a Higgs bundle consists of a vector bundle together with a so-called Higgs field which, interpreted appropriately, can be made to correspond to the Higgs field of particle physics [100, Rk 2.1].

A solution to the self-duality equations on a compact Riemann surface Σ consists of a connection A on a principal G-bundle P over Σ together with a $(1,0)$-form Φ, called a 'Higgs field', on Σ with values in the complex Lie algebra bundle of P, satisfying the equations

$$F(A) = -[\Phi, \Phi^*] \text{ and}$$
$$\overline{\partial}_A \Phi = 0.$$

Here G is a compact Lie group, and Hitchin studies the particular case where $G = \mathrm{SU}(2)$ or $G = \mathrm{SO}(3)$. A solution to the self-duality equations then has the following algebro-geometric interpretation: it corresponds to a holomorphic rank 2 vector bundle E over Σ together with a holomorphic section $\phi \in H^0(\Sigma, \operatorname{End} E \otimes K_\Sigma)$, where K_Σ is the canonical line bundle over Σ. Such a pair (E, ϕ) is said to be stable if for any ϕ-invariant proper subbundle $F \subseteq E$ the inequality on slopes $\mu(F) < \mu(E)$ is satisfied. Hitchin proves, in a result which is key to the remainder of his paper, that a pair (E, ϕ) is gauge equivalent to a solution to the self-duality equations if and only if it is stable. Using this correspondence, the moduli space of solutions to the self-duality equations can equivalently be viewed as the moduli space of stable Higgs bundles.

Hitchin's results triggered the study in a more general algebro-geometric setting of pairs (E, ϕ) on a compact Riemann surface Σ with a fixed line bundle $L \to \Sigma$, where E is a vector bundle over Σ of any rank and degree and ϕ is a global holomorphic section of the vector bundle $\operatorname{End} E \otimes L$. Such pairs are called L-twisted Higgs bundles. A moduli space $\mathcal{M}^{ss}_{r,d}(\Sigma, L)$ for semistable Higgs bundles of a fixed rank and degree was first constructed by Nitsure in 1991, using GIT. It is a quasi-projective scheme containing the moduli space $\mathcal{M}^{s}_{r,d}(\Sigma, L)$ of stable Higgs bundles as an open subscheme [81].

[7] It is this Higgs field which leads to the theoretical prediction of the Higgs boson; its existence was confirmed experimentally in 2012 [37].

The moduli space $\mathcal{M}^s_{r,d}(\Sigma, L)$ is rich in geometric structure from many different mathematical perspectives: holomorphic, differential geometric, symplectic and Riemannian [51], in particular when L is chosen to be the canonical line bundle K_Σ (the resulting objects are simply called Higgs bundles). The study of this moduli space over the past thirty years has led to important new developments in theoretical and mathematical physics: the moduli space of stable Higgs bundles is a hyperkähler manifold with the structure of a completely integrable system[8], a feature which links Higgs bundles to mirror symmetry and Langlands duality [48, 21], as well as to super-symmetric gauge theories and their wall-crossing phenomena [22, 79, 30, 73].

The importance of Higgs bundles is not limited to their impact on physics; the study of Higgs bundles has also led to new mathematical advances. Indeed, the Non-Abelian Hodge Theorem establishes a correspondence between Higgs bundles and representations of fundamental groups of surfaces, and by doing so gives a powerful perspective for the study of such representations [33, 32]. Moreover, the Hitchin integrable system has become an important object of study in a modern branch of representation theory called the Geometric Langlands programme. In particular, the Hitchin system played a crucial role in Ngô's proof in 2010 of the Fundamental Lemma[9], which is central to the Langlands programme.

In light of the rich structure of this moduli space, and of the recent developments in Non-Reductive GIT allowing the construction of moduli spaces for 'unstable' objects, we can therefore ask whether moduli spaces involving unstable Higgs bundles can be constructed, and if so whether they admit a similarly rich structure to the moduli space of stable Higgs bundles.

Aim of the thesis. The aim of this thesis is two-fold. Firstly it is to describe a systematic way of constructing and studying the geometry of moduli spaces in algebraic geometry using both classical and Non-Reductive GIT. Secondly, it is to show how these methods can be applied to the classification problem for Higgs bundles, through the construction of new moduli spaces for Higgs bundles which are not necessarily semistable, and through the study of the geometry of these moduli spaces.

[8] See [22, §2.3] for an introduction to the theory of algebraically completely integrable Hamiltonian systems.
[9] Ngô was awarded the Fields Medal in 2010 for his proof of the Fundamental Lemma.

Structure of the thesis. This thesis is in two parts, to reflect each of the above two aims. Each part is divided into two chapters (Chapters 1 and 2 for Part I and Chapters 3 and 4 for Part II), with the two chapters in Part II mirroring the two chapters in Part I:

- Chapter 1 describes how to construct moduli spaces using GIT (both classical and Non-Reductive), while Chapter 3 constructs new moduli spaces for Higgs bundles using Non-Reductive GIT;
- Chapter 2 focuses on describing the geometry of the classical and Non-Reductive GIT quotients introduced in Chapter 1, while Chapter 4 studies the geometry of the new moduli spaces of Higgs bundles (in the rank 2 case) constructed in Chapter 3 as Non-Reductive GIT quotients.

Both parts of the thesis begin with an overview of their content, and each chapter has its own introduction which includes a detailed summary of the content of the chapter.

Summary and results of the thesis. Chapter 1 begins with a review of classical GIT and its recent generalisation to Non-Reductive GIT, which we split into two cases: when the condition that semistability coincides with stability is satisfied (in both the classical and non-reductive sense), and when it is not. In both cases, a geometric quotient can be obtained for the action of the group on an open subset of the variety. When semistability coincides with stability, this geometric quotient is also projective. When it does not, a sequence of blow-ups can be performed to obtain a variety for which semistability does coincide with stability, and therefore for which a projective geometric quotient can be found, itself a projective completion of the original geometric quotient. Nevertheless, in both cases the application of classical or Non-Reductive GIT can result in an empty semistable locus and hence an empty quotient, which is not helpful for constructing moduli spaces. For this reason we conclude the chapter by describing the 'Projective Completion algorithm'[10], which is an explicit and streamlined version of a procedure used in [11] to find non-empty geometric quotients with projective completions for linear actions of groups with internally graded unipotent radicals. Note that the diagrams and the Projective Completion algorithm presented in this chapter are contained in the Appendix which I wrote for the paper [8].

[10] We note that this algorithm is not a priori deterministic.

Chapter 2 shifts from the problem of constructing quotients to that of studying the geometry of these quotients, a particularly important problem when the quotients represent moduli spaces. Again we begin with a review of existing results which establish inductive formulae for the Poincaré series of classical and Non-Reductive GIT quotients, in the case where semistability coincides with stability, and under the assumption that the initial variety is smooth. We then turn to the case where semistability does not coincide with stability; the aim there is to describe the Poincaré series of the projective completion obtained after performing the relevant sequence of blow-ups (from either classical or Non-Reductive GIT). A formula in the reductive case is established in [64], under the assumption that the stable locus is non-empty. A prerequisite for obtaining this formula is to prove that the centres of the blow-ups are smooth (this ensures that the resulting variety is also smooth and therefore that the formulae for the case where semistability coincides with stability apply). In the non-reductive case, if we wish to obtain a formula for the Poincaré series of the projective completion obtained after performing the blow-up construction, then we must similarly prove that the centres of the blow-ups (from Non-Reductive GIT) are smooth. We achieve this in Theorem 2.2.4, and using this result we are able to establish new formulae for computing the Poincaré series of Non-Reductive GIT quotients when semistabiliy does not coincide with stability, under an assumption analogous to that made in the corresponding reductive case (Theorem 2.3.2 and Corollary 2.3.5).

Chapter 3 applies Non-Reductive GIT to construct new moduli spaces for Higgs bundles. To do so, we work with the moduli stack of twisted Higgs bundles of arbitrary rank and degree on a smooth projective curve, which admits two instability stratifications: the Higgs Harder-Narasimhan stratification, determined by the instability type of the Higgs bundle, and the Harder-Narasimhan stratification, determined by the instability type of the underlying vector bundle. We show that both these stratifications can be refined in such a way that each refined stratum admits a quasi-projective coarse moduli space with a projective completion constructed using Non-Reductive GIT (see Theorems 3.2.2 and 3.3.5). In particular, the application of Non-Reductive GIT produces for any Higgs Harder-Narasimhan type μ (respectively Harder-Narasimhan type τ) a corresponding notion of μ-stability (respectively τ-stability), which coincides with the classical notion of stability when μ (respectively τ) is the trivial Higgs Harder-

Narasimhan (respectively Harder-Narasimhan) type[11]. This extends the well-known result regarding the existence of a quasi-projective coarse moduli space for the substack of stable Higgs bundles, proved by Nitsure in [81] using classical GIT. Nevertheless, the problem of describing these notions explicitly in terms of intrinsic properties of Higgs bundles remains. Progress towards solving this problem is made by providing a complete moduli-theoretic interpretation of the two refined stratifications in the rank 2 case (see Theorem 3.2.9 and Remark 3.3.12), and by obtaining a partial moduli-theoretic interpretation of the refined Harder-Narasimhan stratification for the general rank case (Proposition 3.3.11). We note that the paper [45] is based on the material from this chapter.

Chapter 4 considers the problem of studying the geometry of the moduli spaces of unstable Higgs bundles constructed in Chapter 3, in the case of rank 2 Higgs bundles. We start by considering the case of unstable twisted Higgs bundles of rank 2 on $\mathbb{P}^1$: we explicitly describe the geometry of the moduli spaces (see Theorem 4.1.7) and of the Hitchin fibration for these moduli spaces (see Proposition 4.1.12), thus allowing an explicit comparison with the moduli space of semistable Higgs bundles (see Section 4.1.3). We then turn to the case of curves of arbitrary genus, for which there are two types of moduli spaces to consider given a fixed Higgs Harder-Narasimhan type μ: the moduli space of μ-stable Higgs bundles, and moduli spaces of μ-unstable Higgs bundles. We show that limits at 0 under the Higgs field scaling $\mathbb{C}^*$-action exist for the latter moduli spaces (see Proposition 4.2.2), and we use this property to study the geometry of the moduli spaces and to compute their Poincaré series (see Corollary 4.2.13). For the former moduli space of μ-stable Higgs bundles, for which limits at 0 may not exist, we instead return to the results of Chapter 2. That is, we use the fact that the moduli space of μ-stable Higgs bundles can be constructed as a Non-Reductive GIT quotient to show how the new formulae of Chapter 2 can be applied to compute the Poincaré series of a partial compactification of this moduli space (see Section 4.3.3).

Conventions. In this thesis we work with algebraic varieties defined over an algebraically closed field k of characteristic 0, although we often set $k = \mathbb{C}$ for concreteness. By variety we mean a reduced scheme of finite type over k; in particular we do not assume that a variety is irreducible, and we will make this assumption explicit where necessary.

[11] By trivial Higgs Harder-Narasimhan type we mean the type associated to semistable Higgs bundles. By trivial Harder-Narasimhan type we mean the type associated to Higgs bundles with semistable underlying vector bundle.

Part I

Construction and cohomology of quotients using Non-Reductive Geometric Invariant Theory

Overview

The aim of Part I is to describe a systematic way of constructing and studying the geometry of moduli spaces in algebraic geometry, using classical Geometric Invariant Theory (GIT) and its recent generalisation to Non-Reductive GIT. Indeed, many classification problems in algebraic geometry can be set up in such a way that isomorphism classes of the objects correspond to orbits for a group action on a suitable parameter space. Chapters 1 and 2 describe the two reasons why this correspondence is useful for solving classification problems.

The first reason is that, provided the group action can be linearised[12], GIT (in either its classical or Non-Reductive variant) can then be used to construct a moduli space for the classification problem, together with a projective completion of this moduli space. The purpose of Chapter 1 is to describe how this can be done in general. After summarising the main results of classical and Non-Reductive GIT, we show how the constructions from the two theories can be combined into an algorithm (not a priori deterministic), called the Projective Completion algorithm. Given the linear action of a linear algebraic group on an irreducible projective variety (satisfying the condition of footnote 12), this algorithm produces a non-empty open subset admitting a geometric quotient, together with a projective completion of this quotient.

The second reason is that, if a moduli space for a classification problem can be constructed as a GIT quotient, again either in the classical or Non-Reductive sense, then many tools become available for studying its geometry. These tools include formulae for computing the Poincaré series of the quotient in terms of information about the parameter space, and this space is often easier to describe. The purpose of Chapter 2 is to describe existing formulae for computing the Poincaré series of classical and Non-Reductive GIT quotients, as well as to establish new formulae in the case where semistability does not coincide with stability for non-reductive groups.

[12] If the group is not reductive, then we need the additional assumption that the group has an internally or externally graded unipotent radical (see Section 1.1 or [8]).

Chapter 1

Constructing quotients for linear algebraic group actions

1.0 Introduction

Classical GIT. The main theorem of GIT applies to the linear action of a reductive linear algebraic group G on a projective variety X. Linearity of the action means that the action lifts to an action on an ample line bundle on the variety, which forms part of the initial data. Given this data, GIT produces an open subset X^{ss} of points in X, called the semistable locus, together with a G-equivariant morphism from this locus X^{ss} to a projective variety $X//G$, called the GIT quotient. In addition to being a categorical quotient, this morphism satisfies the property that two orbits are identified in the GIT quotient if and only if their closures meet in the semistable locus (more precisely, the morphism is a good quotient for the action of G on X^{ss}). Thus, if the orbits are not all closed, the GIT quotient may not be an orbit space for the action of G on X^{ss}. However, GIT also produces an open subset $X^s \subseteq X^{ss}$, called the stable locus, such that X^s/G is a quasi-projective orbit space, with $X//G$ as a projective completion (provided X^s is non-empty).

It is important to note that this theorem is not just an existence result: the GIT quotient can be constructed explicitly as the projective variety associated to the algebra of invariants (reductivity of the group ensures that this algebra is finitely generated), and the stable and semistable loci can be computed combinatorially using the so-called Hilbert-Mumford criteria. It is thanks to this property that GIT is an effective tool for constructing moduli spaces in algebraic geometry.

Nevertheless, in order to obtain quotients, GIT requires restricting to the semistable locus and leaving out the complement, called the unstable locus. GIT is not suited to constructing quotients for points in the unstable locus. Fundamentally this is because constructing quotients for the unstable locus is equivalent to constructing a quotient involving a non-reductive group action, a task to which GIT cannot be applied. This shortcoming of GIT was one of the motivations behind the development over the last ten years of a generalisation of GIT, called Non-Reductive GIT. As its name suggests, Non-Reductive GIT applies to actions by groups which are not necessarily reductive.

Non-Reductive GIT. If a linear algebraic group H acts linearly on a projective variety X, and H is not reductive, then many important properties which underpin classical GIT fail: the algebra of invariants may not be finitely generated, and even if it is, the map from the semistable locus X^{ss} to the GIT quotient may not be a good quotient for the action of H on X^{ss}, or even surjective.

In [8], it is shown that provided H has an 'internally graded' unipotent radical and that the action of H on X satisfies a certain condition, then the above problems do not occur and results analogous to those of classical GIT can be obtained. That is, provided the linearisation is amended in a suitable way, then the invariants are finitely generated and the resulting projective quotient $X//H$ is a good quotient for the action of H on an open semistable locus X^{ss}. This semistable locus contains as an open subset a stable locus X^s such that X^s/H is a quasi-projective orbit space with $X//H$ as a projective completion. Moreover, the stable and semistable loci can be described explicitly in terms of Hilbert-Mumford-type criteria. If the condition for the action of H on X is not satisfied, then a sequence of H-equivariant blow-ups can be performed on X to obtain a projective variety $\widehat{X}$ with a linear action of H for which the condition is satisfied, and to which the results of Non-Reductive GIT therefore apply.

In contrast to classical GIT, Non-Reductive GIT also applies to the action of H on the complement of the semistable locus. By induction, the variety X can therefore be stratified in such a way that each stratum admits a quasi-projective orbit space with an explicit projective completion, for which X^s represents the open stratum provided it is non-empty. The first step to obtaining the stratification is therefore to construct, given the action of a linear algebraic group H on a projective variety X, a non-empty open subset admitting a quasi-projective orbit space with an explicit projective completion. This procedure can be described in the form of an algorithm, called the Projective Completion algorithm.

An algorithmic procedure. Given the data consisting of the linear action of a linear algebraic group H with internally graded unipotent radical on a projective variety X, the Projective Completion algorithm combines constructions from classical and Non-Reductive GIT to produce a non-empty open subset of X admitting a quasi-projective orbit space, together with a projective completion of this orbit space which is itself obtained as a classical or Non-Reductive GIT quotient for the action of H on a possibly different projective variety. By induction, this algorithm produces a stratification of the variety X, called its Non-Reductive GIT H-stratification (introduced in [9]), satisfying the property that each stratum admits a quasi-projective orbit space with an explicit projective completion.

If a classification problem admits a variety X which parametrises the objects equipped with some additional data, in such a way that the redundancy is captured by the action of a group H as above, then the non-reductive GIT stratification of X can be thought of as a complete solution to the classification problem, once it is interpreted in terms of properties of the objects parametrised by X. Indeed, every object lies in one of the strata, with each having an associated quasi-projective coarse moduli space with an explicit projective completion. We note that in contrast to the classical GIT case, the projective completion does not a priori have an obvious modular interpretation, since it is constructed not as a quotient of X but as a quotient of a different variety, determined by the Projective Completion algorithm.

Structure of this chapter. The aim of this chapter is to show how constructions from classical and Non-Reductive GIT can be combined to produce a systematic way of constructing non-empty quasi-projective geometric quotients with explicit projective completions, associated to the action of certain linear algebraic groups on projective varieties, a procedure which we call the Projective Completion algorithm. In Section 1.1 we summarise the main results of classical and Non-Reductive GIT in the case where 'semistability coincides with stability' (in both the classical and non-reductive sense). In Section 1.2 we describe the results when this condition is not satisfied, and for which blow-up constructions are required to obtain a variety where the condition is satisfied. These results are combined in Section 1.3 into an algorithm, called the Projective Completion algorithm, which provides a systematic way of constructing non-empty quotients of varieties by linear algebraic groups with internally graded unipotent radical, together with an explicit projective completion. This algorithm leads by induction to the definition of non-reductive GIT stratifications, which correspond to stratifications of projective varieties admitting a linear action by a linear algebraic group (either reductive or with internally graded unipotent radical) and satisfying the property that each stratum admits a quasi-projective geometric quotient with an explicit projective completion.

1.1 When semistability coincides with stability

In this section we summarise the main results of classical and Non-Reductive Geometric Invariant Theory (GIT). As we will see, given the linear action of a reductive group on a projective variety, GIT produces an open subset of the variety, called the semistable locus, which admits a good quotient, called the GIT quotient, that is itself a projective variety. If the assumption that 'semistability coincides with stability' is satisfied, then the GIT quotient is in fact an orbit space for the action of the group on the semistable locus. The main result of [8] is that for a particular class of non-reductive groups (linear algebraic groups with so-called internally graded unipotent radicals), provided an assumption analogous to the assumption in classical GIT that semistability coincides with stability is made, then an effective analogue of classical GIT can be obtained. That is, given the linear action of a group in this class, Non-Reductive GIT produces an open subset of the variety which admits an orbit space that has the structure of a projective variety, as in the case of classical GIT when semistability coincides with stability. Section 1.2 will address the case when this condition is not satisfied.

In Section 1.1.1 we recall the main results and features of classical GIT as in the setting of [76]. In Section 1.1.2 we turn to a special class of non-reductive groups: semi-direct products of unipotent groups with a multiplicative group acting with positive weights on their Lie algebra via the adjoint action. We present the results of [8] which show that results completely analogous to those of classical GIT can be obtained for such groups, provided an assumption analogous to the assumption that semistability coincides with stability is made. Groups of this form are the building blocks of the groups to which Non-Reductive GIT applies: linear algebraic groups with internally graded unipotent radical. In Section 1.1.3 we turn to this more general case and describe the results of [8] which show that GIT can effectively be generalised to the action of such groups, again provided that an assumption analogous to the assumption that semistability coincides with stability is made.

1.1.1 GIT for reductive groups

First introduced by Mumford in the 1960s, Geometric Invariant Theory (GIT) is a powerful tool for constructing and studying quotients of varieties by reductive linear algebraic group actions[1] [76]. In this section we summarise the main results and features of classical GIT; we introduce the GIT-instability stratification and finally we justify why an analogue of GIT for a certain class of non-reductive groups is needed to construct quotients for the GIT-unstable strata.

Initial data. Let X be an irreducible[2] projective variety[3] over an algebraically closed field[4] k of characteristic 0 and let G be a reductive group acting linearly on X with respect to an ample line bundle $L \to X$. By replacing L by a higher power if necessary, we can assume without loss of generality that it is very ample so that $X \subseteq \mathbb{P}(H^0(X, L)^\vee) =: \mathbb{P}^n$ and the action of G on X extends to an action on $\mathbb{P}^n$ via a representation $G \to \mathrm{GL}(n+1, k)$.

Semistability, Mumford-stability and stability. The action of G on L induces an action of G on the homogeneous coordinate ring $k[X] := \bigoplus_{k \geq 0} H^0(X, L^{\otimes k})$ of X, so we can consider the ring of invariant sections $k[X]^G$. A point $x \in X$ is:

(i) *semistable* if there exists an invariant section $\sigma \in k[X]^G$ such that $\sigma(x) \neq 0$;

(ii) *Mumford-stable*[5] if in addition to being semistable (with non-vanishing section say σ), the

[1] As noted by Mumford, constructing moduli spaces for various types of classification problems in algebraic geometry often reduces to a 'special and highly non-trivial case' of the problem of constructing quotients of varieties by group actions [76].

[2] We make the assumption of irreducibility only for convenience. For example, it ensures that the geometric quotient X^s/G of the stable locus is dense (since it is open) inside the GIT quotient $X/\!/G$, which we can therefore call a projective completion of X^s/G. If X is not irreducible, then applying GIT to each irreducible component produces the same result as applying GIT directly to X. This is because reductive groups satisfy the property that invariants always extend, so that if Y is a projective subvariety of X, then the semistable locus in Y for the restricted linearisation coincides with the intersection of Y with the semistable locus in X. The situation is more subtle for Non-Reductive GIT, since invariants for non-reductive groups do not necessarily extend (see Remark 1.1.6).

[3] We work with varieties instead of schemes again for convenience rather than necessity. If the scheme is not reduced then the result of applying GIT is the same as applying GIT to the associated reduced scheme.

[4] We will often identify k with the complex numbers $\mathbb{C}$ in the remainder of the thesis, for the sake of familiarity, although the results hold true for any field k of characteristic 0.

[5] Such points are called 'stable' in [76], but this terminology is no longer standard.

action of G on X_σ, the non-vanishing locus of σ, is closed;

(iii) *stable*[6] if in addition to being Mumford-stable, its stabiliser group is finite.

The set of semistable (respectively Mumford-stable, stable) points defines an open subset $X^{ss} \subseteq X$ (respectively $X^{Ms} \subseteq X$, $X^s \subseteq X$), and we have a chain of inclusions

$$X^s \subseteq X^{Ms} \subseteq X^{ss} \subseteq X.$$

The main theorem of classical GIT concerns the construction of quotients for the linear action of G on each of these open subsets.

Main theorem of classical GIT. Since G is reductive, by Nagata's theorem the algebra of invariants $k[X]^G$ is finitely generated [78]. Thus we can define the projective variety

$$X/\!/G := \operatorname{Proj} \bigoplus_{k \geq 0} H^0(X, L^{\otimes k})^G,$$

called the *GIT quotient* for the linear action of G on X. The inclusion of the invariants into the coordinate ring $k[X]$ induces a rational map $q_G : X \dashrightarrow X/\!/G$ which satisfies the following properties:

(i) its restriction to the semistable locus is a morphism $q_G : X^{ss} \to X/\!/G$ which is a good quotient[7] for the action of G on X, so that set-theoretically $X/\!/G$ is the quotient of X^{ss} by the equivalence given by $x \sim y$ if and only the closures of the G-orbits of x and y meet in X^{ss} (this equivalence is called *S-equivalence*);

(ii) the restriction of q_G to the stable locus $X^s \to q_G(X^s)$ is a geometric quotient[8] for the action of G on X^s, so that set-theoretically $q_G(X^s) = X^s/G$ (the same statement holds true with X^s replaced by X^{Ms}).

Moreover, the (semi)stable locus and GIT quotient are unchanged if the linearisation is replaced by a tensor power.

It is important to note that the reductivity of G is required not just to define $X/\!/G$, but also to ensure that the map $X^{ss} \to X/\!/G$ is surjective and moreover that it is a good quotient.

[6] Such points are called 'properly stable' in [76].

[7] Given a variety U with an action by a group G, a *good quotient* of U by G is a variety Y and a morphism $\pi : U \to Y$ such that (i) π is G-invariant, surjective and affine; (ii) the pull-back map $\phi^* : \mathcal{O}_Y \to \phi_* \mathcal{O}_U$ induces an isomorphism of sheaves $\mathcal{O}_Y \cong \phi_*(\mathcal{O}_U)^G$; (iii) if W_1 and W_2 are closed and disjoint G-invariant subsets of U, then $\pi(W_1)$ and $\pi(W_2)$ are closed and disjoint subsets of Y.

[8] Given a variety U with an action by a group G, a *geometric quotient* of X by G is a good quotient $\pi : U \to Y$ which in addition satisfies the property that Y is an orbit space for the action of G on U.

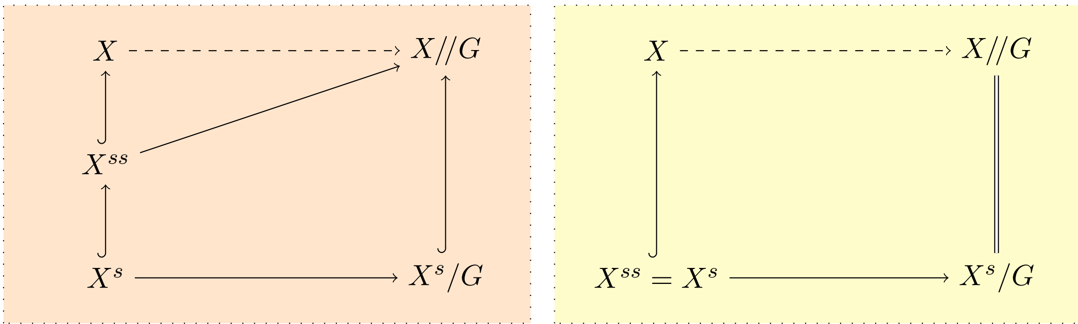

(a) General case (b) When $X^{ss} = X^s$

Figure 1.1: Classical GIT for reductive groups.

This is due to the fact that if G is reductive, then any two disjoint G-invariant subvarieties of X can be separated by an invariant section [23, §1].

Figure 1.1a illustrates the general theorem of classical GIT, while Figure 1.1b illustrates the theorem in the special case where $X^{ss} = X^s$. The situation of Figure 1.1b, namely when semistability coincides with stability, is particularly desirable for two reasons. The first is that in this case, the GIT quotient $X//G$ coincides with the geometric quotient X^s/G, so that the orbit space for the action of G on X^{ss} (i.e. the set-theoretic quotient) has the structure of a projective variety. Projective varieties are one of the most important classes of varieties in algebraic geometry, in particular because they are proper (the algebro-geometric analogue of compactness), and as such a wide range of tools is available to study them. The second advantage of the situation when semistability coincides with stability is that if X is smooth, then the GIT quotient has at worse finite quotient singularities. In particular, the rational cohomology of the GIT quotient $X//G$ is given by the G-equivariant rational cohomology of X^{ss}, which can be computed inductively from the G-equivariant rational cohomology of X, by [64, Thm 8.12] (we will consider the cohomology of GIT quotients in more detail in Chapter 2). In general, strictly semistable points can give rise to much worse singularities and as a result the rational cohomology of the GIT quotient may not coincide with the G-equivariant cohomology of the semistable locus.

Remark 1.1.1 (GIT in the language of stacks)**.** There have been many efforts in recent years to translate classical GIT into the language of algebraic stacks, which is particular suited to the theory of quotients [1, 58, 2, 50, 26, 44, 3]. The first step in this direction is the abstraction to stacks of the notion of a GIT quotient. This was achieved by Alper in [2] which introduces the notion of a good moduli space. A good moduli space of a stack is a morphism from the stack to an algebraic space which abstracts the notion of a GIT quotient, so that if G acts linearly on a projective variety X, then $[X^{ss}/G] \to X//G$ is an example of a good moduli space.

The properties discussed above in the case where semistability coincides with stability can be concisely expressed in the language of stacks. That is, if $X^{ss} = X^s$, then the good moduli space $X//G$ is also a coarse moduli spacefor the quotient stack $[X^{ss}/G]$. Moreover, while in general $[X^{ss}/G]$ is an Artin stack, when $X^{ss} = X^s$ then the quotient stack is a Deligne-Mumford stack. Lastly, if X is smooth, then $[X^{ss}/G]$ is a smooth Deligne-Mumford stack.

To summarise, classical GIT thus provides a way of constructing a projective good quotient for an open subset of X, the semistable locus, such that that this good quotient restricts to a geometric quotient on a (possibly) smaller open subset, the stable (or Mumford-stable) locus. The GIT quotient $X//G$ is then a projective completion of the quasi-projective geometric quotient X^s/G (we recall that we have assumed that X is irreducible, see footnote 2). However, determining what the semistable and stable loci are from their definitions is not straightforward, as this requires computing the invariant sections which is a difficult task in general [19]. The Hilbert-Mumford criterion is an important feature of classical GIT which provides a numerical characterisation of the semistable and stable loci, thus side stepping the need to compute the invariants.

Hilbert-Mumford criterion. The Hilbert-Mumford criterion asserts that

$$X^{(s)s} = \bigcap_{g \in G} gX^{(s)s,T} \tag{1.1}$$

for a fixed maximal torus $T \subseteq G$. In other words, a point is (semi)stable if and only if $g \cdot x$ is (semi)stable for the linear action of T on X for any $g \in G$ (see [23, Rk 3.2.3]). Note that unlike the semistable and stable loci, the Mumford-stable locus is not described by a Hilbert-Mumford-type criterion.

The advantage of this characterisation is that (semi)stability for the linear action of a torus can be determined combinatorially. The action of T on $X \subseteq \mathbb{P}^n$ can be diagonalised so that it acts with weight $\alpha_i \in \mathfrak{t}^\vee$ on each projective coordinate x_i of $\mathbb{P}^n$, where $\mathfrak{t}$ denotes the Lie algebra

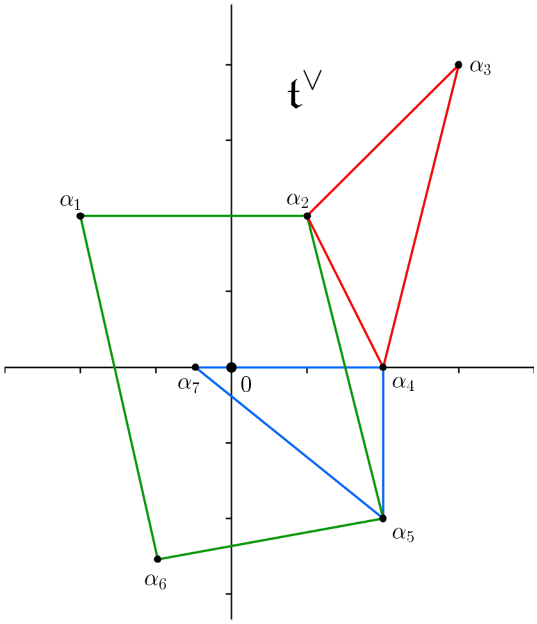

Figure 1.2: Illustration of the Hilbert-Mumford criterion. The small black dots represent the weights $\alpha_i \in \mathfrak{t}^\vee$ determined by the representation $G \to \mathrm{GL}(H^0(X,L)^\vee)$ associated to the linear action of G on X with respect to a very ample line bundle L, and by a choice of maximal torus T. The green polytope represents the convex hull associated to a stable point (since the origin lies in its interior); the blue polytope represents the convex hull associated to a strictly semistable point (since the origin lies on its boundary); the red polytope represents the convex hull associated to an unstable point (since the origin lies outside it).

of T. A point $x = [x_0 : \cdots : x_n] \in X$ is then semistable (respectively stable) for the action of T if and only if the convex hull

$$\mathrm{Conv}\{\alpha_i \mid x_i \neq 0\}$$

contains 0 (respectively contains 0 in its interior). Figure 1.2 illustrates the different cases which can arise.

Remark 1.1.2 (Applications of the Hilbert-Mumford criterion)**.** The Hilbert-Mumford criterion can be used for example to compute the stable and semistable loci for the action of $\mathrm{SL}(n,k)$ on the Grassmannian $G(n,r)$ of r dimensional quotients of k^n, which in turn can be used to determine stability and semistability for vector bundles (see [80, §6 & Chap 5]). Another example of its applications is to the conjugation action on quiver representations [63]. These descriptions would not have been feasible using the method of computing the invariants.

GIT-instability stratification. The above combinatorial characterisation of stability and semistability can be used not only to obtain an explicit description of the semistable and stable loci without having to compute invariants, but also to give structure to the complement of the semistable locus in X, called the *unstable* locus (points which are not semistable are called *unstable*). That is, the unstable locus can be stratified according to 'how unstable' its points are, as measured by the action of a maximal torus. We describe this stratification below, following [64, §12].

By fixing a Weyl group invariant inner product on $\mathfrak{t}$, we can identify $\mathfrak{t}$ with its dual $\mathfrak{t}^\vee$, and thus identify a finite subset of weights $\{\alpha_i\}_{i\in I}$ in $\mathfrak{t}^\vee$ as points in $\mathfrak{t}$. We let Δ_I denote the convex hull of the weights α_i in $\mathfrak{t}$. Given a point $x = [x_0 : \cdots : x_n] \in \mathbb{P}(H^0(X, L)^\vee)$, we define $I_x := \{i \in \{0, \ldots, n\} \mid x_i \neq 0\}$ and let $\Delta_x := \Delta_{I_x}$. For $x \in X$, we call the weights in the set $\{\alpha_i\}_{i\in I_x}$ the *weights of* x. We fix a positive Weyl chamber $\mathfrak{t}_+ \subseteq \mathfrak{t}$ and let $\mathcal{B}$ denote the set of all $\beta \in \mathfrak{t}_+$ such that β is the closest point to the origin of Δ_x for some $x \in X$. Given $\beta \in \mathcal{B} \setminus \{0\}$, we define

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

and

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

In other words, H_β is the hyperplane passing through β and perpendicular to the line through β and the origin, while H_β^+ is the half-space cut out by this hyperplane and which does not contain the origin. Figure 1.3 illustrates these definitions.

We define the closed subvariety

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

and the locally closed subvariety

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

of X. In other words, under the inclusion and isomorphism $X \subseteq \mathbb{P}(H^0(X, L)^\vee) \cong \mathbb{P}^n$, the subvariety Z_β consists of the intersection of X with points $[x_0 : \cdots : x_n] \in \mathbb{P}^n$ such that $x_i = 0$ if $\alpha_i \cdot \beta \neq ||\beta||^2$. Similarly, Y_β consists of the intersection of X with points $[x_0 : \cdots : x_n] \in \mathbb{P}^n$ such that $x_i = 0$ if $\alpha_i \cdot \beta < ||\beta||^2$ and $x_i \neq 0$ for some α_i satisfying $\alpha_i \cdot \beta = ||\beta||^2$. From this perspective we see that there is a natural retraction $p_\beta : Y_\beta \to Z_\beta$ given by $x_i \mapsto x_i$ if $\alpha_i \cdot \beta = ||\beta||^2$ and $x_i \mapsto 0$ otherwise (this is well-defined because given $y \in Y_\beta$, its image $p_\beta(y)$

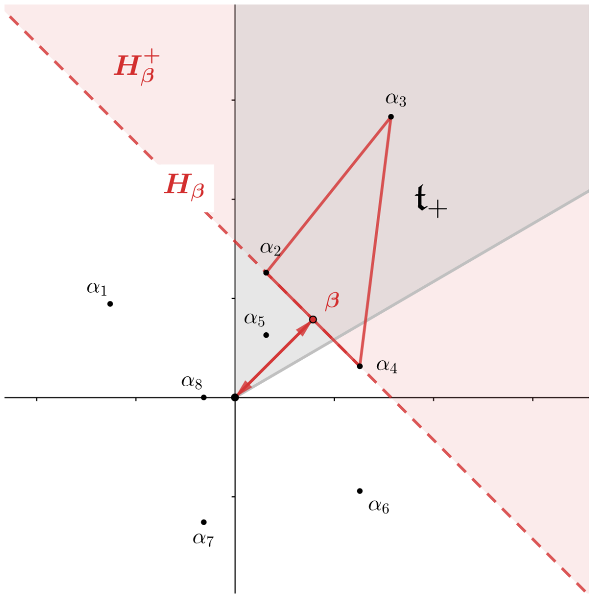

Figure 1.3: Illustration of H_β and H_β^+. The small black dots represent the weights $\alpha_i \in \mathfrak{t}^\vee$ (identified as points in $\mathfrak{t}$ using a Weyl-invariant inner product) determined by the representation $G \to \mathrm{GL}(H^0(X, L)^\vee)$ and a choice of maximal torus T. The grey region represents the chosen positive Weyl chamber $\mathfrak{t}_+$. The red dotted line corresponds to H_β while the half-space shaded in red (including the dotted line) corresponds to H_β^+. The red polytope represents the convex hull associated to an unstable point (since the origin lies outside of it), and $\beta \in \mathfrak{t}_+$ is its closest point to the origin. A point $[x_0 : \cdots : x_n] \in X$ lies in Y_β if the weights α_i corresponding to its non-zero coordinates x_i all lie within H_β^+, with at least one of the weights lying on H_β. The point lies in Z_β if all of the weights lie on the red dotted line.

lies in the closure of the G-orbit of y, and therefore lies in X since X is Zariski-closed in $\mathbb{P}^n$). This is an algebraically locally trivial fibration with contractible fibre (given by an affine space).

Let $\operatorname{Stab}\beta \subseteq G$ denote the stabiliser of $\beta \in \mathfrak{t}^\vee$ under the co-adjoint action of G on the dual $\mathfrak{g}^\vee$ of its Lie algebra $\mathfrak{g}$. Then Z_β is $\operatorname{Stab}\beta$-invariant and the linear action of G on X induces a linear action of $\operatorname{Stab}\beta$ on Z_β. Identifying β as a character of $\operatorname{Stab}\beta$, we let $Z_\beta^{(s)s}$ denote the (semi)stable locus for the above linear action of $\operatorname{Stab}\beta$ on Z_β twisted by the character $-\beta$ (we denote this linearisation by $\mathcal{L}_\beta$, noting that it is ample), and define $Y_\beta^{(s)s} = p_\beta^{-1}\left(Z_\beta^{(s)s}\right)$. Finally, we set

$$S_\beta := GY_\beta^{ss}.$$

The subvarieties S_β form the strata of the GIT-instability stratification of X, given by

$$X = \bigsqcup_{\beta \in \mathcal{B}} S_\beta, \tag{1.2}$$

which satisfies the property that, for each β,

$$\overline{S_\beta} \subseteq S_\beta \sqcup \bigsqcup_{||\beta'||>||\beta||} S_{\beta'}.$$

Moreover, the stratum indexed by 0 coincides with the semistable locus, namely

$$S_0 = X^{ss}.$$

Remark 1.1.3 (GIT instability stratification in the language of stacks)**.** As seen in Remark 1.1.1, the notion of a GIT quotient has been abstracted to stacks through the notion of a good moduli space. More recently, the notion of a GIT-instability stratification has also been generalised to stacks through the notion of a Θ-stratification, introduced by Halpern-Leistner in [44]. The open stratum of a Θ-stratification is the 'semistable' stratum, expected to admit a good moduli space, just as the semistable locus in GIT admits a good quotient; the remaining 'unstable' strata retract onto closed substacks, also expected to admit good moduli spaces, just as the subvariety Y_β^{ss} retracts onto Z_β^{ss} which has a good quotient $Z_\beta /\!/ \operatorname{Stab}\beta$.

Quotients of the unstable strata using classical GIT. The linear action of G on X induces one on the G-equivariant projective completion of S_β given by $\widehat{S}_\beta := G \times_{P_\beta} \overline{Y_\beta^{ss}}$ for any $\beta \neq 0$, where $\overline{Y_\beta^{ss}}$ is the closure of Y_β^{ss} in X. However, the semistable locus $\widehat{S}_\beta^{ss}$ with respect to this linearisation is empty and so the GIT quotient $\widehat{S}_\beta /\!/ G$ is empty. Nevertheless, as described in [53], one can consider a different linearisation on $\widehat{S}_\beta$, called the *canonical linearisation*, for which the semistable locus is non-empty (see [53, Def 2.2.4]). The drawback of considering this different linearisation is that it is in general non-ample, and while GIT is applicable to non-ample linearisations, many results (in particular the Hilbert-Mumford criterion) are only valid if the linearisation is ample [53, §1.5]. Indeed, the GIT quotient constructed using this linearisation coincides with the GIT quotient for the action of $\operatorname{Stab}\beta$ on Z_β with respect to the ample linearisation $\mathcal{L}_\beta$ considered above (see [53, Prop 2.2.8]). It follows that points in Y_β^{ss} are identified in the quotient with their image under the retraction map p_β, which leads to a quotient that is 'too small' since all information about the fibres of the retraction is lost in the quotient. In particular we cannot expect to obtain a geometric quotient, even for an open subset of Y_β^{ss}.

It is for this reason that classical GIT is not suitable for constructing quotients for the unstable strata S_β. This has been one of the motivations behind the development of Non-Reductive GIT, which in particular, and in contrast to classical GIT, can be used to construct quotients for the unstable strata.

From classical to Non-Reductive GIT. The aim of this section is to describe why the problem of constructing geometric quotients for the GIT-unstable strata can be reformulated as a problem involving a non-reductive group action. This requires studying in more detail the structure of the unstable strata S_β for $\beta \neq 0$.

Each $\beta \neq 0$ has a corresponding one-parameter subgroup $\lambda_\beta : \mathbb{G}_m \to T$: it is the one-parameter subgroup associated to the weight vector $q\beta \in \mathfrak{t} \cong \mathfrak{t}^\vee$ where q is the smallest positive rational number such that $q\beta$ has integer entries[9]. We can therefore define for each $\beta \neq 0$ the following parabolic subgroup of G (see [76, Prop 2.6]), which is homotopically equivalent to $\operatorname{Stab}\beta$ (see [64, §14.5]):

$$P_\beta = \left\{ g \in G \;\middle|\; \lim_{t\to 0} \lambda_\beta(t) g \lambda_\beta(t^{-1}) \in G \right\}.$$

By [64, Thm 13.5], for each $\beta \neq 0$ we have that

$$S_\beta \cong G \times_{P_\beta} Y_\beta^{ss}. \tag{1.3}$$

Note that when G is connected, S_β is irreducible if and only if Y_β^{ss} is irreducible.

A consequence of the above isomorphism is that taking a quotient of S_β by G is equivalent to taking a quotient of Y_β^{ss} by P_β. Note that the linear action of G on S_β induces a linear action of P_β on Y_β^{ss}. The key difference is that P_β, as a parabolic subgroup of G, is non-reductive in general, while G is reductive by assumption.

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 (see [59, Lem 4.2.0.2]). Linear algebraic groups satisfying this property are said to have *internally graded unipotent radicals* and results from [8] show that GIT can be successfully extended to allow actions by such groups.

[9] In other words, the one-parameter subgroup λ_β is determined by the property that the derivative of its restriction to S^1, considered as a map $\mathbb{R} \to \mathfrak{t}$, sends 1 to $q\beta$.

In the following two subsections we summarise the results from [8] for semi-direct products of unipotent groups with a positively grading multiplicative group (Section 1.1.2) and for linear algebraic groups with internally graded unipotent radicals (Section 1.1.3), providing diagrams to illustrate them.

1.1.2 GIT for externally graded unipotent groups

Given the linear action of a linear algebraic group on a projective variety, if the group is not reductive then the associated algebra of invariants may not be finitely generated (see [77]), so that there is no obvious analogue of the classical GIT quotient. Moreover, even if the invariants are finitely generated, the induced quotient map is not necessarily a good quotient, nor even surjective[10]. In addition, while reductive quotients of projective and affine varieties are respectively projective and affine, this may no longer be true for non-reductive quotients (see [23, §3.1] for further discussion on the differences between reductive and non-reductive group actions).

Nevertheless, Non-Reductive GIT identifies a certain class of linear algebraic groups for which an effective analogue of classical GIT can be obtained, that is, one which avoids the potential drawbacks of non-reductive group actions mentioned above. The building blocks of this class of linear algebraic groups are so-called externally graded unipotent groups.

Initial data and set-up. 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 the adjoint action of $\mathbb{G}_m$ on $\operatorname{Lie} U$ has positive weights. We call such groups *externally graded unipotent groups*. Suppose that $\widehat{U}$ acts linearly on a projective variety X with ample line bundle L. By taking a higher power if necessary we can assume it to be very ample, so that $X \subseteq \mathbb{P}(V)$ where $V = H^0(X, \mathcal{L})^\vee$. Let $\omega_{\min} = \omega_0 < \omega_1 < \cdots < \omega_{\max}$ denote the weights with which $\lambda(\mathbb{G}_m)$ acts on V.

The linearisation of the action of $\widehat{U}$ on X is *adapted* if $\omega_{\min} < 0 < \omega_1$. If it is not, then by taking a positive tensor power and twisting the linearisation by an appropriate character, we can ensure that the resulting linearisation is adapted. That is, let χ be a rational character of

[10] A simple example of the possible failure of surjectivity, in the affine case, is the action of the subgroup of $\mathrm{SL}(2,k)$ consisting of upper triangular unipotent matrices (which is a unipotent and therefore non-reductive group) on $\mathrm{SL}(2,k)$ via left multiplication: the invariants can be shown to be finitely generated, so that the associated affine variety is $\mathbb{C}^2$, yet the induced quotient map has image only $\mathbb{C}^2 \setminus \{(0,0\}$ (see [59, Ex 3.1.2.9] for more detail).

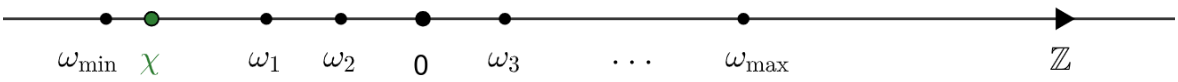

(a) Weights for the action of $\lambda(\mathbb{G}_m)$ on X with respect to a non-adapted linearisation. For c a sufficiently divisible positive integer, the character $c\chi$ lifts to a character of $\widehat{U}$ with U in its kernel. The linearisation can be turned into an adapted linearisation by taking its c-th tensor power and then twisting by $c\chi$, provided χ is chosen so that it is close enough to $\omega_{\min}$.

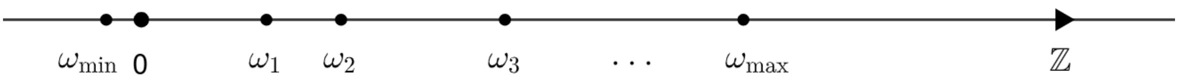

(b) Weights for the action of $\lambda(\mathbb{G}_m)$ on X with respect to an adapted linearisation.

Figure 1.4: Illustration of the distribution of weights for the action of $\lambda(\mathbb{G}_m)$ on X with respect to a non-adapted and adapted linearisation.

$\mathbb{G}_m$ such that

$$\omega_{\min} < \chi < \omega_1. \tag{1.4}$$

For c a sufficiently divisible positive integer, the character $c\chi$ lifts to a character of $\widehat{U}$ with U in its kernel. By twisting the linearisation of the action of $\widehat{U}$ on X with respect to the line bundle $L^{\otimes c}$ by the character $c\chi$, a weight ω_i is replaced by the weight $c(\omega_i - \chi)$. Thus if χ is sufficiently close to $\omega_{\min}$, then

$$c(\omega_{\min} - \chi) < 0 < \omega_1 - \chi \tag{1.5}$$

and so the resulting linearisation is adapted. Figure 1.4 illustrates the positions of the weights for adapted and non-adapted linearisations.

We let $X^{s,\mathbb{G}_m}_{\min+} \subseteq X$ denote the stable locus for the action of $\lambda(\mathbb{G}_m)$ on X with respect to this adapted linearisation. Note that by the theory of variation of GIT [20], this locus is independent of the choice of rational character χ and of the subsequent choice of tensor power c, provided (1.5) is satisfied. This justifies the use of the subscript 'min+' in the notation for this stable locus: it indicates that the linearisation has been twisted by a suitable character and raised to a suitable tensor power so that the origin lies just to the right of the minimum weight. This stable locus can be described explicitly in the following way. Let $V_{\min}$ denote the minimal weight space for the action of $\lambda(\mathbb{G}_m)$ on V. Moreover, set

$$Z_{\min} = X \cap \mathbb{P}(V_{\min}) \text{ and } X^0_{\min} = \{x \in X \mid \lim_{t\to 0} \lambda(t) \cdot x \in Z_{\min}\}.$$

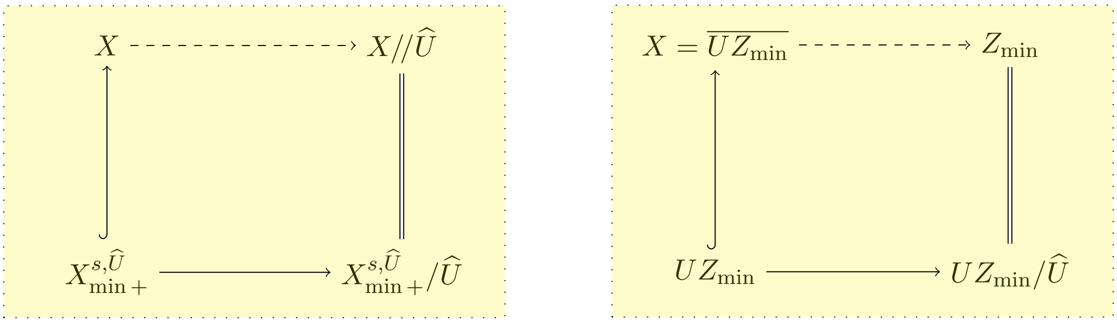

(a) General case.

(b) When $X = \overline{UZ_{\min}}$.

Figure 1.5: GIT for externally graded unipotent groups $\widehat{U} = U \rtimes \mathbb{G}_m$, when ($ss = s \neq \emptyset[\widehat{U}]$) holds.

Then by the Hilbert-Mumford criterion for the linear action of $\mathbb{G}_m$ on X (with respect to the twisted linearisation), we have that

$$X^{s,\mathbb{G}_m}_{\min +} = X^0_{\min} \setminus Z_{\min}.$$

Notation 1.1.4 (Notation for $Z_{\min}$)**.** If $\widehat{U}$ acts on another projective variety Y (i.e. a variety not denoted by X), then we write $Z(Y)_{\min}$ to denote the analogue of $Z_{\min}$ for Y.

We can now state the $\widehat{U}$-theorem, proved in [8], which is the core of Non-Reductive GIT.

The $\widehat{U}$-theorem. The content of the $\widehat{U}$-theorem is that if the linear action of $\widehat{U}$ on X satisfies an additional condition (analogous to the condition that semistability coincides with stability in classical GIT), then after taking a positive tensor power and twisting the linearisation by a suitable rational character, all of the properties of classical GIT in the case when semistability coincides with stability can be recovered. The precise formulation is as follows (see [8, Thm 2.16]).

Theorem 1.1.5 ($\widehat{U}$-theorem)**.** Let $\widehat{U} = U \rtimes \mathbb{G}_m$ where $\mathbb{G}_m$ acts with positive weights on $\operatorname{Lie} U$ via the adjoint action and suppose that $\widehat{U}$ acts linearly on an irreducible projective variety X with respect to a very ample line bundle L. If the condition

$$\operatorname{Stab}_U(z) = \{e\} \text{ for all } z \in Z_{\min} \qquad (ss = s \neq \emptyset[\widehat{U}])$$

is satisfied[11], then after possibly taking a suitable positive tensor power and twisting the linearisation by a suitable rational character so that the resulting linearisation is adapted, we

[11] This condition is analogous to the condition in classical GIT that $X^{ss} = X^s \neq \emptyset$, which is why it is denoted ($ss = s \neq \emptyset[\widehat{U}]$).

have:

(i) the open subvariety

$$X^{s,\widehat{U}}_{\min+} := \bigcap_{u \in U} u X^{s,\mathbb{G}_m}_{\min+} = X^0_{\min} \setminus UZ_{\min} \subseteq X \tag{1.6}$$

has a geometric $\widehat{U}$-quotient, denoted $X/\!/\widehat{U}$, which is a projective variety, so that set-theoretically $X/\!/\widehat{U} = X^{s,\widehat{U}}_{\min+}/\widehat{U}$;

(ii) there exists an $\epsilon > 0$ such that for any rational character χ of $\mathbb{G}_m$ satisfying $\omega_{\min} < \chi < \omega_{\min} + \epsilon < \omega_1$, there exists a sufficiently divisible integer c such that $c\chi$ lifts to a character of $\widehat{U}$ with kernel U and such that after taking the c-th tensor power of the linearisation and twisting it by the character $c\chi$, the invariants $\bigoplus_{k\geq 0} H^0(X, L^{\otimes ck})^{\widehat{U}}$ form a finitely generated algebra and its associated projective variety is isomorphic to $X/\!/\widehat{U}$.

It follows from Theorem 1.1.5 that the projective variety $X/\!/\widehat{U}$ satisfies all of the key properties of a classical GIT quotient in the case where semistability coincides with stability. That is, there exists a projective geometric quotient for the action of $\widehat{U}$ on an open subset $X^{s,\widehat{U}}_{\min+}$ of X, which can be described by a Hilbert-Mumford type criterion (compare (1.6) and (1.1)). Figure 1.5a illustrates this situation; note the similarity with Figure 1.1b. 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 [8, Rk 2.17]). Figure 1.5b illustrates this situation.

Remark 1.1.6 (Irreducibility assumption)**.** In the reductive case, the irreducibility assumption can be made without loss of generality: given a variety X, not necessarily irreducible, classical GIT produces the same result whether applied to each irreducible component separately (the quotients can be glued together to produce a quotient for X), or whether applied to X directly. The reason is that invariant sections defined on closed subvarieties of X always extend to X, thanks to the reductivity of the group, and this ensures that the semistable locus for the restricted linear action of X on any given subvariety coincides with the restriction to the subvariety of the semistable locus for X.

The irreducibility assumption is more subtle for non-reductive groups. Indeed, invariants for non-reductive groups may not always extend from closed subvarieties to the ambient variety.

Nevertheless, under the conditions of Theorem 1.1.5, if Y is a closed subvariety of X then the semistable locus for the restricted linear action of $\widehat{U}$ on Y does indeed coincide with the restriction to Y of the semistable locus for the linear action of $\widehat{U}$ on X, under one condition: the intersection $Y \cap X^0_{\min}$ must be non-empty, or in other words, the minimal weight space for Y must coincide with the minimal weight space for X. Thus a quotient for a reducible variety X can be obtained by patching together quotients for those irreducible components with minimal weight space coinciding with the restriction of the minimal weight space for X. We note that the need to remember further information about the irreducible components can also arise in classical GIT, since certain irreducible components of a variety X with a linear action by a reductive group may have an empty semistable locus for the restricted action of G and therefore would not contribute to the global quotient of X.

Remark 1.1.7 (Strengthening of the $\widehat{U}$-theorem)**.** The condition ($ss = s \neq \emptyset[\widehat{U}]$) of Theorem 1.1.5 above can in fact be weakened (see [8, Rk 2.8]). That is, if $U = U^{(0)} \supseteq U^{(1)} \supseteq \cdots \supseteq U^{(s)} = \{e\}$ is a normal series such that each successive quotient $U^{(j)}/U^{(j+1)}$ is abelian (for example the derived series), then provided the condition that

$$\text{for every } j, \ \dim \mathrm{Stab}_{U^{(j)}}(z) \text{ is constant for all } z \in Z_{\min} \qquad (ss = Ms \neq \emptyset[\widehat{U}])$$

is satisfied[12], the conclusions of Theorem 1.1.5 are still valid (in this case we denote the stable locus by $X^{Ms,\widehat{U}}_{\min+}$ instead). The weakest assumption under which this theorem holds are given in [8, Rk 7.12]. For simplicity, we will also denote this condition by ($ss = Ms \neq \emptyset[\widehat{U}]$). The resulting stronger version of the $\widehat{U}$-theorem is necessary to ensure that Theorem 1.2.3 (the $\widehat{U}$-theorem with blow-ups) holds without any assumptions on the linear action of $\widehat{U}$ on X.

We now turn to the action of linear algebraic groups which are not necessarily externally graded unipotent groups, but which contain such groups. More precisely, we consider groups $H = U \rtimes R$ where U denotes the unipotent radical of H and R a Levi subgroup satisfying that property that there exists a central one-parameter subgroup $\lambda : \mathbb{G}_m \to Z(R)$ acting on $\mathrm{Lie}\, U$ with positive weights. Such groups are said to have *internally graded unipotent radical.*

[12]This situation is analogous to the reductive case when $X^s = \emptyset$ but $\emptyset \neq X^{Ms} = X^{ss}$, which is why we denote the condition ($ss = Ms \neq \emptyset[\widehat{U}]$).

1.1.3 GIT for linear algebraic groups with internally graded unipotent radical

The above two Sections 1.1.1 and 1.1.2 provide all the tools required for constructing quotients of projective varieties by the actions of linear algebraic groups with internally graded unipotent radicals: given the linear action of such a group H on an irreducible projective variety X, the strategy is to quotient in stages, first by the action of $\widehat{U} := U \rtimes \lambda(\mathbb{G}_m)$ (to which results from Section 1.1.2 apply), then by the action of $R/\lambda(\mathbb{G}_m)$ (to which results from Section 1.1.1 apply).

In this section we describe this procedure in more detail, and study its application to the construction of quotients of GIT-unstable strata.

Quotienting in stages. Let $H = U \rtimes R$ be a linear algebraic group with graded unipotent radical acting linearly on X with respect to a very ample line bundle L and let $\lambda : \mathbb{G}_m \to R$ denote the central one-parameter subgroup such that $\lambda(\mathbb{G}_m)$ acts on $\operatorname{Lie} U$ with strictly positive weights. For ease of notation, let R_λ denote the reductive quotient group $R/\lambda(\mathbb{G}_m)$.

When the condition ($ss = s \neq \emptyset[\widehat{U}]$) (or ($ss = Ms \neq \emptyset[\widehat{U}]$)) is satisfied for the action of $\widehat{U} = U \rtimes \lambda(\mathbb{G}_m)$ on X, then by the results of Section 1.1.2, 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). Moreover, if the linearisation is tensored and twisted in a suitable way (that is, raised to a positive c-th tensor power and twisted by the character $c\chi$ of $\widehat{U}$, where χ is an adapted rational character of $\mathbb{G}_m$ that is sufficiently close to the minimal weight $\omega_{\min}$ and c is chosen so that $c\chi$ lifts to a character of $\widehat{U}$ with U in its kernel), then the $\widehat{U}$-invariants are finitely generated and the associated projective variety coincides with the projective quotient $X/\!/\widehat{U}$. This is the first stage of the construction.

For the second stage, we use the fact that the projective variety $X/\!/\widehat{U}$ has an induced action of the reductive group R_λ. Moreover, this action can be linearised in such a way that the pull-back of this linearisation to X under the quotient map coincides with a tensor power of the linearisation for the $\widehat{U}$-action on X (after tensoring and twisting the original linearisation according to the $\widehat{U}$-theorem). We can thus apply classical GIT to further quotient by R_λ and obtain a projective variety

$$X//H := (X//\widehat{U})//R_\lambda$$

which is a good quotient for the action of R_λ on $(X//\widehat{U})^{ss,R_\lambda}$. We also obtain a quasi-projective variety $(X//\widehat{U})^{s,R_\lambda}/R_\lambda$ which is a geometric quotient for the action of R_λ on $(X//\widehat{U})^{s,R_\lambda}$. Note that since the $\widehat{U}$-invariants (with respect to the tensored and twisted original linearisation) and the R_λ-invariants (with respect to the induced linearisation) are finitely generated, we have that the H-invariants are finitely generated for the linearisation of the H-action on X (obtained after taking a suitable c-th tensor power and twisting by an appropriate character). As a result, we obtain that

$$X//H = \operatorname{Proj} \bigoplus_{k \geq 0} (H^0(X, L^{\otimes ck})^{\widehat{U}})^{R_\lambda} = \operatorname{Proj} \bigoplus_{k \geq 0} H^0(X, L^{\otimes ck})^H.$$

By defining

$$X^{(s)s,H}_{\min +} := q_{\widehat{U}}^{-1}((X//\widehat{U})^{(s)s,R_\lambda}) \cap X^{s,\widehat{U}}_{\min +}, \tag{1.7}$$

where $q_{\widehat{U}} : X^0_{\min} \setminus UZ_{\min} \to X//\widehat{U}$ denotes the quotient map, we have that $X//H$ is a good quotient for the action of H on $X^{ss,H}_{\min +}$ and that $(X//\widehat{U})^{s,R_\lambda}/R_\lambda = X^{s,\widehat{U}}_{\min +}/H$ is a geometric quotient for the action of H on $X^{s,H}_{\min +}$.

Although we know from the $\widehat{U}$-theorem that the quotient $X//\widehat{U}$ is independent of the choice of rational character used to twist the initial linearisation, provided it is close enough to the minimal weight $\omega_{\min}$, it does not follow from the method of quotienting that this is also the case for the quotient $X//H$. To show that $X//H$ is similarly canonically determined by the initial linearisation, a Hilbert-Mumford type description of the (semi)stable locus $X^{ss,H}_{\min +}$ is required. This is achieved in [8] by considering the action of a semi-direct product $\widehat{U}_T := U \rtimes T$ where $T \subseteq R$ is a maximal torus containing $\lambda(\mathbb{G}_m)$. The result is the following (see [8, Thm 2.16]):

Theorem 1.1.8 (Hilbert-Mumford criterion for Non-Reductive GIT)**.** Let T be a torus and suppose that $\widehat{U}_T = U \rtimes T$, where $\widehat{U}_T$ has an internally graded unipotent radical, acts linearly on an irreducible projective variety X. Suppose moreover that the condition ($ss = s \neq \emptyset[\widehat{U}]$) is satisfied for this action. Then

$$X^{(s)s,\widehat{U}_T}_{\min +} = \bigcap_{u \in U} u X^{(s)s,T}_{\min +}.$$

More generally, if $H = U \rtimes R$ with internally graded unipotent radical acts linearly on an irreducible projective variety X and the condition ($ss = s \neq \emptyset[\widehat{U}]$) is satisfied, then

$$X^{(s)s,H}_{\min+} = \bigcap_{h \in H} hX^{(s)s,T}_{\min+}$$

for a fixed choice of maximal torus $T \subseteq R$.

By the theory of variation of GIT, there exits an $\epsilon > 0$ such that if χ is a rational character of T with restriction to $\lambda(\mathbb{G}_m)$ satisfying $\omega_{\min} < \chi|_{\lambda(\mathbb{G}_m)} < \omega_{\min} + \epsilon$, then the resulting (semi)stable locus $X^{(s)s,T}_{\min+}$ is independent of the choice of χ. Thus we obtain that the quotient $X/\!/H$ constructed using the method of quotienting in stages is canonically determined by the initial linearisation of the H-action on X.

Figure 1.6 illustrates the above results: Figure 1.6a depicts the general case (it can be viewed as a combination of Figures 1.1a and 1.5a) while Figure 1.6b depicts the case where semistability coincides with stability for the action of R_λ on $X/\!/\widehat{U}$, or equivalently for the action of the torus T on X (it can be viewed as a combination of Figures 1.1b and 1.5a). In the latter case, the projective variety $X/\!/H$ is a geometric quotient for the action of H on $X^{s,H}_{\min+}$.

Remark 1.1.9 (Variation of Non-Reductive GIT)**.** By Theorem 1.1.8, the quotient $X/\!/H$ is independent of the choice of rational character used to amend the linearisation in order to obtain finitely generated invariants, provided its restriction to $\lambda(\mathbb{G}_m)$ is close enough to the minimal weight $\omega_{\min}$ (see Figure 1.4a). Nevertheless, the quotient $X/\!/H$ does not depend only on the initial linearisation. Indeed, it can also depend on the choice of grading one-parameter subgroup $\lambda : \mathbb{G}_m \to Z(R)$. This dependence has been studied in [10] which establishes a 'variation of Non-Reductive GIT', analogous to the classical variation of GIT [97, 20].

Remark 1.1.10 (Non-Reductive GIT for linear algebraic groups with an externally graded unipotent radical)**.** In this section, we have considered only algebraic groups H with an internally graded unipotent radical. Nevertheless, all of the results described can be extended to the case where a linear algebraic group $H = U \rtimes R$ (without necessarily having an internally graded unipotent radical) acts on an irreducible projective variety X in such a way that this action can be extended to the action of some $\hat{H} = H \rtimes \mathbb{G}_m$ such that $\mathbb{G}_m$ acts with positive weights on $\operatorname{Lie} U$. Indeed, in this case good and geometric quotients can be constructed for the

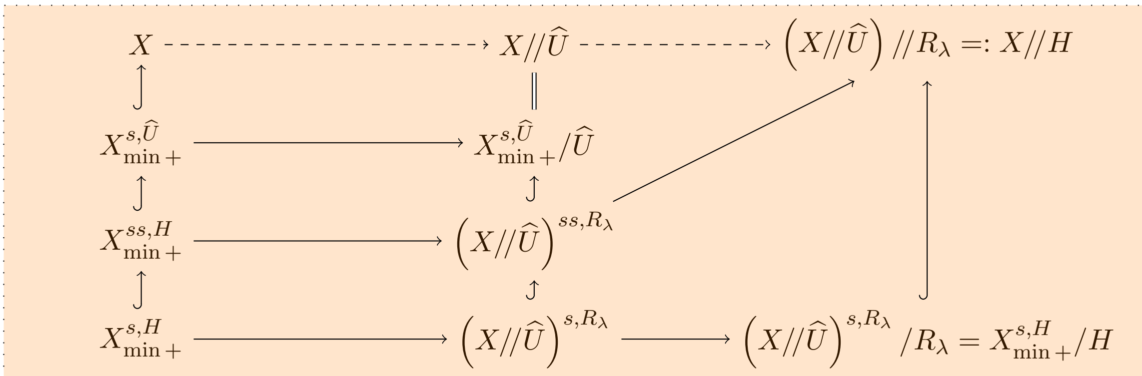

(a) General case.

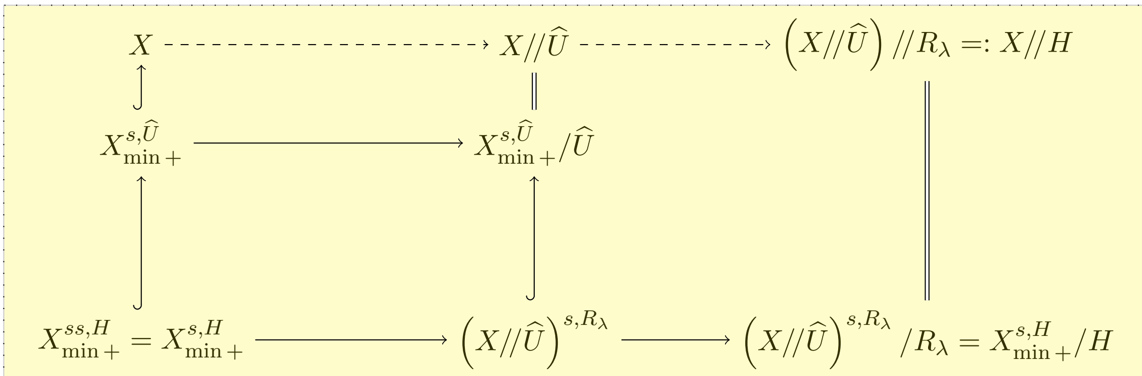

(b) When $\left(X/\!/\widehat{U}\right)^{ss,R_\lambda} = \left(X/\!/\widehat{U}\right)^{s,R_\lambda}$.

Figure 1.6: GIT for linear algebraic groups $H = U \rtimes R$ with internally graded unipotent radical, when ($ss = s \neq \emptyset[\widehat{U}]$) is satisfied. The grading one-parameter subgroup is denoted by $\lambda : \mathbb{G}_m \to Z(R)$, and R_λ denotes the quotient $R/\lambda(\mathbb{G}_m)$.

action of H on open subsets of X by considering instead the action of $\hat{H}$ on the product $X \times \mathbb{P}^1$ (the action of $\hat{H}$ on $\mathbb{P}^1$ is defined as multiplication by the character corresponding to the exact sequence $H \to \hat{H} \to \mathbb{G}_m$). The group $\hat{H}$ has an internally graded unipotent radical, and thus quotients for its action on $X \times \mathbb{P}^1$ can be constructed using the results from this section, and these can be interpreted in turn as quotients for the action of H on X (see [8, §9]).

Application of Non-Reductive GIT to GIT-unstable strata. As seen in Section 1.1.1, taking a quotient of an open subset of a GIT-unstable stratum S_β by the reductive group G is equivalent to taking a quotient of Y_β^{ss} by P_β, where $S_\beta \cong G \times_{P_\beta} Y_\beta^{ss}$. We now briefly describe how the action of P_β on Y_β^{ss} fits into the framework of Non-Reductive GIT described above.

To obtain a linear action of $P_\beta = U_\beta \rtimes L_\beta$ on a projective variety instead of on a quasi-projective variety, we consider the linear action of P_β on the closure $\overline{Y_\beta^{ss}}$ of Y_β^{ss} in X. This provides a setting in which Non-Reductive GIT can be applied, as described in [59, §4.2]. The varieties $Z(\overline{Y_\beta^{ss}})_{\min}$ and $(\overline{Y_\beta^{ss}})^0_{\min}$ are then given by Z_β and Y_β respectively (see [59, Prop 4.2.0.2]).

If the condition ($ss = s \neq \emptyset[\widehat{U}]$) is satisfied for the action of U_β on $\overline{Y_\beta^{ss}}$, then the $\widehat{U}$-theorem produces $X^{s,\widehat{U}_\beta}_{\min+} = Y_\beta \setminus U_\beta Z_\beta$ as the $\widehat{U}_\beta$-stable locus, where $\widehat{U}_\beta := U_\beta \rtimes \lambda_\beta(\mathbb{G}_m)$, which admits a projective geometric $\widehat{U}_\beta$-quotient

$$(Y_\beta \setminus U_\beta Z_\beta)/\widehat{U_\beta} = \overline{Y_\beta^{ss}}/\!/\widehat{U_\beta}.$$

Moreover, there are open subsets $\overline{Y_\beta^{ss}}^{(s)s,P_\beta}_{\min+}$ such that $\overline{Y_\beta^{ss}}^{ss,P_\beta}_{\min+}$ has a good quotient $\overline{Y_\beta^{ss}}/\!/P_\beta$, containing as an open subvariety the quasi-projective geometric quotient $\overline{Y_\beta^{ss}}^{s,P_\beta}_{\min+}/P_\beta$.

If we make the further assumption that $(\overline{Y_\beta^{ss}})^{ss,P_\beta}_{\min+} = (\overline{Y_\beta^{ss}})^{s,P_\beta}_{\min+} \neq \emptyset$ (this is equivalent to requiring that $\mathrm{Stab}_{P_\beta}(x)$ is finite for all $x \in (\overline{Y_\beta^{ss}})^{ss,P_\beta}_{\min+}$, see [12, Def 3.10]), then by [12, Lem 5.6] we have that $Y_\beta^{ss} \setminus UZ_\beta^{ss} = \overline{Y_\beta^{ss}}^{ss,P_\beta}_{\min+}$ and $(Y_\beta^{ss} \setminus UZ_\beta^{ss})/P_\beta = \overline{Y_\beta^{ss}}/\!/P_\beta$.

1.2 When semistability does not coincide with stability

In the previous section, we have seen that classical GIT has an effective analogue for linear algebraic groups with internally graded unipotent radical, provided an assumption analogous to the assumption in classical GIT that semistability coincides with stability is made (we denoted

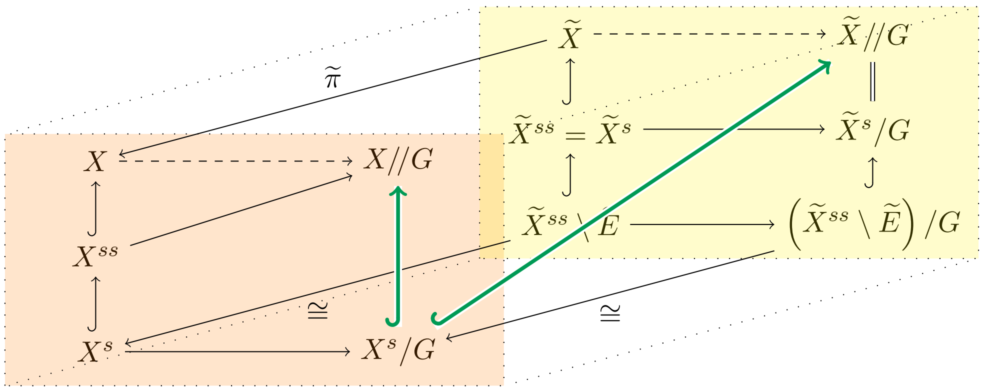

Figure 1.7: Classical GIT for reductive groups when $\emptyset \neq X^s \subsetneq X^{ss}$. The green arrows denote projective completions.

this condition ($ss = s \neq \emptyset[\widehat{U}]$)). If this condition is not satisfied, then it is not currently known in general whether the invariants are finitely generated and whether projective quotients can be obtained. Nevertheless, it is shown in [8] that in this case a sequence of equivariant blow-ups can be performed to obtain a variety for which the condition is satisfied, and thus for which the $\widehat{U}$-theorem can be applied. It is therefore possible to construct a quasi-projective geometric quotient for an open subset of the original variety in terms of this blowing-up procedure. This construction is analogous to the partial desingularisation construction of classical GIT which, given a variety for which semistability does not coincide with stability, provides a sequence of equivariant blow-ups resulting in a variety for which semistability does coincide with stability.

The aim of this section is to describe these results in more detail. In Section 1.2.1 we review the partial desingularisation construction of classical GIT. In Section 1.2.2 we present the results of [8, §7] which show that this construction can be extended to externally graded unipotent groups, so that the $\widehat{U}$-theorem can be extended to cases when the condition ($ss = s \neq \emptyset[\widehat{U}]$) (or ($ss = Ms \neq \emptyset[\widehat{U}]$)) is not satisfied. Finally in Section 1.2.3 we show how both constructions can be combined for linear algebraic groups with internally graded unipotent radical, to extend the results of Section 1.1.3 to the case where ($ss = s \neq \emptyset[\widehat{U}]$) or ($ss = Ms \neq \emptyset[\widehat{U}]$) is not satisfied.

1.2.1 GIT for reductive groups: the partial desingularisation construction

If we are in the situation of Figure 1.1b, that is when $X^s = X^{ss}$, then assuming that X is smooth the GIT quotient can have at worst finite quotient singularities. By contrast, if $X^s \subsetneq X^{ss}$, then the GIT quotient can have much worse singularities, due to the presence of strictly semistable points. The construction of [66] describes a way of obtaining a partial desingularisation of this GIT quotient, which can be viewed as an alternative and less singular projective completion of the geometric quotient X^s/G. This alternative projective completion is obtained as a GIT quotient of a variety for which semistability coincides with stability, as described below. For the construction to produce a variety with a non-empty semistable locus, the assumption that the stable or Mumford-stable locus of X is non-empty is required.

Suppose that $\emptyset \neq X^s \subsetneq X^{ss}$ (if X^s is empty but X^{Ms} is not, then stability should be replaced by Mumford-stability; note that the assumption that X is irreducible ensures that if X^s is non-empty then $X^s = X^{Ms}$). The partial desingularisation construction requires first blowing X up along the closure in X of the G-invariant subvariety of X^{ss} which consists of semistable points with reductive stabiliser group of maximal dimension (all such points are strictly semistable since by assumption $X^s \subsetneq X^{ss}$). The resulting blow-up $X_{(1)}$ has an induced action of G which can be linearised with respect to the ample line bundle obtained by pulling-back a tensor power of the ample line bundle on X and perturbing it by a sufficiently small multiple of the exceptional divisor. If the perturbation is sufficiently small, then the semistable locus $X_{(1)}^{ss}$ for this linear action coincides with the complement of the strict transform $\widetilde{C}$ of the saturation of C, namely $q_G^{-1}(q_G(C))$ where $q_G : X^{ss} \to X/\!/G$ denotes the GIT quotient map (see [66, Rk 7.17] or [88]). Importantly, the maximal dimension of stabiliser group for points in $X_{(1)}^{ss}$ is strictly smaller than that of points in X^{ss}.

The next step is to similarly blow $X_{(1)}$ up along the closure of the subvariety in $X_{(1)}^{ss}$ of points with reductive stabiliser group of maximal dimension, and again remove the strict transform of the saturation of the centre of the blow-up. After a finite number of steps, we obtain a variety $\widetilde{X}$ with a linear G-action and for which the dimension of stabiliser groups of points in $\widetilde{X}^{ss}$ is always 0. Thus $\widetilde{X}^{ss} = \widetilde{X}^s$ (see [66, §1.2]). Moreover, by construction of the blow-ups and of the linearisations, we have that $X^s = \widetilde{\pi}(\widetilde{X}^s \setminus \widetilde{E})$ where $\widetilde{E}$ denotes the exceptional divisor of the composition of blow-up maps $\widetilde{\pi} : \widetilde{X} \to X$. In this way the GIT quotient $\widetilde{X}/\!/G$ is birational to $X/\!/G$ and represents another projective completion of the geometric quotient X^s/G; the birational morphism $\widetilde{X}/\!/G \to X/\!/G$ is a partial desingularisation in the sense that it resolves the singularities of $X/\!/G$ coming from the presence of strictly semistable points. It also does not introduce new singularities, since the variety $\widetilde{X}^{ss}$ is smooth if X is smooth (we will review the proof of this result in Section 2.2.1 of Chapter 2). Figure 1.7 illustrates this blow-up construction. The two green arrows denote the two projective completions of X^s/G, namely $X/\!/G$ and $\widetilde{X}/\!/G$.

The partial desingularisation construction of GIT quotients can be used to describe the topology and geometry of GIT quotients. Indeed, starting from the G-equivariant Betti and Hodge numbers of a smooth projective variety X with a linear G-action such that $\emptyset \neq X^s \subsetneq X^{ss}$, the construction can be used to compute the Betti and Hodge numbers of the desingularised GIT quotient $\widetilde{X}/\!/G$ (we review this procedure in Section 2.3.1 of Chapter 2). This procedure is used for example in [67] as a step in the computation of the intersection Betti numbers of the moduli space of vector bundles of coprime rank and degree on a smooth projective curve.

Remark 1.2.1 (Canonical nature of the partial desingularisation construction)**.** The partial desingularisation construction is canonical in the sense that there is no choice involved for the centres of the blow-ups at each stage. There is however a choice involved for the linearisation at each stage. Nevertheless, at each stage there is a canonically defined chamber in the ample cone (for the linearisations of the action) such that provided the linearisation is chosen within that chamber, the resulting quotient is independent of the choices made at each stage.

Remark 1.2.2 (Generalisation of the partial desingularisation construction to stacks)**.** If the condition $X^{ss} = X^s \neq \emptyset$ is satisfied for the linear action of a reductive group G on a projective variety X then the associated quotient stack $[X^{ss}/G]$ is a Deligne-Mumford stack (in general a quotient stack is only an Artin stack), since the stabiliser groups must be finite (if $X^{ss} = X^{Ms} \neq \emptyset$, then $[X^{ss}/G]$ is a gerbe over a Deligne-Mumford stack instead). Thus in the language of stacks, the partial desingularisation can be viewed as a procedure for turning an Artin stack $[X^{ss}/G]$ into a Deligne-Mumford stack $[\widetilde{X}^{ss}/G]$ (provided the stable locus is non-empty, otherwise a gerbe over a Deligne-Mumford stack provided the Mumford-stable locus is non-empty) via a canonical sequence of blow-ups.

The partial desingularisation construction has in fact been generalised to stacks in [26]: given an Artin stack with a good moduli space (see Remark 1.1.1), a canonical sequence of blow-ups is described for obtaining a Deligne-Mumford stack (or a gerbe over a Deligne-Mumford stack) with a good moduli space which is also a coarse moduli space. The assumption required for the construction to result in a non-empty stack is that the starting good moduli space is 'stable' (see [26, Def 2.5 & Prop 2.6]), which can be viewed as the stack-theoretic analogue of the condition in classical GIT that the stable or Mumford-stable locus is non-empty. The construction given recovers the partial desingularisation construction when the good moduli space comes from a GIT quotient $[X^{ss}/G] \to X//G$ with non-empty stable or Mumford-stable locus.

1.2.2 GIT for externally graded unipotent groups

As reviewed above, in the case of classical GIT, if semistability does not coincide with stability (or Mumford-stability), then by the partial desingularisation construction a sequence of blow-ups can be performed to obtain a variety for which semistability does coincide with stability (or Mumford-stability). The aim of this section is to present the result of [8] which shows that an analogous result holds for the linear action of a positively graded $\widehat{U}$ on X: if the condition ($ss = s \neq \emptyset[\widehat{U}]$) or ($ss = Ms \neq \emptyset[\widehat{U}]$) is not satisfied, then a sequence of blow-ups can be performed to obtain a variety for which one of the two conditions is satisfied. After giving a precise statement of the result in Theorem 1.2.3 below, we describe the constructions required to prove this theorem (Blow-up Construction 1 and Blow-up Construction 2).

The $\widehat{U}$-theorem with blow-ups. The theorem is illustrated in Figure 1.8 and its statement is as follows (see [8, Thm 8.1]):

Theorem 1.2.3 ($\widehat{U}$-theorem with blow-ups)**.** Let $\widehat{U} = U \rtimes \mathbb{G}_m$, where U is a unipotent group and $\mathbb{G}_m$ acts on $\operatorname{Lie} U$ with positive weights, and suppose that $\widehat{U}$ acts linearly on an irreducible projective variety X. Then:

(i) if there is a $z \in Z_{\min}$ such that $\operatorname{Stab}_U(z) = \{e\}$, then there exists a canonical[13] sequence of $\widehat{U}$-equivariant blow-ups of X resulting in an irreducible projective variety $\widehat{\pi} : \widehat{X} \to X$ with a linear action of $\widehat{U}$ (determined by the linear action of $\widehat{U}$ on X) which satisfies the condition ($ss = s \neq \emptyset[\widehat{U}]$). Moreover, the projective geometric quotient $\widehat{X}//\widehat{U}$ is a

[13] We mean canonical in the same sense as for the partial desingularisation construction (see Remark 1.2.1). That is, the centres of the blow-ups at each stage are canonically defined, but there is a choice involved in the linearisation at each stage. Nevertheless, at each stage there is a canonical choice of chamber such that provided the linearisation is chosen within that chamber, the resulting quotient is independent of the choices made.

projective completion of the geometric $\widehat{U}$-quotient of the open subset of X given by

$$X^{\widehat{s},\widehat{U}} := \widehat{\pi}(\widehat{X}^{s,\widehat{U}}_{\min+} \setminus \widehat{E}) = \{x \in X^0_{\min} \setminus UZ_{\min} \mid \operatorname{Stab}_U(x) = \{e\}\},$$

where $\widehat{E}$ denotes the exceptional divisor for the sequence of blow-ups;

(ii) in general, for any choice of normal series $U = U^{(0)} \supseteq U^{(1)} \supseteq \cdots \supseteq U^{(s)} = \{e\}$ of U satisfying the following properties (denoted (†)):

- each subquotient $U^{(i)}/U^{(i+1)}$ is abelian;
- the action of $\mathbb{G}_m$ on $\operatorname{Lie} U^{(i)}/U^{(i+1)}$ has a single weight space;
- there exist natural numbers $r_1, \dots, r_s$ such that if $1 \leq j \leq s-1$ and $z \in X^{\mathbb{G}_m}$ where $\mathbb{G}_m$ acts on the fibre $L^\vee|_z$ with weight strictly less than ω_{r_j}, then $U^{(j)} = U^{(j+1)} \operatorname{Stab}_{U^{(j)}}(z)$ while for generic $z \in X^{\mathbb{G}_m}$ where $\mathbb{G}_m$ acts on $L^\vee|_z$ with weight equal to ω_{r_j}, then $\operatorname{Stab}_{U^{(j)}}(z) \subseteq U^{(j+1)}$,

there exists a sequence of blow-ups of $X \times \mathbb{P}^1 \times \cdots \times \mathbb{P}^1$ (where $\mathbb{P}^1$ is repeated s times) along projective subvarieties which are invariant under a suitably defined action of $\widehat{U} \times \mathbb{G}_m \times \cdots \times \mathbb{G}_m$ (where $\mathbb{G}_m$ is repeated s times) on $X \times \mathbb{P}^1 \times \cdots \times \mathbb{P}^1$, resulting in a projective variety $(X \times \mathbb{P}^1 \times \cdots \times \mathbb{P}^1)\widehat{}$ with a linear action of $\widehat{U} \times \mathbb{G}_m \times \cdots \times \mathbb{G}_m$ which satisfies either ($ss = s \neq \emptyset[\widehat{U}]$) or ($ss = Ms \neq \emptyset[\widehat{U}]$). Moreover, the projective geometric quotient $(X \times \mathbb{P}^1 \times \cdots \times \mathbb{P}^1)\widehat{}/\!/(\widehat{U} \times \mathbb{G}_m \times \cdots \times \mathbb{G}_m)$ is a projective completion of the geometric $\widehat{U}$-quotient of an open subset of X canonically determined by the sequence of blow-ups, and denoted $X^{\widehat{Ms},\widehat{U} \supseteq U^{(0)} \supseteq U^{(1)} \supseteq \cdots \supseteq U^{(s)}}$.

Remark 1.2.4 (Induced linearisation on $\widehat{X}$)**.** The linearisation of the $\widehat{U}$-action on the blown-up variety $\widehat{X}$ is determined by the linearisation of the $\widehat{U}$-action on X in the following way: it is the result of considering, at each stage i of the sequence of blow-ups, the linearisation obtained as a tensor product of the pull-back of the linearisation for the previous stage $i-1$ along the blow-up map with an arbitrarily small rational multiple ϵ_i of the exceptional divisor. The quotient $\widehat{X}/\!/\widehat{U}$ is independent of the finite sequence of choices of ϵ_i for $1 \gg \epsilon_1 \gg \epsilon_2 \gg \cdots$ (see [8, Prop 7.5]).

Remark 1.2.5 (Analogy with the partial desingularisation construction of classical GIT)**.** Although analogous to the partial desingularisation construction from classical GIT, Theorem 1.2.3 differs from it in two ways.

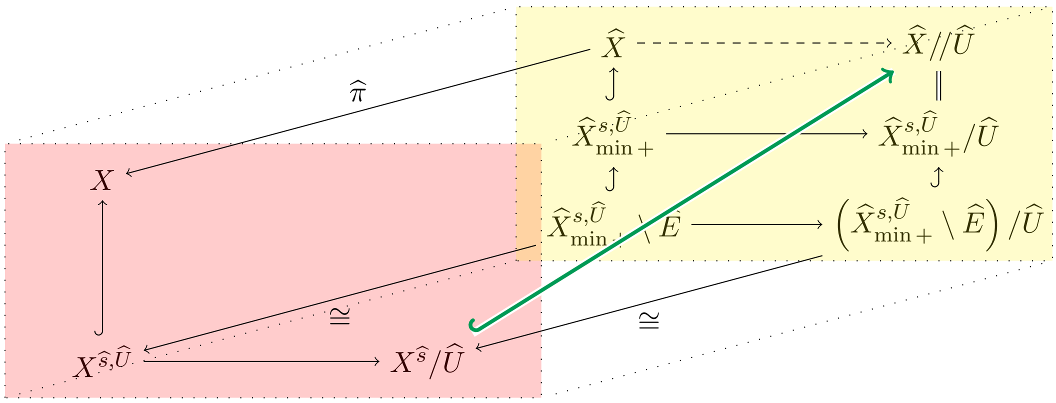

Figure 1.8: The $\widehat{U}$-theorem with blow-ups in the case where there exists a $z \in Z_{\min}$ such that $\mathrm{Stab}_U(z) = \{e\}$ (if there does not exist such a $z \in Z_{\min}$, then we must replace the superscripts 's' by 'Ms' for the action of $\widehat{U}$ on $\widehat{X}$, and the locus $X^{\widehat{s},\widehat{U}}$ by $X^{\widehat{Ms},\widehat{U} \supseteq U^{(0)} \supseteq U^{(1)} \supseteq \cdots \supseteq U^{(s)}}$ where $U = U^{(0)} \supseteq U^{(1)} \supseteq \cdots \supseteq U^{(s)} = \{e\}$ is the filtration used to perform the sequence of blow-ups). The green arrow denotes a projective completion.

Firstly, unlike the partial desingularisation construction which leads to a second projective completion of the geometric quotient X^s/G, birational to the GIT quotient $X/\!/G$, in the non-reductive case there is a priori no candidate projective quotient of X to which $\widehat{X}/\!/\widehat{U}$ (or the projective completion given in the part (ii) of Theorem 1.2.3 above) is birational. Indeed, without the condition of Theorem 1.1.5 (or the weaker conditions of Remark 1.1.7), it is not known at present whether a Non-Reductive GIT quotient can still be constructed for the action of $\widehat{U}$ on X (this would require first proving finite generation of the invariants for a suitable linearisation). Moreover, in the classical case, if X is smooth then $\widetilde{X}$ is also smooth because the centres of the blow-ups can be shown to be smooth. As a result, the quotient $\widetilde{X}/\!/G$ does indeed represent a partial desingularisation of $X/\!/G$, since the singularities arising from strictly semistable points have been resolved, and no new singularities have been introduced by the blow-up construction. Therefore if a projective quotient existed for the action of $\widehat{U}$ on X (assumed to be smooth), to be able to refer to $\widehat{X}/\!/\widehat{U}$ (or $(X \times \mathbb{P}^1 \times \cdots \times \mathbb{P}^1)\widehat{\ }/\!/(\widehat{U} \times \mathbb{G}_m \times \cdots \times \mathbb{G}_m)$) as a partial desingularisation of this quotient, we would have to know that the centres of the blow-ups are smooth. Proving this result is the core of Chapter 2 (see Theorem 2.2.4) where we show that the centres of the blow-ups in the construction corresponding to part (ii) of Theorem 1.2.3 are indeed smooth, provided X is smooth (with the caveat of having to deal with finite quotient singularities, see Remark 2.2.6). In this way, if a method is found for obtaining a projective quotient $X/\!/\widehat{U}$ even when the conditions of Theorem 1.1.5 are not satisfied, then

$(X \times \mathbb{P}^1 \times \cdots \times \mathbb{P}^1)\widehat{/\!/}(\widehat{U} \times \mathbb{G}_m \times \cdots \times \mathbb{G}_m)$ will indeed correspond to a partial desingularisation of $X/\!/\widehat{U}$.

Secondly, it is important to note that the partial desingularisation construction of classical GIT requires an assumption on the linear action of G on X, namely that the stable locus or the Mumford-stable locus is non-empty. The reason for this assumption is that if this is not the case, then the construction can result in a variety $\widetilde{X}$ for which the semistable locus is empty and thus the resulting GIT quotient is also empty. The assumption of the part (i) of Theorem 1.2.3 is the non-reductive analogue of the assumption in the classical case that the stable locus is non-empty. Under this assumption, a blow-up construction completely analogous to the partial desingularisation construction can be used to prove Theorem 1.2.3 (see Blow-up Construction 1 below). Nevertheless, this assumption is not required to obtain a non-empty projective completion, as indicated by the part (ii) of Theorem 1.2.3, in contrast to the partial desingularisation construction. That is, even if there is no $z \in Z_{\min}$ such that $\mathrm{Stab}_U(z) = \{e\}$, then a blow-up construction leading to a non-empty quotient can still be found.

Thus although Non-Reductive GIT does not at present provide a projective quotient for the action of $\widehat{U}$ on an open subset of X when the conditions of Theorem 1.1.5 are not satisfied, in the absence of these conditions, Theorem 1.2.3 can be used instead to construct a quasi-projective $\widehat{U}$-quotient of an open subset of X with a projective completion of this geometric quotient obtained as a Non-Reductive GIT quotient.

The blow-up constructions. We now describe the constructions used to prove Theorem 1.2.3. Blow-up Construction 1 describes the construction corresponding to part (i), while Blow-up Construction 2 describes that corresponding to part (ii).

We start by introducing the notation which we will use to denote the centres of the blow-ups (we will use this notation in Sections 2.2 and 2.3 of Chapter 2 as well).

Notation 1.2.6. Given the action of a group H on a variety Y, for each $d \in \mathbb{N}$ we define

$$C_d(Y,H) := \{y \in Y \,|\, \dim \mathrm{Stab}_H(y) = d\}$$

and

$$d_{\max}(Y,H) := \max\{\dim \mathrm{Stab}_H(y) \mid x \in Y\}.$$

Moreover, we let $C_{\max}(Y,H) := C_{d_{\max}(Y,H)}(Y,H)$ to simplify notation.

Blow-up Construction 1 (When there exists a $z \in Z_{\min}$ such that $\mathrm{Stab}_U(z) = \{e\}$)**.** We describe here the construction of $\widehat{\pi} : \widehat{X} \to X$ given in [8, Def 8.4]. As noted in the second point of Remark 1.2.5, it is the direct non-reductive analogue of the partial desingularisation construction of classical GIT.

The first step is to blow X up along the closure of $C_{\max}(X^0_{\min}, \widehat{U})$ in X. The key result is that the maximal dimension of unipotent stabiliser groups for points in the $X^0_{\min}$ for the blown-up space is strictly smaller than that for X. Thus by repeating this procedure finitely many times we obtain a variety $\widehat{X}$ for which ($ss = s \neq \emptyset[\widehat{U}]$) is satisfied. Note that this construction is completely analogous to the partial desingularisation construction in which X is blown-up along the locus of points with maximal dimension reductive stabiliser group.

The description of the open subset $X^{\widehat{s},\widehat{U}} := \widehat{\pi}(\widehat{X}^{s,\widehat{U}}_{\min+} \setminus \widehat{E})$ of X given in Theorem 1.2.3, as points in $X^0_{\min} \setminus UZ_{\min}$ with trivial unipotent stabiliser group, follows from [8, Prop 8.8]) (in the reductive case, this corresponds to the fact that $X^s = \widetilde{\pi}(\widetilde{X}^s \setminus \widetilde{E})$ where $\widetilde{\pi} : \widetilde{X} \to X$ is the result of the partial desingularisation construction and $\widetilde{E}$ the corresponding exceptional divisor).

Note that if the condition that there exists a $z \in Z_{\min}$ such that $\mathrm{Stab}_U(z) = \{e\}$ is not satisfied, then this sequence of blow-ups can still be performed but will result instead in a variety $\widehat{X}$ for which points in $X^0_{\min}$ have constant-dimensional unipotent stabiliser groups, rather than trivial unipotent stabiliser groups. However, having constant-dimensional unipotent stabiliser groups in U does not imply that ($ss = Ms \neq \emptyset[\widehat{U}]$) is necessarily satisfied (unless U is abelian and $\mathbb{G}_m$ acts on U with a single positive weight). As a result, Theorem 1.1.5 may not apply for the action of $\widehat{U}$ on $\widehat{X}$. For this reason, when the condition that there exists a $z \in Z_{\min}$ such that $\mathrm{Stab}_U(z) = \{e\}$ is not satisfied, we must instead use the more general construction, described in Blow-up Construction 2 below, which by contrast will always result in a variety satisfying either ($ss = s \neq \emptyset[\widehat{U}]$) or ($ss = Ms \neq \emptyset[\widehat{U}]$).

Remark 1.2.7 (Towards a stack-theoretic formulation of Blow-up Construction 1)**.** As seen in Remark 1.2.2, the partial desingularisation construction of classical GIT has been generalised to stacks in [26]. We expect that Blow-up Construction 1 could similarly be formulated in the language of stacks, by starting with a stack-theoretic generalisation of the map $[X^0_{\min}/U] \to Z_{\min}$

induced by the retraction map $X^0_{\min} \to Z_{\min}$, instead of starting with a good moduli space from an Artin stack. The condition that there is a point in $Z_{\min}$ with trivial unipotent stabiliser ensures that if we perform the saturated blow-ups from [26] (but blowing up along points with maximal unipotent stabiliser dimension, rather than maximal stabiliser dimension), the resulting stack is non-empty. Thus in the stack-theoretic setting the condition required for the application of Blow-up Construction 1 would represent the analogue of the condition required in [26] that the initial stack contains a stable point, a condition which ensures that the resulting stack is non-empty.

Remark 1.2.8 (Alternative Blow-up Construction 1)**.** We present here an alternative to Blow-up Construction 1, one which similarly results in a variety $\widehat{X}$ with an action of $\widehat{U}$ satisfying ($ss = s \neq \emptyset[\widehat{U}]$), under the same assumption that there exists a point in $Z_{\min}$ with trivial unipotent stabiliser. This alternative construction is one which was presented in an earlier version of [8], and consists in blowing $X^0_{\min}$ (and its equivalent at each stage of the construction) up along the locus of points with maximal dimension stabiliser in $\widehat{U}$, rather than in U. Thus the centre of the first blow-up is given by $C_{\max}(X^0_{\min}, \widehat{U})$ instead of $C_{\max}(X^0_{\min}, U)$. We note that there is an equality

$$C_{\max}(X^0_{\min}, \widehat{U}) = C_{\max}(X^0_{\min}, U) \cap UZ_{\min}. \tag{1.8}$$

Indeed, if $x = u \cdot z \in UZ_{\min}$ has maximal dimension stabiliser in U, then $u^{-1} \cdot x = z$, so that $u^{-1} \cdot x \in Z_{\min}$ and hence x is fixed by $u^{-1}\mathbb{G}_m u^{-1}$. Thus $\dim \mathrm{Stab}_{\widehat{U}}(x) = \dim \mathrm{Stab}_U(x) + 1$ and so x lies in $C_{\max}(X^0_{\min}, \widehat{U})$. For the inclusion, we note that if $x \in C_{\max}(X^0_{\min}, \widehat{U})$, then it must have the same unipotent stabiliser dimension as $p(x)$. This is because $\mathbb{G}_m$ normalises U, which implies that $\dim \mathrm{Stab}_U(p(x)) \geq \dim \mathrm{Stab}_U(x)$ (see [8, Rk 5.7]). Moreover, since $p(x) \in Z_{\min}$ is fixed by $\mathbb{G}_m$, then x must also be fixed by a one-parameter subgroup of $\widehat{U}$, which will be a conjugate of the grading $\mathbb{G}_m$ in $\widehat{U}$ and thus we obtain that $x \in UZ_{\min}$.

We will use this alternative Blow-up Construction 1 rather than Blow-up Construction 1 in Section 2.2 of Chapter 2 to show that if X is smooth then the centre of the non-reductive GIT blow-up is smooth (see Theorem 2.2.4). The advantage is that it then becomes sufficient to show that the intersection of $C_{\max}(X^0_{\min}, U)$ with $UZ_{\min}$ is smooth, rather than all of $C_{\max}(X^0_{\min}, U)$. Since this alternative Blow-up Construction 1 similarly results in a variety $\widehat{X}$ with a linear

$\widehat{U}$-action for which ($ss = s \neq \emptyset[\widehat{U}]$) is satisfied, there is no disadvantage in working with this construction instead.

Blow-up Construction 2 (General case)**.** If the condition that there exists a point $z \in Z_{\min}$ such that $\mathrm{Stab}_U(z) = \{e\}$ is not satisfied, a suitable sequence of blow-ups still exists. As indicated in Theorem 1.2.3, the blow-up construction relies on a choice of normal series

$$U = U^{(0)} \supseteq U^{(1)} \supseteq \cdots \supseteq U^{(s)} = \{e\}$$

for U satisfying the conditions of (†). Note that such a filtration can always be found given our assumptions on $\widehat{U}$.

Given a normal series of U satisfying (†), the idea of the construction is to perform Blow-up Construction 1 inductively for each $U^{(j)}/U^{(j+1)}$, starting with $U^{(s-1)}/U^{(s)}$, so that at each stage a quotient by a group of the form $U^{(j-1)}/U^{(j)} \rtimes \mathbb{G}_m$ can be constructed. Using the procedure of quotienting in stages introduced in Section 1.1.3 above, we can interpret the resulting variety as a quotient by $\widehat{U}$. However, for this inductive procedure to work, we must replace X by a product of X with s copies of $\mathbb{P}^1$, and $\widehat{U}$ by a product of $\widehat{U}$ with s copies of $\mathbb{G}_m$. This is to ensure that at each stage of the inductive process we are taking the quotient by a positively graded extension of a unipotent group (see [8, §5] for more detail).

Remark 1.2.9 (A conjectural 'wall and chamber' picture for Blow-up Construction 2)**.** We expect that, just as in the reductive case and in the case of Blow-up Construction 1, there should exist at each stage of the construction a canonical chamber such that provided the linearisation is chosen in that chamber, the resulting quotient in the blow-up is independent of all the choices made. This is more difficult to establish than in the reductive case or the case of Blow-up Construction 1, because of the need to twist the linearisation by a character each time a $\widehat{U}$-quotient is taken (as per Theorem 1.1.5), and thus remains conjectural at present.

1.2.3 GIT for linear algebraic groups with internally graded unipotent radicals

The partial desingularisation construction of reductive GIT can be combined with its non-reductive analogue given in Theorem 1.2.3 to produce a general result for constructing quasi-projective geometric quotients of projective varieties by linear actions of algebraic groups with

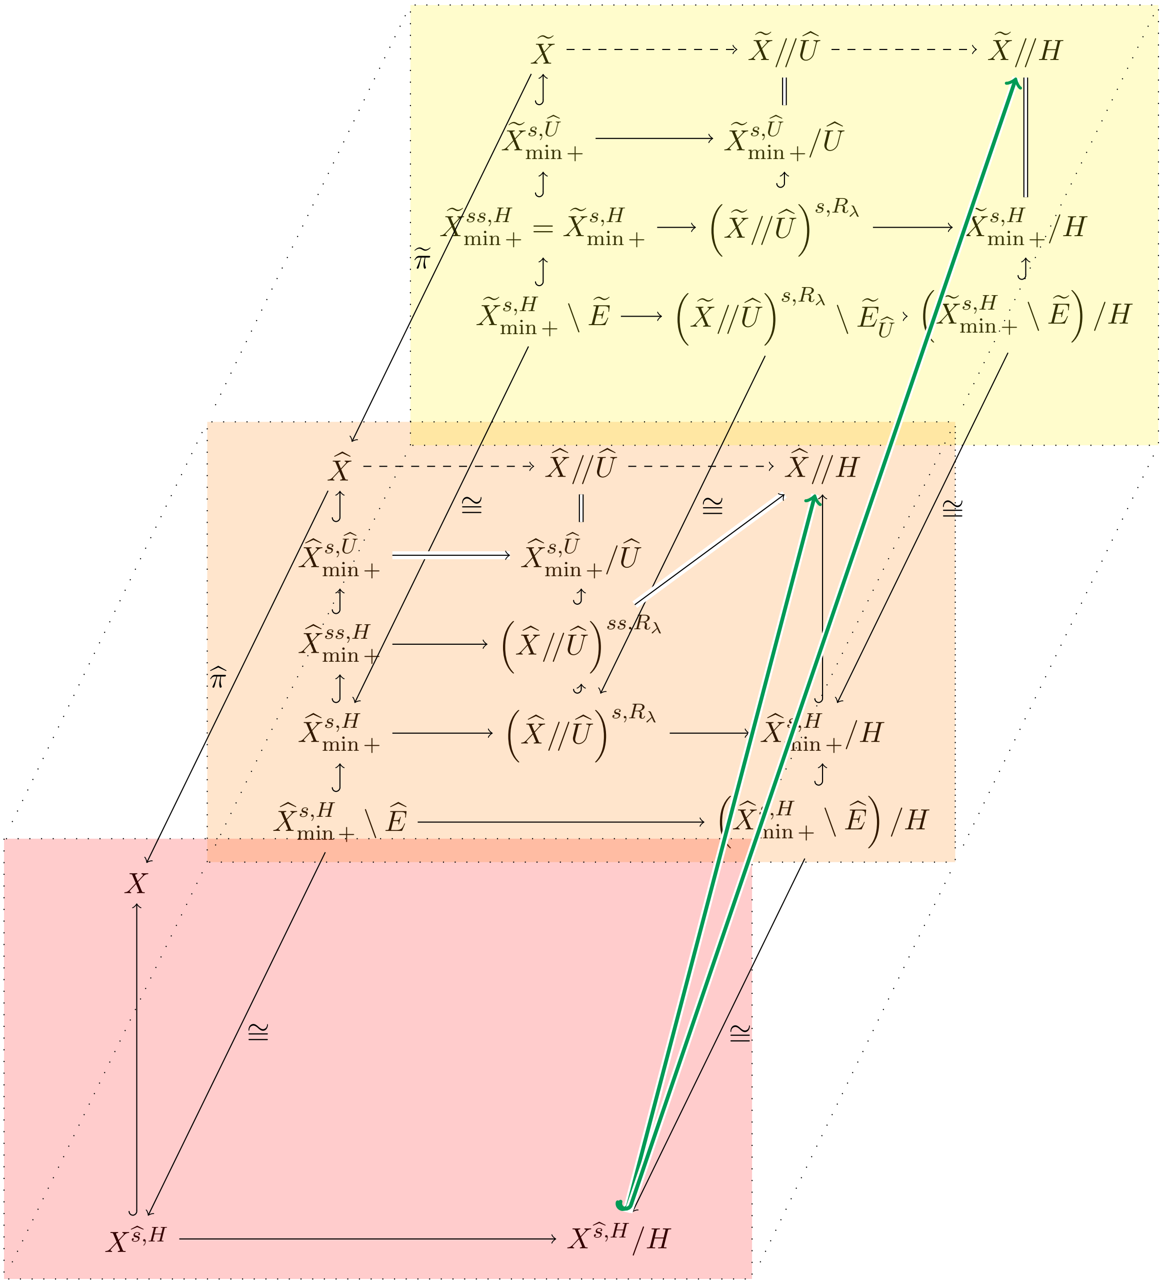

Figure 1.9: GIT for linear algebraic groups with internally graded unipotent radical. We assume that the linear action of $H = U \rtimes R$ on X satisfies the property that there exists a point in $Z_{\min}$ with trivial unipotent stabiliser group, so that Blow-up Construction 1 applies. The grading one-parameter subgroup is denoted by $\lambda : \mathbb{G}_m \to Z(R)$, and R_λ denotes the quotient $R/\lambda(\mathbb{G}_m)$.

internally graded radical, with explicit projective completions. That is, we can first perform blow-ups for the action of $\widehat{U}$ on X to obtain a projective variety $\widehat{X}$ for which a projective geometric quotient $\widehat{X}/\!/\widehat{U}$ exists (as per Theorem 1.1.5 and Remark 1.1.7), and furthermore we can ensure that these blow-ups are H-equivariant. Then, we can perform the partial desingularisation construction for the action of R_λ on $\widehat{X}/\!/\widehat{U}$, to obtain a projective variety $\widetilde{X}$ admitting a projective geometric quotient $\widetilde{X}/\!/H$. The statement of the theorem is as follows (see [8, Thm 2.20]) and is illustrated in Figure 1.9, where for simplicity we assume that Blow-up Construction 1 applies so that we need not replace X by a product of X with copies of $\mathbb{P}^1$ and H by a product of H with copies of $\mathbb{G}_m$.

Theorem 1.2.10 (Constructing quotients and projective completions for actions by linear algebraic groups with internally graded unipotent radical)**.** Let $H = U \rtimes R$ be a linear algebraic group with graded unipotent radical and let $\lambda : \mathbb{G}_m \to R$ denote the central one-parameter subgroup acting on $\operatorname{Lie} U$ with strictly positive weights. Suppose that H acts linearly on an irreducible projective variety X with respect to a very ample line bundle L. Then:

(i) for any choice of normal series $U = U^{(0)} \supseteq U^{(1)} \supseteq \cdots \supseteq U^{(s)} = \{e\}$ of U satisfying (†) (see Theorem 1.2.3), there exists a sequence of blow-ups of $X_{\mathbb{P}^1} := X \times \mathbb{P}^1 \times \cdots \times \mathbb{P}^1$ (where $\mathbb{P}^1$ is repeated s times) along projective subvarieties which are invariant under a suitably defined action of $H_{\mathbb{G}_m} := H \times \mathbb{G}_m \times \cdots \times \mathbb{G}_m$ (where $\mathbb{G}_m$ is repeated s times) on $X_{\mathbb{P}^1}$, resulting in a projective variety $\widehat{X}_{\mathbb{P}^1}$ with a linear action of $\widehat{U}_{\mathbb{G}_m} := \widehat{U} \times \mathbb{G}_m \times \cdots \times \mathbb{G}_m$ which satisfies either ($ss = s \neq \emptyset[\widehat{U}]$) or ($ss = Ms \neq \emptyset[\widehat{U}]$). The projective variety $\widehat{X}_{\mathbb{P}^1}/\!/\widehat{U}_{\mathbb{G}_m}$ is a projective geometric quotient for the linear action of $\widehat{U}_{\mathbb{G}_m}$ on $(\widehat{X}_{\mathbb{P}^1})^{ss,\widehat{U}}_{\min+}$, and $\widehat{X}_{\mathbb{P}^1}/\!/H_{\mathbb{G}_m}$ is a projective good quotient for the action of $H_{\mathbb{G}_m}$ on $(\widehat{X}_{\mathbb{P}^1})^{ss,H}_{\min+}$;

(ii) there is a further sequence of blow-ups along $H_{\mathbb{G}_m}$-invariant projective subvarieties resulting in a projective variety $\widetilde{X}_{\mathbb{P}^1}$ satisfying the same conditions as $\widehat{X}_{\mathbb{P}^1}$ for the action of $\widehat{U}_{\mathbb{G}_m}$ and such that $\widetilde{X}_{\mathbb{P}^1}/\!/H_{\mathbb{G}_m}$ is a projective geometric quotient for the action of $H_{\mathbb{G}_m}$ on $(\widetilde{X}_{\mathbb{P}^1})^{s,H}_{\min+}$;

(iii) the blow-up maps $\hat{\pi} : \widehat{X}_{\mathbb{P}^1} \to X_{\mathbb{P}^1}$ and $\widetilde{\pi} : \widetilde{X}_{\mathbb{P}^1} \to X_{\mathbb{P}^1}$ are isomorphisms over an $H_{\mathbb{G}_m}$-invariant open subvariety of $(X_{\mathbb{P}^1})^0_{\min} \setminus UZ(X_{\mathbb{P}^1})_{\min}$ and this open subvariety admits a quasi-projective geometric $H_{\mathbb{G}_m}$-quotient which can be identified via $\widehat{\pi}$ and $\widetilde{\pi}$ as an open subvariety of the projective varieties $\widehat{X}_{\mathbb{P}^1}/\!/H_{\mathbb{G}_m}$ and $\widetilde{X}_{\mathbb{P}^1}/\!/H_{\mathbb{G}_m}$ respectively.

Remark 1.2.11 (The resulting open subset of X admitting a geometric H-quotient)**.** Given a linear action of H on X, Theorem 1.2.10 (iii) identifies an open subset of $X_{\mathbb{P}^1}$ (the image of the complement of the exceptional divisor in $(\widetilde{X}_{\mathbb{P}^1})^{s,H}_{\min+}$ under the blow-down map $\widehat{\pi}$) admitting a geometric $H_{\mathbb{G}_m}$-quotient with $\widehat{X}_{\mathbb{P}^1}/\!/H_{\mathbb{G}_m}$ and $\widetilde{X}_{\mathbb{P}^1}/\!/H_{\mathbb{G}_m}$ as projective completions. By considering the embedding of X inside $X_{\mathbb{P}^1}$ given by pairing X with the point [1:1] for each $\mathbb{P}^1$ and by restricting this geometric quotient to X, we obtain a geometric quotient for the action of H on an open subset of X given by the intersection with X of the open subset of $X_{\mathbb{P}^1}$ admitting a geometric $H_{\mathbb{G}_m}$-quotient (see [8, Rk 2.12 & §9]).

Notation 1.2.12 (Notation for the open subset of X admitting a geometric H-quotient)**.** The open subset of X described in Remark 1.2.11 above depends on the initial linearisation of the action of H on X, on the choice of normal series of U and a priori (see Remark 1.2.9) on the choices of linearisation made at each stage of Blow-up Construction 2. We let $\mathcal{L}^*$ encode all of this data[14], and denote the resulting open subset of X by $X^{\widehat{Ms},H,\mathcal{L}^*}$. The superscript '$\widehat{Ms}$' indicates that it is an open subset obtained by applying Blow-up Construction 2 (this is the reason for the '$\widehat{\ }$' notation), which can be used in cases where Blow-up Construction 1 cannot, and in such cases results in a variety satisfying the condition ($ss = Ms \neq \emptyset[\widehat{U}]$) rather than ($ss = s \neq \emptyset[\widehat{U}]$) (this is the reason for the 'Ms' notation).

1.3 Algorithm for constructing geometric quotients and projective completions

As covered in Section 1.1.3, if a linear algebraic group H with internally graded unipotent radical acts linearly on an irreducible projective variety X, then by Theorem 1.2.10 there exists an open subset $X^{\widehat{Ms},H,\mathcal{L}^*} \subseteq X$ such that $X^{\widehat{Ms},H,\mathcal{L}^*}/H$ is a quasi-projective geometric quotient for the action of H on $X^{\widehat{Ms},H,\mathcal{L}^*}$. Moreover, this geometric quotient has a projective completion $\widetilde{X}_{\mathbb{P}_1}/\!/H_{\mathbb{G}_m}$ which is itself a projective geometric quotient for the action of $H_{\mathbb{G}_m}$ on an open subset of the variety $\widetilde{X}_{\mathbb{P}_1}$, obtained as a result of performing Blow-up Construction 2 to the

[14]In classical GIT the only data required to define stability is the choice of a linearisation, often denoted $\mathcal{L}$. This is why we have chosen the notation $\mathcal{L}^*$, to indicate that it contains more than just the data of the linearisation.

action of H on X. However, just as in classical GIT where the stable and semistable loci can be empty, leading to empty geometric and GIT quotients, the geometric quotient $X^{\widehat{Ms},H,\mathcal{L}^*}/H$ and its projective completion $\widetilde{X}_{\mathbb{P}_1}//H_{\mathbb{G}_m}$ could be empty. The purpose of this section is to describe a generalisation of Theorem 1.2.10 which ensures that given a linear action of H on X, a non-empty geometric quotient and projective completion is always obtained. The proof is in the form of an algorithm, called the Projective Completion algorithm; it is based on the inductive construction from [9].

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 Projective Completion algorithm produces a non-empty open subset $S_0(X, H, \lambda, \mathcal{L})$ of X which admits a quasi-projective geometric quotient by the action of H, together with a projective completion $\mathrm{PC}(S_0(X, H, \lambda, \mathcal{L})/H)$ of this geometric quotient.

This Projective Completion algorithm is based on [9]. It will be described in Section 1.3.2 and is illustrated in Figure 1.11. It relies on another stand-alone algorithm called the Replacement algorithm, which we describe in Section 1.3.1 and which is illustrated in Figure 1.10. The Replacement algorithm also appears on the left-hand side of Figure 1.11.

1.3.1 The Replacement algorithm

Given the linear action of a linear algebraic group H with internally graded unipotent radical on an irreducible projective variety X, if the grading one-parameter subgroup acts trivially on X (this implies by the positive grading of $\mathbb{G}_m$ on $\operatorname{Lie} U$ that U also acts trivially), and moreover if the semistable locus for the residual reductive action is empty, then the classical GIT quotient is empty. The aim of the Replacement algorithm is to describe a way of replacing such a variety by a different but related variety (with a linear action by a different but related linear algebraic group, that also has an internally graded unipotent radical) for which either the grading one-parameter subgroup acts non-trivially, or if it does act trivially then for which the semistable locus for the induced reductive action is non-empty. This procedure is important for obtaining non-empty quotients, which is the goal of the Projective Completion algorithm described in Section 1.3.2 below.

The Replacement algorithm is illustrated in Figure 1.10 and is described in words below.

Replacement algorithm. Suppose that $(X, H, \lambda, \mathcal{L})$ encodes the data consisting of the linear 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 inner 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 because of the positive grading of $\mathbb{G}_m$ on $\operatorname{Lie} U$, and so X is already a geometric quotient for the action of $\widehat{U}$ on X. Thus it suffices to consider the action of R_λ 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 1.10, and can be described in words as follows:

Step 1. 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 1.1.1, 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.

Step 2. If $(\overline{Y_\beta^{ss}}, P_\beta, \lambda_\beta, \mathcal{L}|_{\overline{Y_\beta^{ss}}})$ satisfies ($\star$), then we set $(X', H', \lambda', \mathcal{L}') = (\overline{Y_\beta^{ss}}, P_\beta, \lambda_\beta, \mathcal{L}|_{\overline{Y_\beta^{ss}}})$ and the algorithm terminates here.

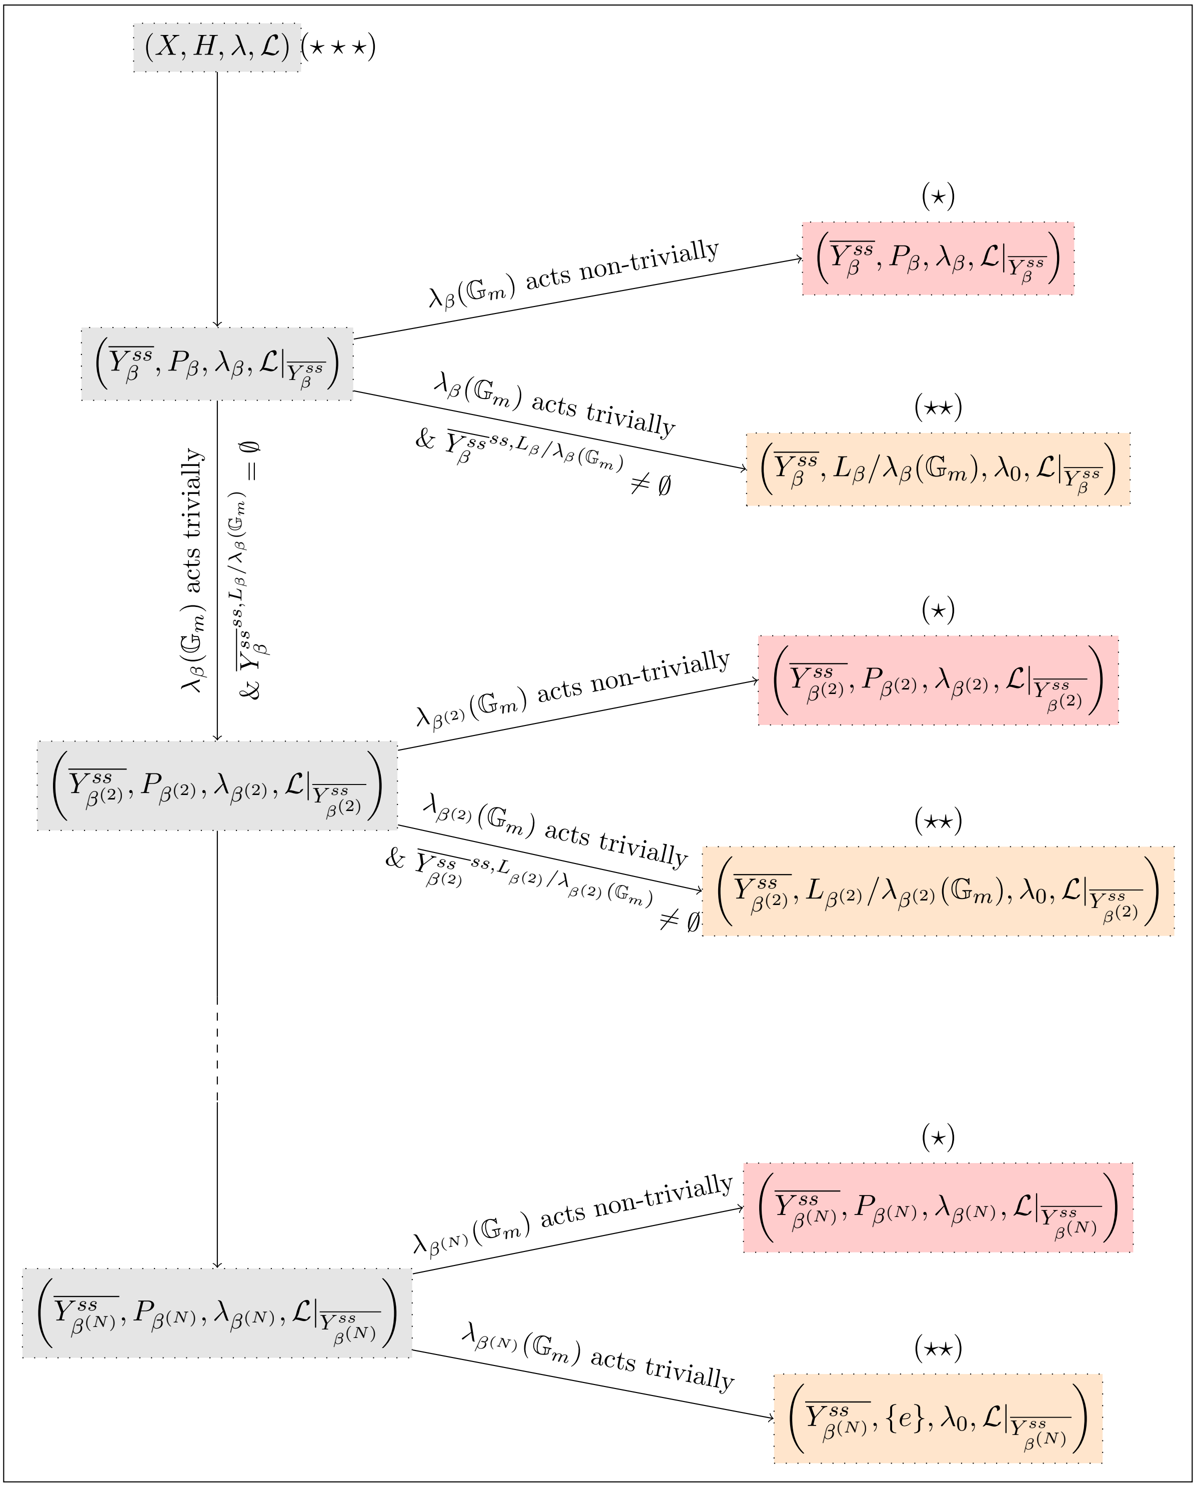

Figure 1.10: Replacement algorithm. The algorithm starts in the top left grey rectangle, and proceeds by following the arrows according to the conditions specified above and below the arrows. The algorithm terminates when a red or orange rectangle is reached.

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 a stratum $S_{\beta^{(2)}} := R_\lambda Y^{ss}_{\beta^{(2)}}$ with $\beta^{(2)} \neq 0$ which is open in $\overline{Y^{ss}_\beta}$ and therefore irreducible since $\overline{Y^{ss}_\beta}$ is irreducible. 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.

Step 3. Repeat Step 2, but this time 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 $\boldsymbol{k+1}$. Repeat Step k, starting with the 4-tuple $(\overline{Y^{ss}_{\beta^{(k-1)}}}, P_{\beta^{(k-1)}}, \lambda_{\beta^{(k-1)}}, \mathcal{L}|_{\overline{Y^{ss}_{\beta^{(k-1)}}}})$ instead of $(\overline{Y^{ss}_{\beta^{(k)}}}, P_{\beta^{(k)}}, \lambda_{\beta^{(k)}}, \mathcal{L}|_{\overline{Y^{ss}_{\beta^{(k)}}}})$. 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 $(\overline{Y^{ss}_{\beta^{(k+1)}}}, P_{\beta^{(k+1)}}, \lambda_{\beta^{(k+1)}}, \mathcal{L}|_{\overline{Y^{ss}_{\beta^{(k+1)}}}})$.

Proposition 1.3.1. 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. □

1.3.2 The Projective Completion algorithm

This algorithm is illustrated in Figure 1.11 and described in words below.

Projective Completion algorithm. Given a 4-tuple $(X, H, \lambda, \mathcal{L})$ encoding the data consisting of the linear action of a linear algebraic 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 $S_0(X, H, \lambda, \mathcal{L}) \subseteq X$ such that $S_0(X, H, \lambda, \mathcal{L}) \to S_0(X, H, \lambda, \mathcal{L})/H$ is a geometric for the action of H on $S_0(X, H, \lambda, \mathcal{L})$; and

ii) an irreducible projective variety X' with an action by a group $H' = U' \rtimes R'$ with internally graded unipotent radical satisfying ($ss = s \neq \emptyset[\widehat{U}]$) and such that $X'/\!/H' = X'^{(M)s,H}_{\min+}/H'$ is a projective geometric quotient for the action of H' on $S_0(X, H, \lambda, \mathcal{L})$ and a projective completion of $S_0(X, H, \lambda, \mathcal{L})/H$. This projective completion is denoted $\mathrm{PC}(S_0(X, H, \lambda, \mathcal{L})/H)$.

To simplify notation, we denote the open subset $S_0(X, H, \lambda, \mathcal{L})$ by $S_0(X, H)$, leaving implicit the dependence on the linearisation and on the choice of grading one-parameter subgroup.

Remark 1.3.2 (Familiar case)**.** In the familiar case where a reductive group G (in this case the grading one-parameter subgroup is the trivial one) acts linearly on an irreducible projective variety X with non-empty stable locus, the algorithm produces the GIT-stable locus X^s as the G-invariant open subset $S_0(X, G)$, and the GIT quotient $\widetilde{X}/\!/G$, resulting from the partial desingularisation construction, as the projective completion.

The description of the algorithm below relies on the conditions $(\star)$, $(\star\star)$ and $(\star\star\star)$ introduced in Section 1.3.1 above.

Case 1. Suppose that $(X, H, \lambda, \mathcal{L})$ satisfies $(\star)$. We start in the red square of the right-hand rectangle in Figure 1.11. According to the diagram two paths can be taken: 1 or 2. Both

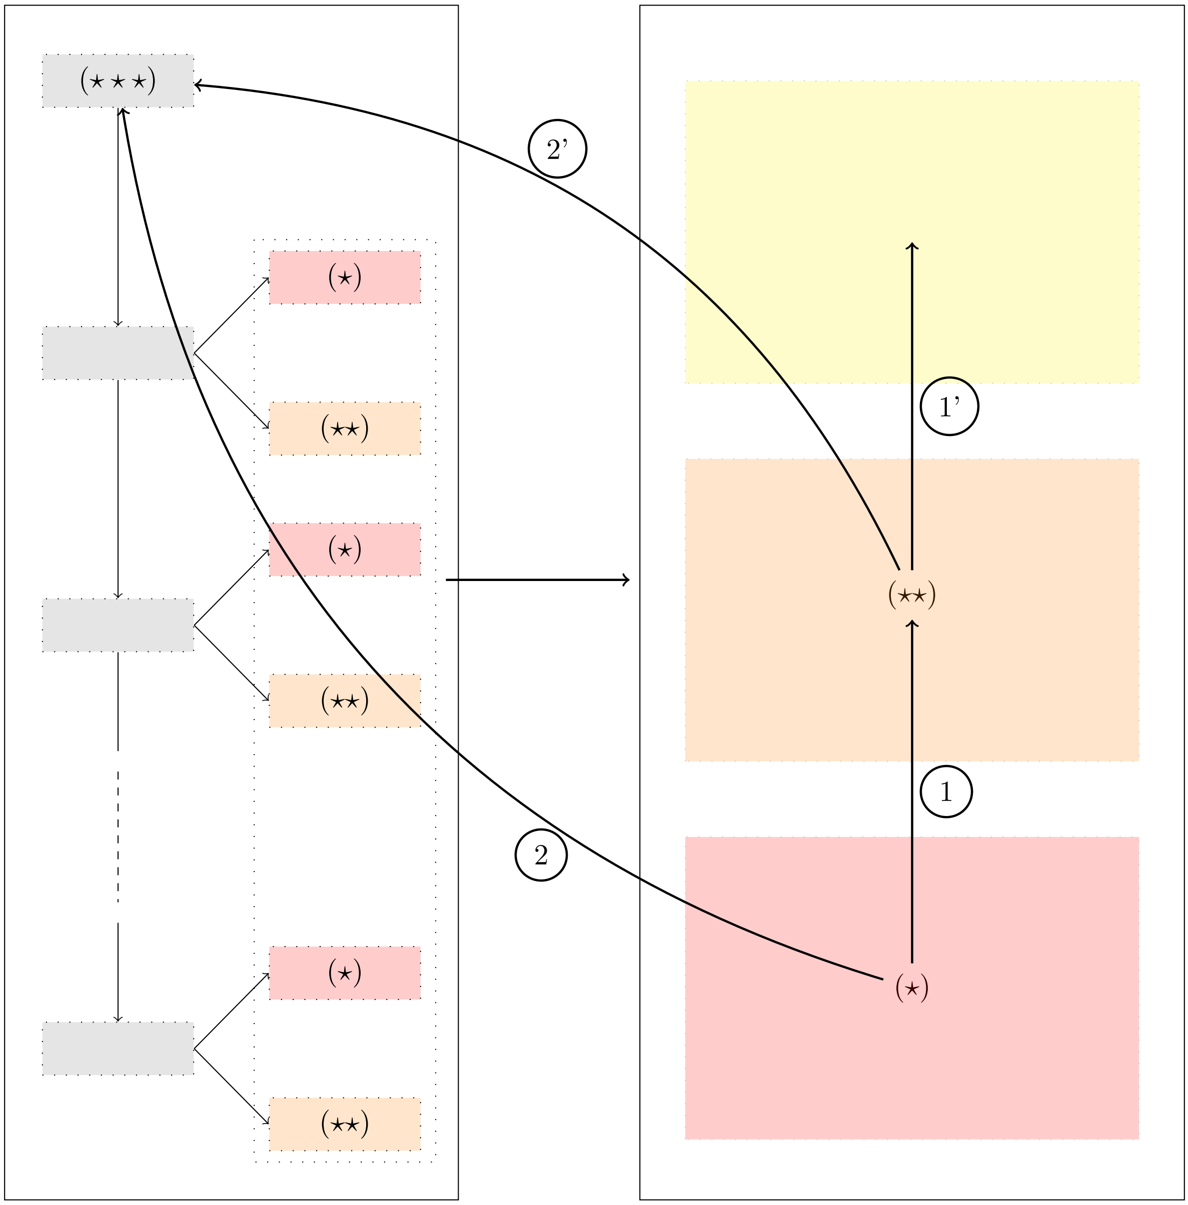

Figure 1.11: Projective Completion algorithm for the action of a linear algebraic group with internally graded unipotent radical on an irreducible projective variety. The left-hand side corresponds to the Replacement algorithm, described in Figure 1.10. The sequence of blow-ups corresponding to each of the four paths 1, 2, $1'$ and $2'$ is described in Figure 1.12. Given a 4-tuple $(X, H, \lambda, \mathcal{L})$, the algorithm starts in the grey rectangle in the top left or in the red or orange boxes 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. If at any point in the algorithm we find ourselves on the grey rectangle 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 box on the right-hand side: if the output is red we move to the red box, if it is orange we move to the orange box. Figure 1.13 depicts the different paths which can be taken within the algorithm and illustrates why the algorithm must terminate in the yellow box after a finite number of steps (Proposition 1.3.3).

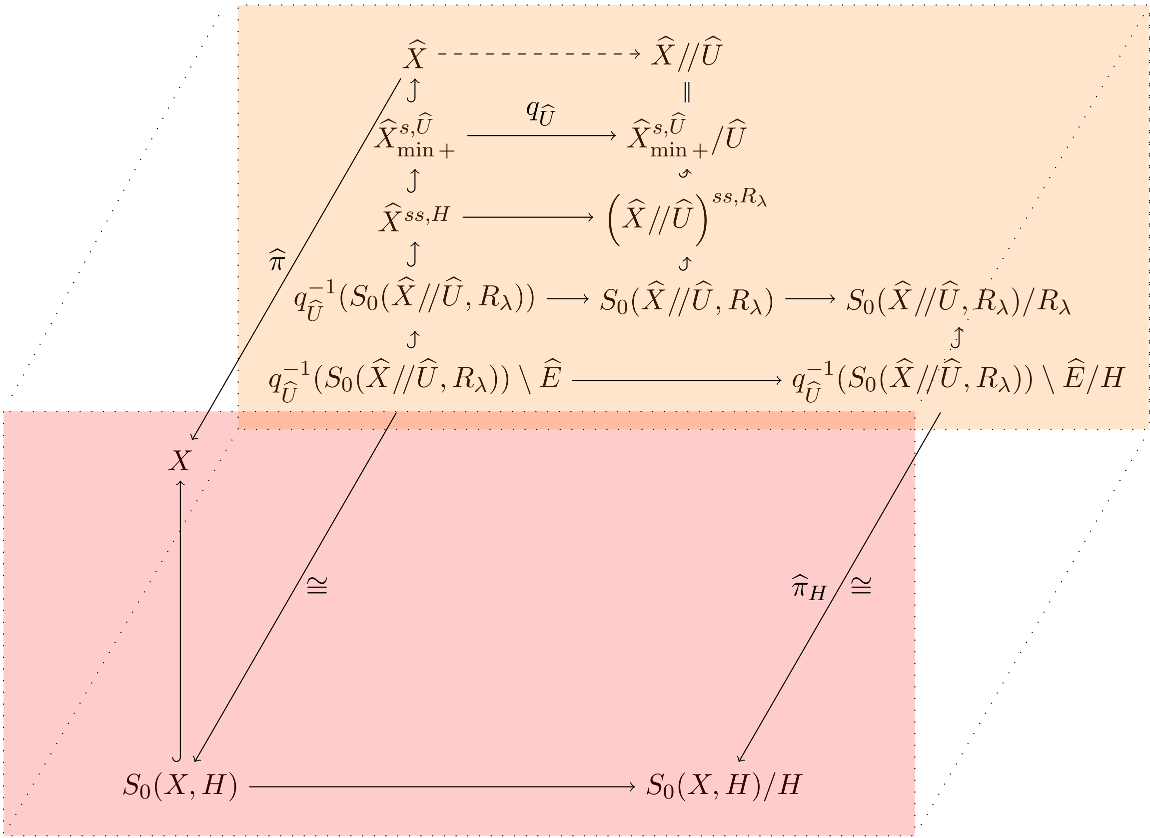

(a) Blow-up corresponding to Path 1 in Figure 1.11 (when $\widehat{X} \neq \overline{U\widehat{Z}_{\min}}$)

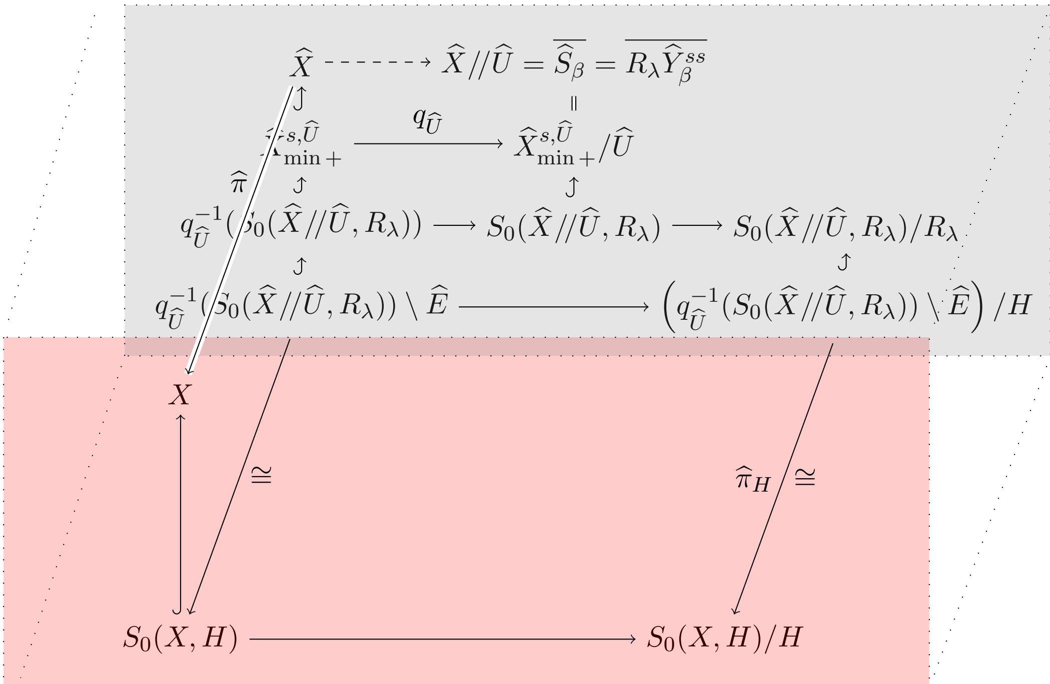

(b) Blow-up corresponding to Path 2 in Figure 1.11

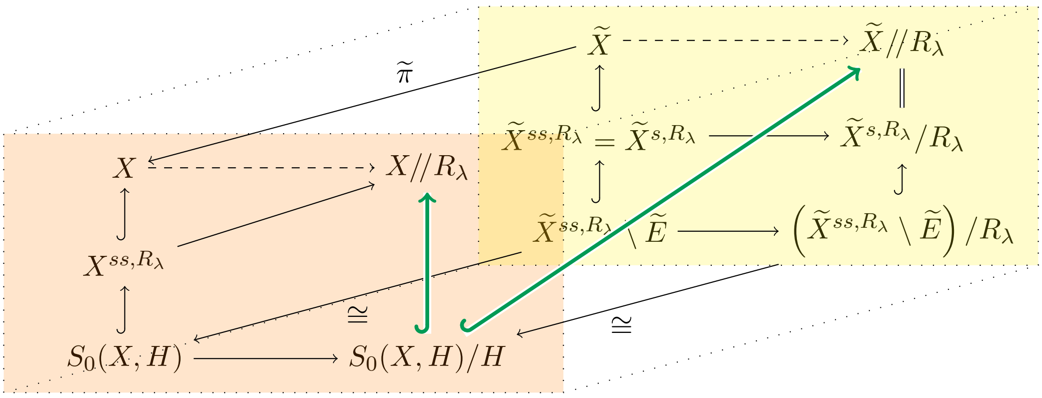

(c) Blow-up corresponding to Path $1'$ in Figure 1.11 (when $\widetilde{X} \neq \overline{U\widetilde{Z}_{\min}}$)

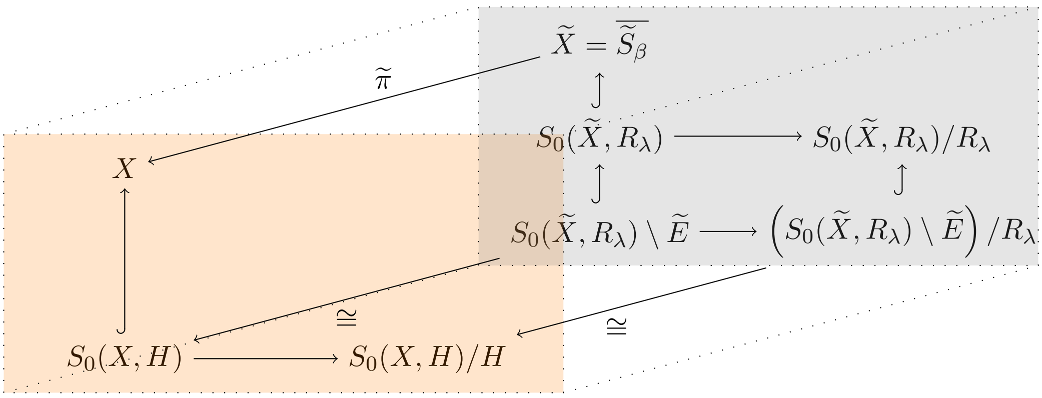

(d) Blow-up corresponding to Path $2'$ in Figure 1.11 (when $\widetilde{X} \neq \overline{U\widetilde{Z}_{\min}}$)

Figure 1.12: Blow-ups from the Projective Completion algorithm

correspond to performing a sequence of $\widehat{U}$-blow-ups, as per Theorem 1.2.3 (we recall that this construction, more precisely Blow-up Construction 2, depends on a number of choices, in addition to the choice of linearisation $\mathcal{L}$, see Notation 1.2.12). 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 linear action of $R_\lambda = R/\lambda(\mathbb{G}_m)$, and the semistable locus for this linear 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.

Case 1.1. If $(\widehat{X}/\!/\widehat{U})^{ss,R_\lambda} \neq \emptyset$, then path 1 must be taken, as illustrated in Figure 1.12a. By assumption, the 4-tuple $(\widehat{X}/\!/\widehat{U}, R_\lambda, \lambda_0, \widehat{\mathcal{L}})$ satisfies $(\star\star)$, where λ_0 denotes the trivial one-parameter subgroup of R_λ and $\widehat{\mathcal{L}}$ denotes a choice of linearisation for the induced action of R_λ on $\widehat{X}/\!/\widehat{U}$. 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 $S_0(\widehat{X}/\!/\widehat{U}, R_\lambda) \subseteq \widehat{X}/\!/\widehat{U}$ and the projective completion $\mathrm{PC}(S_0(\widehat{X}/\!/\widehat{U}, R_\lambda)/R_\lambda)$. Having done so, we set

$$S_0(X,H) := \widehat{\pi}\left(q_{\widehat{U}}^{-1}(S_0(\widehat{X}/\!/\widehat{U}, R_\lambda)) \setminus \widehat{E}\right) \tag{1.9}$$

where $\widehat{E}$ is the exceptional divisor corresponding to the sequence of $\widehat{U}$-blow-ups $\widehat{\pi} : \widehat{X} \to X$. Moreover, we define

$$\mathrm{PC}\left(S_0(X,H)/H\right) := \mathrm{PC}\left(S_0(\widehat{X}/\!/\widehat{U}, R_\lambda)/R_\lambda\right).$$

Case 1.2. If $(\widehat{X}/\!/\widehat{U})^{ss,R_\lambda} = \emptyset$, then path 2 must be taken, as illustrated in Figure 1.12b. 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 $S_0(\widehat{X}/\!/\widehat{U}, R_\lambda) \subseteq \widehat{X}/\!/\widehat{U}$ and the variety $\mathrm{PC}(S_0(\widehat{X}/\!/\widehat{U}, R_\lambda))$. Having done so, as in Case 1.1 we set

$$S_0(X,H) := \widehat{\pi}\left(q_{\widehat{U}}^{-1}(S_0(\widehat{X}/\!/\widehat{U}, R_\lambda)) \setminus \widehat{E}\right) \tag{1.10}$$

and define

$$\mathrm{PC}\left(S_0(X,H)/H\right) := \mathrm{PC}(S_0(\widehat{X}/\!/\widehat{U}, R_\lambda)/R_\lambda).$$

Case 2. Suppose that $(X, H, \lambda, \mathcal{L})$ satisfies $(\star\star)$. We start in the orange box of the right-hand side in Figure 1.11. According to the diagram two paths can be taken: $1'$ or $2'$. Both correspond to performing the sequence of blow-ups for the partial desingularisation construction associated to the linear action of R_λ on X and described in Section 1.1.1. 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 (we let $\widetilde{\mathcal{L}}$ denote the chosen linearisation) 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 linear action determines the path to be taken, as described in the following subcases.

Case 2.1. If $\widetilde{X}^{ss,R_\lambda} \neq \emptyset$, then path $1'$ must be taken, as illustrated in Figure 1.12c. Taking this path corresponds to applying the partial desingularisation construction 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

$$S_0(X, H) := \widetilde{\pi}\left(\widetilde{X}^{ss,R_\lambda} \setminus \widetilde{E}\right), \tag{1.11}$$

where $\widetilde{E}$ is the exceptional divisor corresponding to the blow-up $\widetilde{\pi} : \widetilde{X} \to X$, and

$$\mathrm{PC}\left(S_0(X, H)/H\right) := \widetilde{X}/\!/R_\lambda.$$

By assumption we have that $\widetilde{X}/\!/R_\lambda = \widetilde{X}^{ss,R_\lambda}/R_\lambda$ is a projective geometric quotient for the action of R_λ on $\widetilde{X}^{ss,R_\lambda}$. The algorithm terminates here.

Case 2.2. If $\widetilde{X}^{ss,R_\lambda} = \emptyset$, then path $2'$ must be taken, as illustrated in Figure 1.12d. Taking this path corresponds to applying the partial desingularisation construction 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 $(\widetilde{X}, R_\lambda, \lambda_0, \widetilde{\mathcal{L}})$ in order to construct the open set $S_0(\widetilde{X}, R_\lambda) \subseteq \widetilde{X}$ and the projective completion $\mathrm{PC}(S_0(\widetilde{X}, R_\lambda)/R_\lambda)$. Having done so, we set

$$S_0(X, H) := \widetilde{\pi}\left(S_0(\widetilde{X}, R_\lambda) \setminus \widetilde{E}\right)$$

where $\widetilde{E}$ is the exceptional divisor corresponding to the blow-up $\widetilde{\pi} : \widetilde{X} \to X$, and

$$\mathrm{PC}\left(S_0(X, H)/H\right) := \mathrm{PC}\left(S_0(\widetilde{X}, R_\lambda)/R_\lambda\right).$$

Case 3. 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 $(\star\star)$.

Case 3.1. 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 1.10). We then proceed to Case 1 applied to the 4-tuple $(X', H', \lambda', \mathcal{L}')$ in order to construct the open set $S_0(X', H') = S_0(\overline{Y^{ss}_{\beta(k)}}, P_{\beta(k)}) \subseteq \overline{Y^{ss}_{\beta(k)}}$ and the projective completion $\mathrm{PC}(S_0(X', H')/H')$. Having done so, we set

$$S_0(X,H) := R\left(L_\beta\left(\cdots\left(L_{\beta(k-1)}\left(S_0(\overline{Y^{ss}_{\beta(k)}}, P_{\beta(k)}) \cap Y^{ss}_{\beta(k)}\right)\cdots\right) Y^{ss}_{\beta(2)}\right) \cap Y_{\beta^{ss}}\right).$$

Moreover, we define

$$\mathrm{PC}\left(S_0(X,H)/H\right) := \mathrm{PC}\left(S_0(\overline{Y^{ss}_{\beta(k)}}, P_{\beta(k)})/P_{\beta(k)}\right).$$

Case 3.2. If $(X', H', \lambda', \mathcal{L}')$ satisfies $(\star\star)$ instead, then it is of the form $(\overline{Y^{ss}_{\beta(k)}}, P_{\beta(k)}/\lambda_{\beta(k)}(\mathbb{G}_m), \lambda_0, \mathcal{L}|_{\overline{Y^{ss}_{\beta(k)}}})$ 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 $S_0(X', H') = S_0(\overline{Y^{ss}_{\beta(k)}}, P_{\beta(k)}/\lambda_{\beta(k)}(\mathbb{G}_m)) \subseteq \overline{Y^{ss}_{\beta(k)}}$ and the variety $\mathrm{PC}(S_0(X', H')/H')$. Having done so, we set

$$S_0(X,H) = R\left(L_\beta\left(\cdots\left(L_{\beta(k-1)}\left(S_0(\overline{Y^{ss}_{\beta(k)}}, P_{\beta(k)}/\lambda_{\beta(k)}(\mathbb{G}_m)) \cap Y^{ss}_{\beta(k)}\right)\cdots\right) Y^{ss}_{\beta(2)}\right) \cap Y_{\beta^{ss}}\right)$$

and define

$$\mathrm{PC}\left(S_0(X,H)/H\right) := \mathrm{PC}\left(S_0(\overline{Y^{ss}_{\beta(k)}}, P_{\beta(k)}/\lambda_{\beta(k)}(\mathbb{G}_m))/(P_{\beta(k)}/\lambda_{\beta(k)}(\mathbb{G}_m))\right).$$

Proposition 1.3.3. The Projective Completion algorithm terminates after a finite number of steps.

Proof. Suppose that the initial 4-tuple is $(X, H, \lambda, \mathcal{L})$. By construction, the Projective Completion 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 $(\widehat{X}/\!/\widehat{U}, R_\lambda, \lambda_0, \widehat{\mathcal{L}})$ (or by $(\widehat{Z}_{\min}, R_\lambda, \lambda_0, \widehat{\mathcal{L}}|_{Z_{\min}})$ if $\widehat{X}^{s,\widehat{U}}_{\min+} = \emptyset$). 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

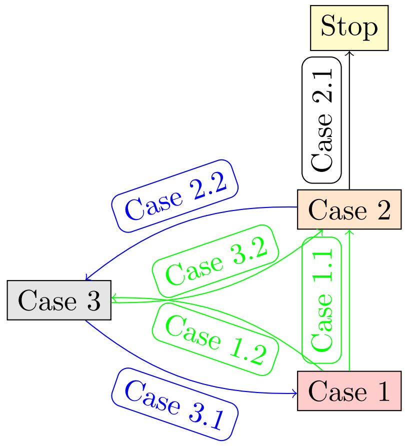

Figure 1.13: Proof that the Projective Completion algorithm must terminate. A green (respectively blue) arrow indicates that the dimension of the group decreases (respectively does not increase) upon passing from one case to the other. Since each loop contains at least one green arrow, when going through a loop the dimension of the group must always decrease.

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, for which the 4-tuple must be replaced 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 move between Cases 1, 2, and 3 infinitely many times. In Figure 1.13, 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 does not increase). However, as the diagram shows, every such closed loop contains at least one green arrow. Thus each time we circulate through a loop, the dimension of the group must decrease. If we did this ad infinitum, 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. □

Remark 1.3.4 (The Projective Completion algorithm is not deterministic)**.** The Projective Completion algorithm is not a deterministic algorithm: the initial data $(X, H, \lambda, \mathcal{L})$ does not uniquely determine the output. Indeed, many choices are involved throughout the algorithm. These include the choice of a suitable filtration of the unipotent radical whenever Blow-up Construction 2 is applied, as well as the choice of a linearisation at each stage of the construction. Nevertheless, we expect that there should exist, at each step of the algorithm where choices must be made, a suitable 'chamber' in a space of 'generalised linearisations' (which would encode the data needed to proceed to the next step of the algorithm) such that provided the 'generalised linearisation' is chosen within this chamber at each step, the resulting open subset $S_0(X, H, \lambda, \mathcal{L})$ and projective completion of $S_0(X, H, \lambda, \mathcal{L})/H$ are independent of the choices made throughout. This conjectural picture would be a generalisation of that described in Remark 1.2.9.

Remark 1.3.5 (Incorporating cohomological calculations into the Projective Completion algorithm)**.** In Chapter 2 we will develop formulae for computing the Poincaré series of classical and Non-Reductive GIT in terms of the equivariant Poincaré series of the initial variety and of certain of its subvarieties. As will be discussed in Remark 2.3.4, the results we obtain in Chapter 2 show that we could feasibly incorporate into the Projective Completion algorithm the calculation of the Poincaré series of the projective completion obtained as the output, in terms of the equivariant Poincaré series of the initial variety and of certain subvarieties, assuming that the initial variety is smooth (see also Remark 2.2.8).

1.3.3 Non-reductive GIT stratifications

By induction, the Projective Completion algorithm provides an algorithmic way of stratifying irreducible projective varieties acted upon by linear algebraic groups with internally graded unipotent radical, in such a way that each stratum admits a geometric quotient with an associated projective completion, as per Theorem 1.3.6 below (see [9, Thm. 1.1]).

Theorem 1.3.6. Let $H = U \rtimes R$ be a linear algebraic group with internally graded unipotent radical acting linearly on an irreducible projective variety X. Then, given the choice of an invariant inner product on $\operatorname{Lie} H$, there exists a stratification

$$X = \bigsqcup_{\gamma \in \Gamma} S_\gamma \tag{1.12}$$

of X by H-invariant quasi-projective subvarieties S_γ, called a *non-reductive GIT stratification of* X, satisfying the following properties:

(i) there exists a partial ordering on the index set Γ such that for every $\gamma \in \Gamma$,

$$\overline{S_\gamma} \subseteq \bigsqcup_{\gamma' \geq \gamma} S_{\gamma'};$$

(ii) each stratum S_γ has a quasi-projective geometric H-quotient S_γ/H and an associated projective completion $\mathrm{PC}(S_\gamma/H)$;

(iii) it is a refinement of the Bialynicki-Birula stratification[15] of X associated to the action of the grading one-parameter subgroup $\lambda : \mathbb{G}_m \to \mathrm{Z}(R)$;

(iv) if $H = G$ is a reductive group, then the stratification is a refinement of the GIT-instability stratification associated to the linear action of G on X and to the choice of inner product on $\operatorname{Lie} G$.

The main application of Theorem 1.3.6 is to classification problems in algebraic geometry which can be reduced to the problem of constructing the quotient of a variety by a linear algebraic group action. In particular, and as a consequence of Theorem 1.3.6 (iv), the theorem can be used to obtain complete solutions to classification problems involving reductive group actions. Indeed, for such classification problems, classical GIT provides a moduli space for stable objects; Non-Reductive GIT can be used to construct moduli spaces for the remaining strictly semistable and unstable objects. It is this perspective which will allow in Chapter 3 the construction of moduli spaces for unstable Higgs bundles.

[15]Given the linear action of $\mathbb{G}_m$ on a smooth projective variety X, the corresponding *Bialynicki-Birula stratification of* X is a canonical decomposition of X into $\mathbb{G}_m$-invariant locally closed subvarieties such that any such subvariety contains exactly one irreducible component of $X^{\mathbb{G}_m}$ (the fixed point locus for the $\mathbb{G}_m$-action) and retracts onto this component in such a way that the fibre of the retraction is $\mathbb{G}_m$-invariant and isomorphic to a vector space [13, Thm 4.3].

Chapter 2

Cohomology of quotients for linear algebraic group actions

2.0 Introduction

The geometry of moduli spaces via GIT. One of the most important applications of GIT is to the construction of moduli spaces when addressing classification problems in algebraic geometry. This application is possible because for many classification problems the construction of a moduli space can be reduced to the problem of constructing the quotient of a parameter space by a group. As we have seen in Chapter 1, when the group is either reductive or has an internally graded unipotent radical and the action can be linearised, then GIT (classical or Non-Reductive as appropriate) can be used to construct a quasi-projective orbit space for an explicitly determined open subset of the parameter space, as well as a projective completion of this orbit space. Moreover, in the best case where 'semistability coincides with stability' (in either the classical or non-reductive sense), the orbit space is in fact projective – a desirable feature for a moduli space.

But the application of GIT to classification problems in algebraic geometry is not limited to the construction of algebro-geometric moduli spaces: once GIT has been used to construct a moduli space, it can then be used to study the geometry of this space. This second step is paramount because the usefulness of moduli spaces in addressing classification problems lies not just in the algebro-geometric structure they give to the collection of objects they classify, but more so in the information that can be extracted from the study of their geometry. Indeed, the geometry of a moduli space encapsulates the essential features of the objects it classifies, and as a result information about the objects can be more easily extracted by studying the geometry of their moduli space, rather than by studying the objects themselves.

An important tool for describing the geometry of a space is its cohomology, and this is particularly useful for moduli spaces[1]. Although describing the cohomology of a moduli space is difficult in general, the task is made easier when the moduli space can be constructed using GIT. The reason is that for moduli spaces that are GIT quotients, the corresponding parameter spaces are often 'well-behaved spaces whose properties are relatively easy to understand', as noted in [71]. The task then reduces to relating the cohomology (we will always work with rational cohomology) of the parameter space to the cohomology of the GIT quotient, and this can be achieved through careful analysis of the structure of GIT quotients and their construction.

When semistability coincides with stability: the classical case. This analysis was first achieved by Kirwan in [64] for classical GIT in the case where semistability coincides with stability. It provides a method for computing the Poincaré series[2] of classical GIT quotients in terms of those of the parameter spaces when semistability coincides with stability. That is, given the linear action of a reductive group G on an irreducible smooth projective variety X, it is shown that the GIT-instability stratification is an equivariantly perfect stratification, so that the equivariant Poincaré series[3] of the semistable locus of X^{ss} can be expressed in terms of the equivariant Poincaré series of X and of the unstable strata. The former is often known since X is typically a well-behaved space, and results from [64] show that the latter can be determined inductively from the equivariant Poincaré series of semistable loci for the actions of lower-dimensional reductive groups on lower-dimensional smooth varieties.

When semistability coincides with stability, calculating the equivariant Poincaré series of

[1] For example, if we are interested in understanding a particular subclass of the objects classified, then one fruitful approach is to study the cohomological class of the corresponding geometric subvariety of the moduli space.

[2] The (rational) Poincaré series $P_t(X)$ of a topological space X is defined by

$$P_t(X) = \sum_{i \geq 0} t^i \dim H^i(X, \mathbb{Q}).$$

[3] If a topological space X admits the action of a group G, then the (rational) equivariant Poincaré series $P_t^G(X)$ of X is defined by

$$P_t^G(X) = \sum_{i \geq 0} t^i \dim H_G^i(X, \mathbb{Q}).$$

Here $H_G^i(X, \mathbb{Q})$ denotes the i-th equivariant rational cohomology group of X, which is defined by

$$H_G^i(X, \mathbb{Q}) := H^i(EG \times_G X; \mathbb{Q})$$

where EG is the total space of the universal bundle over BG, the classifying space of G.

the semistable locus X^{ss} is equivalent to calculating the Poincaré series of the GIT quotient $X//G$. Indeed, in this case the action of G on X^{ss} has at worst finite stabiliser groups so that the equivariant rational cohomology of the semistable locus is isomorphic to the rational cohomology of the GIT quotient. Thus results from [64] provide an inductive formula for computing the Poincaré series of GIT quotients of irreducible smooth projective varieties by the linear actions of reductive groups.

When semistability coincides with stability: the non-reductive case. As we have seen in Chapter 1, if an assumption analogous to the assumption from classical GIT that semistability coincides with stability (a condition denoted by ($ss = s \neq \emptyset[\widehat{U}]$)) is made, then GIT can effectively be extended to actions of linear algebraic groups with internally graded unipotent radical[4]. As in the classical case, we can therefore seek to describe the geometry of the resulting Non-Reductive GIT quotient in terms of that of the parameter space. This is achieved in the recent paper [12] by Bérczi and Kirwan.

If a graded extension $\widehat{U}$ of a unipotent group U acts linearly on an irreducible smooth projective variety X, it is shown in [12] that the equivariant Poincaré series of the semistable locus can be computed via stratification, as in the reductive case. However the situation is simpler in the non-reductive case: the formula is not inductive. This is because for the $\widehat{U}$-action there is a distinguished subvariety $Z_{\min}$ of X which carries much of the cohomological information of $X//\widehat{U}$.

If the condition ($ss = s \neq \emptyset[\widehat{U}]$) is satisfied, then the stabilisers of points in the semistable locus are finite and thus, as in the classical case, the Poincaré series of the GIT quotient coincides with the equivariant Poincaré series of the semistable locus. As a result, [12] provides a formula for computing the Poincaré series of $X//\widehat{U}$ in terms of the Poincaré series of $Z_{\min}$. More generally, if a linear algebraic group $H = U \rtimes R$ with internally graded unipotent radical acts linearly on an irreducible smooth projective variety X, so that ($ss = s \neq \emptyset[\widehat{U}]$) is satisfied for the action of $\widehat{U} \subseteq H$ and so that semistability coincides with stability for the induced action of $R/\lambda(\mathbb{G}_m)$ on $X//\widehat{U}$, then the formulae from [64] and [8] can be combined to produce an

[4] A key difference nevertheless is that the linearisation must be amended in a suitable way to obtain finite generation of the invariants.

inductive formula for computing the Poincaré series of the GIT quotient $X//H$[5].

When blow-ups are required: the reductive case. In both the reductive and non-reductive cases considered above, the assumption that 'semistability coincides with stability' is essential for identifying the cohomology of the GIT quotient with the cohomology of the semistable locus, since it ensures that stabilisers are at most finite. When this condition is not satisfied, the isomorphism no longer holds.

Nevertheless, as we have seen in Section 1.2 of Chapter 1, when the condition that 'semistability coincides with stability' (in either the classical or non-reductive sense) is not satisfied, then both in the classical and non-reductive case a sequence of blow-ups can be performed to obtain a variety for which the condition is satisfied[6]. Hence for the resulting blown-up variety, the formulae described above apply. By combining these formulae with formulae for the cohomology of blow-ups, it is then reasonable to try to obtain a formula for the Poincaré series of the GIT quotient of the blown-up variety in terms of the equivariant Poincaré series of the initial variety and of certain distinguished subvarieties.

This was first achieved in the classical case by Kirwan in [66]. Given the action of a reductive group G on an irreducible smooth projective variety X, and given $\widetilde{X}$ the variety resulting from the partial desingularisation construction[7], a formula is obtained for the Poincaré series of the quotient $\widetilde{X}//G$ in terms of the G-equivariant Poincaré series of X^{ss} and of the centres of the blow-ups, which can in turn be computed inductively from the equivariant Poincaré series of X. The prerequisite for these calculations is proving that the centres of the blow-ups occurring in the partial desingularisation construction are smooth, so that $\widetilde{X}$ is also smooth.

Later in [68], this formula was extended to provide the intersection Poincaré series of the original GIT quotient $X//G$ in terms of the G-equivariant Poincaré series of X^{ss}, by making use of the induced blow-down map $\widetilde{X}//G \to X//G$.

[5]In fact, in [12] Bérczi and Kirwan delve further into the analysis of the cohomology of the GIT quotient $X//H$: in addition to giving a formula for its Poincaré series, they also describe the rational cohomology ring of $X//H$ in terms of the rational cohomology ring of the GIT quotient $X//T$ where T is a maximal torus in H.

[6]The condition that the stable locus is non-empty (in the classical case), or that the analogous condition in the non-reductive case is satisfied, is necessary for the blow-up construction to terminate in this way.

[7]Again the assumption that the stable locus is non-empty must be made.

When blow-ups are required: the non-reductive case. In the non-reductive case, we have seen in Section 1.2 of Chapter 1 that when the condition that 'semistability coincides with stability' (namely ($ss = s \neq \emptyset[\widehat{U}]$)) is not satisfied for the action of an externally graded unipotent group $\widehat{U}$ on an irreducible projective variety X, then a sequence of blow-ups can be performed to obtain a variety $\widehat{X}$ for which ($ss = s \neq \emptyset[\widehat{U}]$) is satisfied[8]. Thus we obtain a projective quotient $\widehat{X}/\!/\widehat{U}$, and we can ask whether, when X is smooth, the Poincaré series of $\widehat{X}/\!/H$ can be computed in terms of the cohomological information about X, as can be done in the reductive case.

The aim of this chapter is to show that such a computation can be achieved. As in the classical case, the key result required to enable these computations is that the centres of the blow-ups occurring in the blow-up constructions of Non-Reductive GIT are smooth if the initial variety is smooth. The proof of the corresponding result in the classical case does not generalise in a straightforward way to the non-reductive case because it relies heavily on the reductivity of the group. As a result, in the non-reductive setting we must devise a set-up which allows us to reduce the proof to the reductive case. Once this result is proved, it can be used to obtain explicit formulae for the Poincaré series of blown-up Non-Reductive GIT quotients.

Structure of this chapter In Section 2.1 we summarise the results concerning the computation of the Poincaré series of GIT quotients, both in the classical and Non-Reductive setting, when the group acts on an irreducible smooth projective variety such that 'semistability coincides with stability' is satisfied. In Sections 2.2 and 2.3 we consider the more general case where semistability does not coincide with stability. Section 2.2 focuses on the centres of the blow-ups occurring in the blow-up constructions. After reviewing the proof that they are smooth in the classical case, we prove that they are also smooth in the non-reductive case (see Theorem 2.2.4). In Section 2.3 we then show how, thanks to this result, formulae can be obtained for the Poincaré series of the blown-up quotients in terms of cohomological information about X (see Theorem 2.3.2 and Corollary 2.3.5). Again we first review the reductive case, and then show how analogous computations can be made in the non-reductive setting.

[8] In this non-reductive setting, the assumption that there exists a point in $Z_{\min}$ with trivial stabiliser must be made. This condition is analogous to the condition in classical GIT that the stable locus is non-empty.

2.1 When semistability coincides with stability

As we have seen in Section 1.1 of Chapter 1, in both classical and Non-Reductive GIT the optimal scenario occurs when the condition that 'semistability coincides with stability'[9] is satisfied. That is to say, if this condition is satisfied then the projective GIT quotient, obtained from the (finitely generated) algebra of invariants[10], is an orbit space with at worst finite stabiliser groups for the action of the group on an open subset (the semistable locus) of the initial projective variety which can be described explicitly. In particular, if the initial variety is smooth, then the GIT quotient will have at worst finite quotient singularities. This outcome is particularly desirable for describing the cohomology of GIT quotients. Indeed, in this case the rational cohomology of the GIT quotient (either classical or Non-Reductive) coincides with the equivariant rational cohomology of the semistable locus, and thus it suffices to compute the latter.

The aim of this Section 2.1 is to summarise known results about the computation of the equivariant Poincaré series of the semistable loci appearing in both classical and Non-Reductive GIT, and to present the resulting formulae for the Poincaré series of the GIT quotient when 'semistability coincides with stability'. In Section 2.1.1 we consider the reductive case, in Section 2.1.2 the externally graded unipotent case and finally in Section 2.1.3 the general case for linear algebraic groups with internally graded unipotent radical.

2.1.1 Cohomology of reductive quotients

Given the linear action of a reductive group G on an irreducible smooth projective variety X, a method is introduced in [64] for inductively computing the G-equivariant Poincaré series of the semistable locus X^{ss}. When semistability coincides with stability this method gives the Poincaré series of the GIT quotient $X//G$. The method of [64] can be summarised as follows.

Given a choice of Weyl-invariant inner product on the Lie algebra of a maximal torus of G, it can be shown using techniques from symplectic geometry (see [64, Thm 5.4]) that the associated GIT-instability stratification (see Section 1.1.1) $X = \bigsqcup_{\beta\in\mathcal{B}} S_\beta$ is equivariantly perfect. By

[9]In the classical case, this condition means that the semistable locus is equal to the stable locus (i.e. $X^{ss} = X^s$); in the externally graded unipotent case it means that points in $Z_{\min}$ have trivial unipotent stabiliser (i.e. the condition denoted ($ss = s \neq \emptyset[\widehat{U}]$)); in the general case for linear algebraic groups with internally graded unipotent radical, it is a combination of these two conditions, namely that ($ss = s \neq \emptyset[\widehat{U}]$) is satisfied as well as the condition that the semistable and stable loci coincide for the residual reductive action on the $\widehat{U}$-quotient.

[10]We recall that in the non-reductive case, the linearisation may need to be suitably amended to obtain finite generation of invariants, see Theorem 1.1.5.

definition (see [64, §2.16]), this means that

$$P_t^G(X) = P_t^G(X^{ss}) + \sum_{\beta \in \mathcal{B} \setminus \{0\}} t^{2d(\beta)} P_t^G(S_\beta), \tag{2.1}$$

where $d(\beta)$ is the complex codimension of S_β in X. For simplicity we have assumed that all of the connected components for a given S_β have the same dimension (so that $d(\beta)$ is well-defined). If this is not the case, the formula needs modification (see [64, §8.12]).

The Poincaré series $P_t^G(X)$ and $P_t^G(S_\beta)$ can be further simplified. For the former, we can use the fact that if X is an irreducible smooth projective variety acted upon by a reductive group G, then $P_t^G(X) = P_t(X)P_t(BG)$ (see [64, Prop 5.8]). For the latter, we must recall the structure of the unstable strata S_β. As seen in Section 1.1.1, there is an isomorphism $S_\beta \cong G \times_{P_\beta} Y_\beta^{ss}$ where Y_β^{ss} retracts onto Z_β^{ss} via a locally trivial fibration with contractible fibre compatible with the action of $\operatorname{Stab}\beta$ (by definition Z_β^{ss} consists of the semistable locus for the linear action of $\operatorname{Stab}\beta$ on Z_β induced by that of G on X). Moreover, P_β is homotopically equivalent to $\operatorname{Stab}\beta$. As a result, for each stratum S_β there is an isomorphism of rational cohomology groups $H_G^*(S_\beta, \mathbb{Q}) \cong H^*_{\operatorname{Stab}\beta}(Z_\beta^{ss}, \mathbb{Q})$ from which it follows that $P_t^G(S_\beta) = P_t^{\operatorname{Stab}\beta}(Z_\beta^{ss})$. Thus we obtain an inductive formula for the G-equivariant Poincaré series of X^{ss}:

$$P_t^G(X^{ss}) = P_t(X)P_t(BG) + \sum_{\beta \in \mathcal{B} \setminus \{0\}} t^{2d(\beta)} P_t^{\operatorname{Stab}\beta}(Z_\beta^{ss}). \tag{2.2}$$

If moreover the semistable and stable loci coincide, the action of G on X^{ss} has at worst finite stabiliser groups so that the equivariant rational cohomology of the semistable locus coincides with the rational cohomology of the GIT quotient. Thus we have that

$$P_t(X/\!/G) = P_t^G(X^{ss}) = P_t(X)P_t(BG) + \sum_{\beta \in \mathcal{B} \setminus \{0\}} t^{2d(\beta)} P_t^{\operatorname{Stab}\beta}(Z_\beta^{ss}). \tag{2.3}$$

Remark 2.1.1 (When X is singular and semistability coincides with stability)**.** The case where X is singular and $X^{ss} = X^s$ is treated in [69, §2]: it is shown that a version of (2.3) can be obtained in the singular case by considering intersection Poincaré series instead. Intersection cohomology is considered instead because it is better suited to the study of singular projective varieties than ordinary cohomology (in particular it satisfies a version of Poincaré duality, which is not the case for ordinary cohomology of singular varieties). However, in the singular case the inductive formula is more complicated because it must take into account how the singularities meet the various strata. Nevertheless if they meet transversely then (2.3) is valid simply by considering intersection Poincaré series instead.

2.1.2 Cohomology of externally graded unipotent quotients

A similar approach to that described in Section 2.1.2 is adopted by Bérczi and Kirwan in [12] to obtain a formula for the Poincaré series of Non-Reductive GIT quotients of the form $X/\!/\widehat{U}$ when ($ss = s \neq \emptyset[\widehat{U}]$) is satisfied. However, the situation is simpler in this case than in the reductive case: thanks to the existence of a distinguished subvariety $Z_{\min}$ of X which carries cohomological information about $X^0_{\min}$, an explicit, rather than inductive, formula can be obtained. The approach of [12] can be summarised as follows.

Suppose that an externally graded unipotent group $\widehat{U}$ acts linearly on an irreducible smooth projective variety X, such that the condition ($ss = s \neq \emptyset[\widehat{U}]$) is satisfied. Then by Theorem 1.1.5 (the $\widehat{U}$-theorem), there exists a projective quotient $X/\!/\widehat{U}$ which is a geometric quotient for the action of $\widehat{U}$ on the semistable locus $X^{ss,\widehat{U}}_{\min +} := X^0_{\min} \setminus UZ_{\min}$. By [12, Cor 5.4], the stratification

$$X^0_{\min} = X^{ss,\widehat{U}}_{\min +} \sqcup UZ_{\min} \tag{2.4}$$

of $X^0_{\min}$ is equivariantly perfect so that

$$P^{\widehat{U}}_t(X^0_{\min}) = P^{\widehat{U}}_t(X^{ss,\widehat{U}}_{\min +}) + t^{2d} P^{\widehat{U}}_t(UZ_{\min})$$

where $d = \dim X - \dim U - \dim Z_{\min}$ is the complex codimension of $UZ_{\min}$ in $X^0_{\min}$. Moreover, since the condition ($ss = s \neq \emptyset[\widehat{U}]$) is satisfied, the stabiliser groups of points in $X^{ss,\widehat{U}}_{\min +}$ are finite and thus as in the classical case we have that $P_t(X/\!/\widehat{U}) = P^{\widehat{U}}_t(X^{ss,\widehat{U}}_{\min +})$.

The fact that $X^0_{\min}$ retracts onto $Z_{\min}$ and that $\widehat{U}$ is homotopically equivalent to the grading $\mathbb{G}_m$ allows further simplification of the formula for $P^{\widehat{U}}_t(X^{ss,\widehat{U}}_{\min +})$. Indeed, it implies that $P^{\widehat{U}}_t(X^0_{\min}) = P^{\mathbb{G}_m}_t(Z_{\min})$, and since $\mathbb{G}_m$ acts trivially on $Z_{\min}$, there is an equality

$$P^{\mathbb{G}_m}_t(Z_{\min}) = P_t(Z_{\min}) P_t(B\mathbb{G}_m) = P_t(Z_{\min}) \frac{1}{1-t^2}.$$

Thus replacing and rearranging the terms of (2.4), we obtain that

$$P_t(X/\!/\widehat{U}) = P_t(Z_{\min}) \frac{1-t^{2d}}{1-t^2}. \tag{2.5}$$

This formula contrasts with (2.3) in the reductive case where an inductive procedure is needed to compute $P_t(X//G) = P_t^G(X^{ss})$; for the $\widehat{U}$-action there is a distinguished subvariety $Z_{\min}$ of X which carries much of the cohomological information of $X//\widehat{U}$.

2.1.3 Cohomology of non-reductive quotients

As we have seen in Section 1.1.3 of Chapter 1, if $H = U \rtimes R$ is a linear algebraic group with internally unipotent radical U acting linearly on an irreducible projective variety X, such that ($ss = s \neq \emptyset[\widehat{U}]$) is satisfied, then a GIT quotient $X//H$ can be constructed in stages, first by quotienting by $\widehat{U}$, and then by quotienting by the residual reductive group $R/\lambda(\mathbb{G}_m)$ where $\lambda(\mathbb{G}_m) \subseteq Z(R)$ is the grading one-parameter subgroup.

If we moreover assume that semistability coincides with stability for the action of $R/\lambda(\mathbb{G}_m)$ on $X//\widehat{U}$, then the formulae (2.3) and (2.5) from the above Sections 2.1.1 and 2.1.2 for the Poincaré series of classical and $\widehat{U}$-GIT quotients can be combined to produce an inductive formula for the Poincaré series of GIT quotients by linear algebraic groups H with internally graded unipotent radical. Indeed, under this condition we obtain:

$$\begin{aligned} P_t(X//H) &= P_t\left((X//\widehat{U})//R/\lambda(\mathbb{G}_m)\right) = P_t^{R/\lambda(\mathbb{G}_m)}\left((X//\widehat{U})^{ss,R/\lambda(\mathbb{G}_m)}\right) && (2.6)\\ &= P_t^{R/\lambda(\mathbb{G}_m)}(X//\widehat{U})P_t(R/\lambda(\mathbb{G}_m)) + \sum_{\beta\in\mathcal{B}\setminus\{0\}} t^{2d(\beta)}P_t^{\mathrm{Stab}\,\beta}(Z(X//\widehat{U})_\beta^{ss}) && \text{by (2.3)}\\ &= P_t(Z_{\min})P_t(R/\lambda(\mathbb{G}_m))\frac{1-t^{2d}}{1-t^2} + \sum_{\beta\in\mathcal{B}\setminus\{0\}} t^{2d(\beta)}P_t^{\mathrm{Stab}\,\beta}(Z(X//\widehat{U})_\beta^{ss}) && \text{by (2.5).}\end{aligned}$$

The above formula for $P_t(X//H)$ requires understanding the GIT-instability stratification for the action of $R/\lambda(\mathbb{G}_m)$ on the intermediate quotient $X//\widehat{U}$ (the choice of an invariant inner product on H induces one on $R/\lambda(\mathbb{G}_m)$), which may be difficult in practice (for example if the intermediate quotient has no obvious modular interpretation, in the case where we are using GIT to construct a moduli space).

For this reason, a different approach is given in [12] for computing the Poincaré series of the GIT quotient $X//H$, resulting in a formula which depends on information only about X, rather than about the intermediate quotient $X//\widehat{U}$ as well. However, this approach requires a stronger assumption on the action of H on X: in addition to assuming that the condition ($ss = s \neq \emptyset[\widehat{U}]$) is satisfied and that semistability coincides with stability for the action of $R/\lambda(\mathbb{G}_m)$ on $X//\widehat{U}$, we must further assume that semistabiliy coincides with stability for the action of $R/\lambda(\mathbb{G}_m)$ on $Z_{\min}$ (this action is well-defined since $\lambda(\mathbb{G}_m)$ acts trivially on $Z_{\min}$).

Under this assumption, by [12, Lem 5.6] we have that

$$X//H = \left(p^{-1}(Z_{\min}^{ss,R/\lambda(\mathbb{G}_m)}) \setminus U Z_{\min}^{ss,R/\lambda(\mathbb{G}_m)}\right)/H.$$

The stratification approach from Section 2.1.2, this time applied to $p^{-1}(Z_{\min}^{ss,R/\lambda(\mathbb{G}_m)})$ instead of $X^0_{\min}$, can then be used to obtain that

$$P_t(X//H) = P_t(Z_{\min}//R/\lambda(\mathbb{G}_m))\frac{1-t^{2d}}{1-t^2}. \tag{2.7}$$

The Poincaré series $P_t(Z_{\min}//R/\lambda(\mathbb{G}_m))$ can in turn be computed using (2.3). The advantage of (2.7) over (2.6) is that instead of having to describe the GIT-instability stratification for the action of $R/\lambda(\mathbb{G}_m)$ on $X//\widehat{U}$, it suffices to understand the GIT-instability stratification for the action of $R/\lambda(\mathbb{G}_m)$ on $Z_{\min}$, thus bypassing the need to study the intermediate quotient $X//\widehat{U}$.

2.2 When semistability does not coincide with stability: smoothness of the centres of the blow-ups

When the condition 'semistability coincides with stability' is not satisfied in either the classical or non-reductive case, then as we have seen in Section 1.2 a sequence of blow-ups can be performed to produce a variety for which the condition is satisfied. Under suitable conditions on the initial action, the GIT quotient obtained in the blow-up can be thought of as a projective completion of the geometric quotient of the stable locus of the initial variety.

The aim of this section and Section 2.3 below is to describe formulae for the Poincaré series of the GIT quotient in the blow-up in terms of cohomological information about the initial variety, assumed to be smooth. In order to obtain these formulae, it is necessary to first show that if we start with a smooth variety, then the variety resulting from the blow-up construction is also smooth. This ensures that the GIT quotient of the blow-up has at worst finite quotient singularities, and thus that its rational cohomology coincides with the equivariant rational cohomology of the semistable locus.

In the case of classical GIT, smoothness of the centres of the blow-ups appearing in the partial desingularisation construction was proven by Kirwan in [66]. The aim of this Section 2.2

is to prove the analogous result for the blow-up construction of Non-Reductive GIT. That is, we prove Theorem 2.2.4 which states that the centres of the blow-ups of Blow-up Construction 2 are smooth if the initial variety is smooth. As with many results of classical GIT, the proof that the centres of the blow-ups for the partial desingularisation construction are smooth relies heavily on the reductivity of the group. As a result, the proof does not extend in a straightforward way to the non-reductive case. Instead, an elaborate construction is required to reduce to the reductive setting.

We start by reviewing in Section 2.2.1 the proof of smoothness in the reductive setting. The proof in the non-reductive setting is split into two parts: Section 2.2.2 and Section 2.2.3. In Section 2.2.2 we prove the result in the special case where $U = \mathbb{G}_a$ (see Theorem 2.2.2), and in Section 2.2.3 we prove the result in the general case (see Theorem 2.2.4). We prove the general result by reducing to proving it under a simplifying assumption (see Proposition 2.2.7). This is the assumption that there exists a point in $X^0_{\min}$ fixed by the whole unipotent radical, and we can prove the result under this assumption by generalising the proof of Theorem 2.2.2.

2.2.1 The reductive case

In this section we summarise the proof given in [66] that the centres of the blow-ups of the partial desingularisation construction of classical GIT are smooth if the initial variety is smooth. This is indeed a prerequisite to obtaining a formula for the Poincaré series of the GIT quotient of the blow-up in terms of cohomological information about X; we introduce this formula in Section 2.3.1.

The centres of the blow-ups. Suppose that a reductive group G acts linearly on an irreducible smooth projective variety X such that the stable locus is non-empty and such that semistability does not coincide with stability. Then, as we have seen in Section 1.2.1, the partial desingularisation construction can be performed. The construction results in a projective variety $\widetilde{X}$ with a linear action of G for which semistability coincides with stability and such that the stable locus is non-empty.

To explain how smoothness of the centres of the blow-ups is established, we give a more detailed description of the construction than that given in Section 1.2.1. Let $d_{\max}$ denote the maximal dimension amongst the dimensions of connected reductive subgroups of G arising as the stabiliser groups of points in X^{ss}. Then, let $\mathcal{R}(d_{\max}) = \{R_1, \ldots, R_k\}$ denote a set of conjugacy classes of these reductive subgroups of G. The reductivity of G ensures that this set is finite (see [75]). The partial desingularisation construction requires first blowing X^{ss} up along $G(X^{R_1} \cap X^{ss})$; we note that the subvariety $G(X^{R_1} \cap X^{ss})$ is closed in X^{ss} by [66, Lem 5.11][11]. The resulting blow-up $X_{(1)}$ has a linear action of G, induced by a perturbation (by a small multiple of the exceptional divisor) of the pull-back of the linearisation of the G-action on X via the blow-down map. Importantly, there is no subgroup of G conjugate to R_1 which fixes a point in $X^{ss}_{(1)}$ (see [66, Lem 6.1]).

The next step is to blow $X^{ss}_{(1)}$ up along $G(X^{R_2}_{(1)} \cap X^{ss}_{(1)})$, and so on for increasing i. Once this has been done for all representatives $R_1, \ldots, R_k$, we repeat the entire process for the new maximal dimension of connected reductive stabiliser group for points in $X^{ss}_{(k)}$. Since this new maximal dimension is strictly smaller than $d_{\max}$ (this follows from [66, Lem 6.1] referred to above), after repeating the process a finite number of times we obtain the desired variety $\widetilde{X}$ with zero-dimensional connected reductive stabiliser groups in $\widetilde{X}^{ss}$.

Remark 2.2.1 (Analogy with the non-reductive case)**.** By [66, Lem 8.2], the resulting variety $X_{(1)}$ (and by iteration the resulting variety $\widetilde{X}$) is independent of the order given to the representatives in $\mathcal{R}(d_{\max})$. This relies on the fact that if R and R' lie in different conjugacy classes, then $GX^R \cap GX^{R'} \cap X^{ss} = \emptyset$. As a result, if we let

$$C_{\max}(X^{ss}, G) := \{x \in X^{ss} \mid \dim \mathrm{Stab}_G(x) = d_{\max}\},$$

then we have an equality

$$C_{\max}(X^{ss}, G) = \bigsqcup_{R \in \mathcal{R}(d_{\max})} G(X^R \cap X^{ss}). \tag{2.8}$$

Thus the result of blowing X^{ss} up along $G(X^R \cap X^{ss})$ for each $R \in \mathcal{R}(d_{\max})$ iteratively is the same as blowing X^{ss} up along $C_{\max}(X^{ss}, G)$.

This perspective makes the analogy with the blow-up constructions of Non-Reductive GIT clearer. Indeed, in the non-reductive case, the variety is blown up iteratively along loci of points

[11]To be exact, the partial desingularisation construction of [66] requires blowing X up along the closure of $G(X^{R_1} \cap X^{ss})$ inside X, after having resolved its singularities. There is a linear G-action on this blow-up of X: it is given by a perturbation by a small multiple of the exceptional divisor of the pull-back of the linearisation of the G-action on X. If the perturbation is small enough, then the pre-image of X^{ss} in the blow-up coincides with the blow-up $X_{(1)}$ of X^{ss} along $G(X^{R_1} \cap X^{ss})$.

with a fixed dimension of unipotent stabilisers, rather than iteratively along the sweep of fixed point sets for specific unipotent stabiliser groups as can be done in the reductive setting. The reason why the latter method does not work in the unipotent case is that the set of conjugacy classes of subgroups of a unipotent group of a fixed dimension may not be finite, in contrast to the reductive case (see [75] referred to above).

Smoothness of the centres of the blow-ups. The proof that the centres of the blow-ups of the partial desingularisation construction, which are all of the form $G(X^R \cap X^{ss})$ for R a connected reductive subgroup of G acting on an irreducible smooth projective variety X, are smooth is achieved in two steps: first by proving that $X^R \cap X^{ss}$ is smooth, then by proving that its G-sweep is smooth.

For the first step, since $X^R \cap X^{ss}$ is open in X^R it suffices to show that X^R is smooth. But it is a well-known result that if a reductive group R acts on a smooth projective variety X, then the fixed point set[12] X^R is a closed smooth subvariety of X (see for example [57, Prop 1.3]).

For the second step, namely to show that $G(X^R \cap X^{ss})$ is smooth given that $X^R \cap X^{ss}$ is smooth, the following isomorphism induced by the group action map is established (see [66, Cor 5.10]):

$$G \times_N (X^R \cap X^{ss}) \cong G(X^R \cap X^{ss}) \tag{2.9}$$

where N is the normaliser in G of R and $G \times_N (X^R \cap X^{ss}) = (G \times (X^R \cap X^{ss}))/N$ with N acting by right multiplication by the inverse on G and by the standard left action on X^{ss}. Indeed, since $G \times (X^R \cap X^{ss})$ is smooth and N acts freely on this product, the quotient is also smooth. This argument concludes the proof that the centres of the blow-ups involved in the partial desingularisation construction are smooth if the initial variety is smooth.

To extend this result to the non-reductive case, we will prove a non-reductive analogue of the first step described above, namely that if an externally graded unipotent group $\widehat{U}$ acts linearly on an irreducible smooth projective variety X, then $X^U \cap X^0_{\min}$ is smooth (see Theorem 2.2.2). Just as the smoothness of the centres of the blow-ups of the partial desingularisation construction can be proved from the smoothness of $X^R \cap X^{ss}$, as per the second step described above, we will show that the smoothness of the centres of the blow-ups from Non-Reductive GIT can be

[12]The scheme-theoretic structure of fixed point sets for group actions on schemes is given in [28].

proved from the smoothness of $X^U \cap X^0_{\min}$ (see the proof of Theorem 2.2.4). Doing so requires a different strategy to the second step described above, since $C_{\max}(X^0_{\min}, \widehat{U})$ cannot be written as a finite union of sets of the form $\widehat{U}(X^{U'} \cap X^0_{\min})$ where U' is a representative of a conjugacy class for subgroups of $\widehat{U}$ of a fixed dimension arising as the stabiliser group of a point in $X^0_{\min}$ (indeed, the set of such conjugacy classes may not be finite, as mentioned in Remark 2.2.1 above).

2.2.2 The simplest non-reductive case: when $U = \mathbb{G}_a$

In the above Section 2.2.1 we outlined the proof of the smoothness of the centres of the blow-ups for the partial desingularisation construction of classical GIT. The $\widehat{U}$-theorem with blow-ups (Theorem 1.2.3) involves a blow-up construction which is analogous to this partial desingularisation construction (see Remark 2.2.1): the variety is also blown up along loci of points with maximal dimension stabiliser group. The aim of the present Section 2.2.2 and Section 2.2.3 below is to prove that the centres of these non-reductive blow-ups are also smooth if the initial variety is smooth. In this section we prove the result in the case where $U = \mathbb{G}_a$ (see Theorem 2.2.2) while in Section 2.2.3 below we generalise the proof to obtain the result in the general case (see Theorem 2.2.4).

Given the linear action of $\widehat{U} := \mathbb{G}_a \rtimes \mathbb{G}_m$ on an irreducible projective variety X, if the condition ($ss = s \neq \emptyset[\widehat{U}]$) or ($ss = Ms \neq \emptyset[\widehat{U}]$) is not satisfied, then there must exist a point in $Z_{\min}$ with trivial stabiliser in $\mathbb{G}_a$. Indeed, if not then all points in $X^0_{\min}$ are fixed by $\mathbb{G}_a$, in which case ($ss = Ms \neq \emptyset[\widehat{U}]$) is satisfied and we need not perform blow-ups to obtain a projective quotient. Thanks to the existence of such a point in $Z_{\min}$, we can apply the alternative Blow-up Construction 1 (see Remark 1.2.8); we wish to show that the centres of the blow-ups involved in this construction are smooth if X is smooth. That is, we want to show that if X is smooth, then $C_{\max}(X^0_{\min}, \widehat{U})$ is smooth. In fact, we prove the following more general result:

Theorem 2.2.2 (Smoothness of the centres of the blow-ups when $U = \mathbb{G}_a$)**.** Let $\widehat{U} = \mathbb{G}_a \rtimes \mathbb{G}_m$ where $\mathbb{G}_m$ acts on $\operatorname{Lie} \mathbb{G}_a$ with a positive weight. Suppose that $\widehat{U}$ acts linearly on an irreducible projective variety X and let $x \in C_{\max}(X^0_{\min}, \widehat{U})$. Then $C_{\max}(X^0_{\min}, \widehat{U})$ is smooth at x if X is smooth at x. In particular, if X is smooth then $C_{\max}(X^0_{\min}, \widehat{U})$ is a smooth subvariety of X.

Remark 2.2.3 (Simpler description of $C_{\max}(X^0_{\min}, \widehat{U})$ when $U = \mathbb{G}_a$)**.** When $U = \mathbb{G}_a$, if all points in $X^0_{\min}$ have trivial stabiliser in $\mathbb{G}_a$, then $C_{\max}(X^0_{\min}, \widehat{U})$ coincides with $UZ_{\min}$, which is smooth if X is smooth. Thus we can assume that there exists a point in $X^0_{\min}$ fixed by $\mathbb{G}_a$, in which case $C_{\max}(X^0_{\min}, \widehat{U})$ is equal to $X^{\mathbb{G}_a} \cap Z_{\min}$. Hence to prove Theorem 2.2.2 it suffices to show that $X^{\mathbb{G}_a} \cap Z_{\min}$ is smooth.

Our proof of Theorem 2.2.2 relies on the following set-up:

Set-up (Representations of $\widehat{U}$ and of $\mathrm{GL}(2;\mathbb{C})$)**.** Let $\widehat{U} := \mathbb{G}_a \rtimes \mathbb{G}_m$ where $\mathbb{G}_m$ acts with a positive weight w on $\mathrm{Lie}\,\mathbb{G}_a$ via the adjoint action. The coadjoint action of $\widehat{U}$ on $(\mathrm{Lie}\,\widehat{U})^\vee$ gives a representation $\rho : \widehat{U} \to \mathrm{GL}((\mathrm{Lie}\,\widehat{U})^\vee)$. The restriction of ρ to $\mathbb{G}_a$ is an embedding because of the positive grading of $\mathbb{G}_m$ on $\mathrm{Lie}\,\mathbb{G}_a$. Moreover, we can choose a basis for $(\mathrm{Lie}\,\widehat{U})^\vee$ to obtain an isomorphism $\mathrm{GL}((\mathrm{Lie}\,\widehat{U})^\vee) \cong \mathrm{GL}(2;\mathbb{C})$ such that the image of $(u,t) \in \widehat{U}$ under ρ is given by the matrix $\begin{pmatrix} 1 & u \\ 0 & t^{-w} \end{pmatrix}$. Here we have used that $\mathbb{G}_m$ acts trivially on $(\mathrm{Lie}\,\mathbb{G}_m)^\vee$ since $\mathbb{G}_m$ is abelian.

Let $\sigma : \mathrm{GL}(2) \to \mathrm{GL}(\mathrm{Sym}^2(\mathbb{C}^2))$ denote the standard representation of $\mathrm{GL}(2;\mathbb{C})$ on $\mathrm{Sym}^2(\mathbb{C}^2)$, which under a suitable identification of $\mathrm{GL}(\mathrm{Sym}^2(\mathbb{C}^2))$ with $\mathrm{GL}(3;\mathbb{C})$ is given by

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \mapsto \begin{pmatrix} a^2 & 2ab & b^2 \\ ac & bc+ad & bd \\ c^2 & 2cd & d^2 \end{pmatrix}.$$

We let $\widetilde{\sigma} : \mathrm{Mat}_{2\times 2}(\mathbb{C}) \to \mathrm{Mat}_{3\times 3}(\mathbb{C})$ denote the natural extension of this map.

The composition $\sigma \circ \rho$ gives a representation of $\widehat{U}$ on $\mathrm{Sym}^2(\mathbb{C}^2)$, and we consider the representation obtained by twisting the representation $\sigma\circ\rho$ by the character of $\rho(\widehat{U})$ corresponding to the restriction of the determinant character of $\mathrm{GL}(2;\mathbb{C})$. This ensures that the image of $\sigma\circ\rho$ lies in $\mathrm{SL}(3;\mathbb{C})$ after twisting. We let $\widetilde{\rho}$ denote the resulting representation of $\widehat{U}$ on $\mathrm{Sym}^2(\mathbb{C}^2) \cong \mathbb{C}^3$; it is given by

$$(u,t) \mapsto \begin{pmatrix} t^w & 2ut^w & u^2t^w \\ 0 & 1 & u \\ 0 & 0 & t^{-w} \end{pmatrix}.$$

The above representations play an important role in the proof of Theorem 2.2.2 below.

Proof of Theorem 2.2.2. By Remark 2.2.3, it suffices to show that $X^{\mathbb{G}_a} \cap Z_{\min}$ is smooth at any point at which X is smooth. We define the projective variety $W := \mathbb{P}(\mathrm{End}(\mathrm{Sym}^2(\mathbb{C}^2)) \oplus \mathbb{C})$, which we identify with the variety $\mathbb{P}(\mathrm{Mat}_{3\times 3}(\mathbb{C})\oplus\mathbb{C})$ by fixing an isomorphism $\mathrm{Sym}^2(\mathbb{C}^2)\cong\mathbb{C}^3$. The variety W admits a linear action of $\widehat{U}$ given by $(u,s)\cdot[M{:}v]=[M\widetilde{\rho}(u,s)^{-1}{:}sv]$ for any $(u,s)\in\widehat{U}$ and $[M:v]\in W$, where $\widetilde{\rho}:\widehat{U}\to\mathrm{GL}(3;\mathbb{C})$ is as per Set-up 2.2.2. By definition of the action of $\widehat{U}$ on W, we have that

$$Z(W)_{\min}=\left\{\left[\begin{pmatrix} a & 0 & 0\\ b & 0 & 0\\ c & 0 & 0\end{pmatrix}:0\right]\;\middle|\; a,b,c\in\mathbb{C}\right\}. \tag{2.10}$$

Moreover, points in $Z(W)_{\min}$ have trivial unipotent stabiliser groups. Thus the condition ($ss=s\neq\emptyset[\widehat{U}]$) is satisfied for the linear action of $\widehat{U}$ on W.

Now let $X':=W\times X$. Then X' has a natural $\widehat{U}$-action induced by that on W and on X, and we can consider the linearisation of this $\widehat{U}$-action given by taking the tensor product of the pull-back to X' of the linearisation on W with the pull-back to X' of the linearisation on X. We start by constructing a projective geometric quotient for the linear action of $\widehat{U}$ on X' using the results from Non-Reductive GIT introduced in Section 1.1.2. Let $Z'_{\min}$ denote the analogue for X' of $Z_{\min}$, so that $Z'_{\min}=Z(W)_{\min}\times Z_{\min}$. Since points in $Z(W)_{\min}$ have trivial unipotent stabiliser groups, it follows that points in $Z'_{\min}$ also have trivial unipotent stabiliser groups. Thus by Theorem 1.1.5, after twisting the linearisation of the $\widehat{U}$-action on X' and taking a sufficiently large tensor power, we obtain a projective geometric quotient

$$\pi: X'^{0}_{\min}\setminus UZ'_{\min}\to Y':=X'/\!/\widehat{U}.$$

If X' is smooth at $x\in X'^{0}_{\min}\setminus UZ'_{\min}$, then Y' has at worse a finite quotient singularity at $\pi(x)$. Moreover, if X' is smooth at $([M{:}1],x)\in X'^{0}_{\min}\setminus UZ'_{\min}$, then Y' is smooth at $\pi(x)$. Indeed, any such pair has trivial stabiliser group in $\widehat{U}$, and thus its image under the quotient map cannot inherit a finite quotient singularity.

We let Y denote the subvariety of Y' given by points of the form $\widehat{U}\cdot([M{:}v],x)$ where M lies in the image of $\widetilde{\sigma}$ (see Set-up 2.2.2). Then Y is a geometric quotient for the action of $\widehat{U}$ on the intersection of $X'^{0}_{\min}\setminus UZ'_{\min}$ with the locus in X' consisting of pairs of the form $([M{:}v],x)$ where M is in the image of $\widetilde{\sigma}$. Similarly to Y', we have that Y is smooth at points of the form $\widehat{U}\cdot([M{:}v],x)$ with $v\neq 0$ and such that X is smooth at x.

We now fix a point $x_0\in X^{\mathbb{G}_a}\cap Z_{\min}$ such that X is smooth at x_0. Our aim is to show that $X^{\mathbb{G}_a}\cap Z_{\min}$ is smooth at x_0. Our strategy for doing so is to relate $X^{\mathbb{G}_a}\cap Z_{\min}$ in a small neighbourhood of x_0 to the fixed point set in Y for the maximal torus T in $\mathrm{GL}(2;\mathbb{C})$, in a small neighbourhood of a corresponding point $y_0 := \widehat{U} \cdot ([M_0{:}1], x_0)$ for a suitably chosen matrix M_0. The advantage of doing so is that T is reductive, and we can therefore establish using classical results that Y^T is smooth at y_0.

To achieve this aim, we need in particular to define for any $x \in Z_{\min}$ a corresponding point $y = \widehat{U} \cdot ([M_0{:}1], x) \in Y'$ such that y is fixed by $T \subseteq \mathrm{GL}(2;\mathbb{C})$ if and only if x is fixed by $\mathbb{G}_a$. As we will now show, any matrix M_0 of the form

$$M_0 := \widetilde{\sigma}\left(\begin{pmatrix} a & b \\ 0 & 0 \end{pmatrix}\right) = \begin{pmatrix} a^2 & 2ab & b^2 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

with a and b non-zero achieves this aim. That is, we show that given $x \in X^0_{\min}$, the point $y := \widehat{U} \cdot ([M_0{:}1], x)$ lies in Y^T if and only if x lies in $X^{\mathbb{G}_a} \cap Z_{\min}$, given M_0 of the above form.

First we must check that $([M_0{:}1], x)$ lies in $X'^0_{\min} \setminus UZ'_{\min}$, so that y is well-defined as a point in Y. Since the entry of M_0 in the first column is non-zero, it follows from the description of $Z(W)_{\min}$ given in (2.10) that $M_0 \in W^0_{\min}$. Thus we need only check that $([M_0{:}1], x)$ does not lie in $UZ'_{\min} = U(Z(W)_{\min} \times Z_{\min})$. But this follows from the fact that $x \in Z_{\min}$ by assumption, whereas $M_0 \notin Z(W)_{\min}$ since both a and b are non-zero.

Thus $y = \widehat{U} \cdot ([M_0{:}1], x)$ is a well-defined orbit in Y. By definition of the action of $\mathrm{GL}(2;\mathbb{C})$ on Y and of the action of $\widehat{U}$ on X, we have that y is fixed by T if and only if for every $t_1, t_2 \in \mathbb{G}_m$, there exists an element $(u,s)^{-1} \in \widehat{U}$ such that $\left[\sigma\left(\begin{pmatrix} t_1 & 0 \\ 0 & t_2 \end{pmatrix}\right) M_0{:}1\right] = \left[M_0\widetilde{\rho}(u,s){:}s^{-1}\right]$ and such that $(u,s)^{-1} \cdot x = x$. In other words, $y \in Y^T$ if and only if for every $t_1, t_2 \in \mathbb{G}_m$ there exists an element $(u,s)^{-1} \in \mathrm{Stab}_{\widehat{U}}(x)$ such that

$$\sigma\left(\begin{pmatrix} t_1 & 0 \\ 0 & t_2 \end{pmatrix}\right) M_0 = sM_0\widetilde{\rho}(u,s). \tag{2.11}$$

By definition of σ and $\widetilde{\rho}$, the equality of (2.11) is equivalent to the following:

$$\begin{pmatrix} t_1^2a^2 & t_1^2 2ab & t_1^2 b^2 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} = \begin{pmatrix} s^{w+1}a^2 & 2us^{w+1}a^2 + 2abs & u^2s^{w+1}a^2 + 2abus + s^{-w+1}b^2 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

Equality of the first two non-zero entries of the matrices implies that $(u,s) \in \mathrm{Stab}_{\widehat{U}}(x)$ must satisfy $s^{w+1} = t_1^2$ and $u = \frac{b}{a}\left(\frac{t_1^2 - s}{t_1^2}\right)$. Under these assumptions on u and s, we obtain that the third non-zero entries of the matrices (namely the top right entry) also coincide; this can be seen by observing that $m_{13} = \frac{m_{12}^2}{4m_{11}}$ where m_{11}, m_{12} and m_{13} denote the first, second and third entries respectively of the above equal matrices.

Thus we have that $y = \widehat{U} \cdot ([M_0 : 1], x)$ lies in Y^T if and only if for every $t_1, t_2 \in \mathbb{G}_m$, there exists a pair $(u,s)^{-1} \in \mathrm{Stab}_{\widehat{U}}(x)$ such that $s^{w+1} = t_1^2$ and $u = \frac{b}{a}\left(\frac{t_1^2 - s}{t_1^2}\right)$. We now show that the latter condition is equivalent to the condition that x lies in $X^{\mathbb{G}_a} \cap Z_{\min}$. For the implication, we use the fact that $(u,s)^{-1} = ((u,1)\cdot(0,s)^{-1}) = (0,s^{-1})\cdot(-u,1) = (-s^w u, 1)\cdot(0,s)$ (this follows from the grading of $\mathbb{G}_m$ on $\mathrm{Lie}\,\mathbb{G}_a$) and that $x \in Z_{\min}$ by assumption, so that $(u,s)^{-1}\cdot x = x$ if and only if $(-s^w u, 1)\cdot x = x$. Since

$$s^w u = \frac{b}{a}\left(\frac{s^w - s^{w+1}}{t_1^2}\right) = \frac{b}{a}\left(\frac{t_1^2 s^{-1} - t_1^2}{t_1^2}\right) = \frac{b}{a}(s^{-1} - 1),$$

we obtain that for every $t_1 \in \mathbb{G}_m$, the point x is fixed by $(-\frac{b}{a}(s^{-1}-1), 1) \in \mathbb{G}_a$ where $s^{w+1} = t_1^2$. It follows that $x \in X^{\mathbb{G}_a} \cap Z_{\min}$. Conversely, we note that if $x \in X^{\mathbb{G}_a} \cap Z_{\min}$, then clearly $(u,s) \in \mathrm{Stab}_U(x)$ for any (u,s) satisfying $s^{w+1} = t_1^2$ and $u = \frac{b}{a}\left(\frac{t_1^2 - s}{t_1^2}\right)$. This shows that $y = \widehat{U} \cdot ([M_0{:}1], x)$ lies in Y^T if and only if x lies in $X^{\mathbb{G}_a} \cap Z_{\min}$.

To conclude the proof, we consider again the point $x_0 \in X^{\mathbb{G}_a} \cap Z_{\min}$ fixed above, and set $y_0 := \widehat{U} \cdot ([M_0{:}1], x_0) \in Y$ for a fixed choice of $M_0 = \widetilde{\sigma}\left(\begin{pmatrix} a & b \\ 0 & 0 \end{pmatrix}\right)$ with a and b non-zero. It follows from the above analysis that $y_0 \in Y^T$. We note that Y is smooth at y_0 (since the scalar coordinate is non-zero and we have assumed that X is smooth at x_0), and therefore Y is smooth in a small neighbourhood of y_0. Since T is reductive, by [57, Prop 1.3] we have that Y^T is also smooth in a neighbourhood of y_0. Moreover, the above analysis also shows that the preimage of this neighbourhood under the geometric quotient map π is a neighbourhood of $([M_0{:}1], x_0)$ contained in $W \times (X^{\mathbb{G}_a} \cap Z_{\min})$, which must also be smooth since π is a geometric quotient. Therefore $X^{\mathbb{G}_a} \cap Z_{\min}$ must be smooth at x_0. □

2.2.3 The general non-reductive case

The aim of this section is to generalise Theorem 2.2.2 to the case where U has dimension greater than one. That is, we prove that the centres of the blow-ups from the $\widehat{U}$-theorem with blow-ups (Theorem 1.2.3) are smooth, provided the variety we start with is smooth.

Suppose that $\widehat{U} = U \rtimes \mathbb{G}_m$ where U is a unipotent algebraic group and $\mathbb{G}_m$ acts on $\mathrm{Lie}\,U$ with positive weights. Suppose moreover that $\widehat{U}$ acts linearly on an irreducible projective variety X in such a way that neither ($ss = s \neq \emptyset[\widehat{U}]$) nor ($ss = Ms \neq \emptyset[\widehat{U}]$) is satisfied, so that blow-ups must be performed to obtain a projective quotient. The blow-up construction which applies without any further assumptions on the action is Blow-up Construction 2, thus it is for this construction that we will prove that the centres of the blow-ups are smooth. We recall that Blow-up Construction 2 relies on a choice of a normal series $U = U^{(0)} \supseteq U^{(1)} \supseteq \cdots \supseteq U^{(s)} = \{e\}$ satisfying the properties of (†) (see the second bullet point of Theorem 1.2.3). The construction then consists in combining the alternative Blow-up Construction 1 with the procedure of quotienting in stages. More precisely, the first step is to blow $X \times \mathbb{P}^1 \times \cdots \times \mathbb{P}^1$ (with $\mathbb{P}^1$ repeated $s-1$ times) up along points in $\left(X \times \mathbb{P}^1 \times \cdots \times \mathbb{P}^1\right)^0_{\min}$ with maximal dimension stabiliser in $\widehat{U}^{(s-1)}$, that is, along the subvariety $C_{\max}((X \times \mathbb{P}^1 \times \cdots \times \mathbb{P}^1)^0_{\min}, \widehat{U}^{(s-1)})$. At each subsequent stage of the construction, we are blowing up a variety Y, with an action by a semi-direct product $\widehat{U'}$ of an abelian unipotent group U' with a $\mathbb{G}_m$ acting on the Lie algebra of the unipotent group with a single positive weight, along the subvariety $C_{\max}(Y^0_{\min}, \widehat{U}')$.

By construction, if we start off with a smooth variety X, then the variety Y occurring at each stage will also be smooth (or at worse have finite quotient singularities), since it is a geometric quotient of a smooth variety. Thus to prove that the centres of the blow-ups defined in Blow-up Construction 2 are smooth, it suffices to prove the following theorem:

Theorem 2.2.4 (The centres of the blow-ups are smooth)**.** Let $\widehat{U} := U \rtimes \mathbb{G}_m$ where U is an abelian unipotent group and $\mathbb{G}_m$ acts with a single positive weight on $\operatorname{Lie} U$ via the adjoint action. Suppose that $\widehat{U}$ acts linearly on an irreducible projective variety X. Then $C_{\max}(X^0_{\min}, \widehat{U})$ is smooth at $x \in C_{\max}(X^0_{\min}, \widehat{U})$ if X is smooth at x. In particular, if X is smooth then $C_{\max}(X^0_{\min}, \widehat{U})$ is a closed smooth subvariety of $X^0_{\min}$.

Remark 2.2.5 (Closedness of $C_{\max}(X^0_{\min}, \widehat{U})$)**.** The fact that the locus $C_{\max}(X^0_{\min}, \widehat{U})$ is closed in $X^0_{\min}$ follows from the standard result regarding upper semi-continuity of dimensions (in this case applied to stabiliser dimension), refer to [39, §13.1] for example.

Remark 2.2.6 (Dealing with finite quotient singularities)**.** The reason for studying the centres of the blow-ups from Blow-up Construction 2 is to show that provided the initial variety X is smooth, then the formula (2.5) from Section 2.1.2 applies to the action of $\widehat{U}$ on the resulting blown-up variety $\widehat{X}$. Since the action of $\widehat{U}$ on $\widehat{X}$ satisfies the condition that ($ss = s \neq \emptyset[\widehat{U}]$), then to be able to apply (2.5) only requires showing that $\widehat{X}$ is smooth, or that it has at worst finite quotient singularities (since we are working with rational cohomology). Theorem 2.2.4 is not quite sufficient on its own to prove this. Indeed, as noted above before the statement of Theorem 2.2.4, if we start with a smooth variety, then the variety at any given stage is not necessarily smooth: it may have finite quotient singularities. And while Theorem 2.2.4 implies that the centres of the blow-ups are smooth at any stage where the variety is smooth, it does not a priori apply at a stage where the variety has a finite quotient singularity. Thus Theorem 2.2.4 does not suffice in general to conclude that if X is smooth, then $\widehat{X}$ is smooth or has at worst finite quotient singularities.

Nevertheless, if the filtration of $\widehat{U}$ satisfying (†) has only two terms (i.e. $s = 1$), then it does follow from Theorem 2.2.4 that the variety $\widehat{X}$ resulting from Blow-up Construction 2 is also smooth, since in this case the construction does not require taking intermediate quotients (these intermediate quotients are indeed the only way in which finite quotient singularities can arise). We note that this condition on the filtration will be satisfied in our application to Higgs bundles of rank 2, which is considered in Section 4.3 of Chapter 4. In fact, this application was our motivation for obtaining the results in this Chapter 2.

In general, to show that if X is smooth then the variety $\widehat{X}$ resulting from applying Blow-up Construction 2 is smooth or has at worst finite quotient singularities, we would need to generalise Theorem 2.2.4 by replacing the smoothness condition on X by the condition that it has at worst finite quotient singularities. The expected conclusion from this generalisation would be that the centre of the blow-up does not have any singularities, other than those finite quotient singularities inherited from the ambient variety X.

A natural way to formulate this more general statement is in the language of irreducible smooth Deligne-Mumford stacks, which locally can be identified as quotient stacks for the action of a finite group on a smooth variety. We expect that our proof of Theorem 2.2.4 will carry over to this setting by working locally. More precisely, we would show smoothness of the centre of the blow-up inside the Deligne-Mumford stack at any given point by identifying a neighbourhood of the point as a quotient stack of a smooth variety by a finite group, and applying Theorem 2.2.4 to this variety. We hope to prove this more general result in future work.

Our strategy for proving Theorem 2.2.4 is to reduce to proving the result in a special case, namely in the case where $d_{\max}(X^0_{\min}, \widehat{U}) = \dim \widehat{U}$. Under this assumption, we have that $C_{\max}(X^0_{\min}, \widehat{U}) = X^{\widehat{U}} \cap X^0_{\min} = X^U \cap UZ_{\min} = X^U \cap Z_{\min}$, so that it suffices to show that $X^U \cap Z_{\min}$ is smooth at points at which X is smooth. This is the content of Proposition 2.2.7 below.

Proposition 2.2.7. Let $\widehat{U} := U \rtimes \mathbb{G}_m$ where U is an abelian unipotent group and $\mathbb{G}_m$ acts with a single positive weight on $\operatorname{Lie} U$ via the adjoint action. Suppose that $\widehat{U}$ acts linearly on an irreducible projective variety X and let $x \in X^U \cap Z_{\min}$. Then the subvariety $X^U \cap Z_{\min}$ is smooth at x if X is smooth at x.

Before proving Proposition 2.2.7, we first show how Theorem 2.2.4 can be deduced from Proposition 2.2.7.

Proof of Theorem 2.2.4 assuming Proposition 2.2.7. By (1.8) of Remark 1.2.8, we have an equality $C_{\max}(X^0_{\min}, \widehat{U}) = C_{\max}(X^0_{\min}, U) \cap UZ_{\min}$. Suppose that X is smooth at a point $x_0 \in C_{\max}(X^0_{\min}, U) \cap UZ_{\min}$. Our aim is to show that $C_{\max}(X^0_{\min}, U) \cap UZ_{\min}$ is smooth at x_0.

Let $U' = \operatorname{Stab}_U(x_0)$. Since U is abelian, we can choose a complementary subgroup $U'^{\perp}$ of U' in U, so that $U = U' \times U'^{\perp}$. Then, for x in a small enough neighbourhood N_0 of x_0 in $X^0_{\min}$ (which we can take to be invariant under $U'^{\perp}$), we have that $U'^{\perp}$ is complementary to $\operatorname{Stab}_U(x)$ if and only if $x \in C_{\max}(X^0_{\min}, U)$. Moreover the action of $U'^{\perp}$ on N_0 has trivial stabilisers, so that N_0 is contained in the stable locus for the action of $U'^{\perp}$ (by [8, Thm 8.16 (1)]). As a result, by [8, Prop 7.1] we obtain a geometric quotient $\pi : N_0 \to Y := N_0/U'^{\perp}$, which has an induced linear action of $\widehat{\overline{U}} = \widehat{U}/U'^{\perp} \cong U/U'^{\perp} \rtimes \mathbb{G}_m$.

We now show that given $x \in N_0$ and $y = \pi(x) \in Y$, we have that $x \in C_{\max}(X^0_{\min}, U) \cap UZ_{\min}$ if and only if $y \in Y^{\widehat{\overline{U}}}$. If $x \in N_0$, then by the construction of N_0 we can assume that $U'^{\perp}$ is complementary to $\operatorname{Stab}_U(x)$. Thus if $x \in C_{\max}(X^0_{\min}, U) \cap UZ_{\min}$ then $y = \pi(x)$ must be fixed by $U/U'^{\perp}$. This follows from the fact that there is an equality

$$\operatorname{Stab}_{U/U'^{\perp}}(\pi(x)) = (U'^{\perp} \operatorname{Stab}_U(x))/U'^{\perp}.$$

In fact, y is fixed by the whole group $\widehat{\overline{U}} = \widehat{U}/U'^{\perp}$, since x is also fixed by a $\mathbb{G}_m$ inside $\widehat{U}$. To show the converse, suppose that $x \in N_0$ does not have maximal dimension stabiliser group in U.

Again using the fact that $\mathrm{Stab}_{U/U'^{\perp}}(\pi(x)) = (U'^{\perp}\,\mathrm{Stab}_U(x))/U'^{\perp}$, we see that y cannot be fixed by all of $U/U'^{\perp}$, and therefore y is not fixed by $\widehat{\widehat{U}}$. This shows that $x \in C_{\max}(X^0_{\min}, U) \cap UZ_{\min}$ if and only if $y \in Y^{\widehat{\widehat{U}}}$.

Since π is a geometric quotient with trivial stabilisers, it follows that $C_{\max}(X^0_{\min}, U) \cap UZ_{\min}$ is smooth at x_0 if and only if $Y^{\widehat{\widehat{U}}}$ is smooth at y_0. We note that since by assumption X is smooth at x_0, we also have that Y is smooth at y_0.

By [8, pp 33-34], there exists an integer $s > 0$ such that Y admits a $\widehat{\widehat{U}}$-equivariant locally closed embedding into the projective space $\mathbb{P}(W^\vee)$ where $W = H^0(X, L^{\otimes s})^U$ (here L denotes the ample line bundle on X with respect to which the action is linearised). By [8, Lem 7.6], the image of the locally closed embedding $Y \to \mathbb{P}(W^\vee)$ is contained in $\mathbb{P}(W^\vee)^0_{\min}$, and this map restricts to a closed embedding $Y \to \mathbb{P}(W^\vee)^0_{\min}$. We can therefore identify Y as a subvariety of $\mathbb{P}(W^\vee)^0_{\min}$. We let $\overline{Y}$ denote the closure of Y in this projective space. Then $\overline{Y}$ is a projective variety with a linear action of the externally graded unipotent group $\widehat{\widehat{U}}$.

Since $y_0 \in Y$, we have that $y_0 \in \overline{Y}^0_{\min} = \overline{Y} \cap \mathbb{P}(W^\vee)^0_{\min}$. Moreover, from above we have that $y_0 \in Y^{\widehat{\widehat{U}}}$ since $x_0 \in C_{\max}(X^0_{\min}, U) \cap UZ_{\min}$. Therefore $y_0 \in \overline{Y}^{\widehat{\widehat{U}}} \cap \overline{Y}^0_{\min} = \overline{Y}^{U/U'^{\perp}} \cap Z(\overline{Y})_{\min}$. Since Y is smooth at y_0, the projective variety $\overline{Y}$ is also smooth at y_0. We can therefore apply Proposition 2.2.7 to conclude that $\overline{Y}^{U/U'^{\perp}} \cap Z(\overline{Y})_{\min} = \overline{Y}^{\widehat{\widehat{U}}} \cap \overline{Y}^0_{\min}$ is smooth at y_0. And since $Y^{\widehat{\widehat{U}}} \cap \overline{Y}^0_{\min}$ is an open subset of $\overline{Y}^{\widehat{\widehat{U}}}$, it follows that $Y^{\widehat{\widehat{U}}}$ is smooth at y_0. This implies by our argument that $C_{\max}(X^0_{\min}, \widehat{U})$ is smooth at x_0. □

Thus to establish Theorem 2.2.4 it remains only to prove Proposition 2.2.7.

Proof of Proposition 2.2.7. We note that the statement we must show is the generalisation of Theorem 2.2.2 to the case where $\dim U > 1$. We prove this by generalising the proof of Theorem 2.2.2.

For the set-up, we let $d = \dim U$, and identify U with $\mathbb{G}_a^d$, using the fact that U is abelian. Under this identification, we can consider the embedding $\widehat{U} \cong \mathbb{G}_a^d \rtimes \mathbb{G}_m \hookrightarrow \widehat{\mathbb{G}}_a \times \cdots \times \widehat{\mathbb{G}}_a$ given by $(u_1, \ldots, u_d, s) \mapsto ((u_1, s), \ldots, (u_d, s))$. We note that the product $\widehat{\mathbb{G}}_a \times \cdots \times \widehat{\mathbb{G}}_a$ can be identified with $\widehat{U}_T := U \rtimes T$ where $T = \mathbb{G}_m^d$.

In the proof of Theorem 2.2.2, we defined the projective variety $W := \mathbb{P}(\mathrm{Mat}_{3\times 3}(\mathbb{C}) \oplus \mathbb{C})$. We now consider the product W^d of d copies of W, on which we define a linear action of $\widehat{\mathbb{G}}_a \times \cdots \times \widehat{\mathbb{G}}_a$

as follows: the element $((u_1, s_1), \ldots, (u_d, s_d))$ acts on the i-th factor of W^d via (u_i, s_i) as defined in the proof of Theorem 2.2.2, that is, via multiplication on the right by $\widetilde{\rho}(u_i, s_i)^{-1}$ on matrices, and by multiplication by s_i on the scalar coordinate. Just as in the $d = 1$ case, we obtain that the condition ($ss = s \neq \emptyset[\widehat{U}]$) is satisfied for this action.

Each copy of $\widehat{\mathbb{G}}_a$ inside the product $\widehat{\mathbb{G}}_a \times \cdots \times \widehat{\mathbb{G}}_a$ acts on X via the restriction to $\widehat{\mathbb{G}}_a \subseteq \widehat{U}$ of the $\widehat{U}$-action on X; we note that a point in X is fixed by $\widehat{U}$ if and only if it is fixed by $\widehat{\mathbb{G}}_a \times \cdots \times \widehat{\mathbb{G}}_a$. We consider the natural linear action of $\widehat{\mathbb{G}}_a \times \cdots \times \widehat{\mathbb{G}}_a$ on the product $X' = W^d \times X$ induced by the linear action on each factor, which also satisfies ($ss = s \neq \emptyset[\widehat{U}]$). Thus we have a projective geometric quotient

$$\pi : X'^{ss,\widehat{U}_T}_{\min +} \to Y' := X' /\!/ \widehat{U}_T.$$

Generalising the proof from the $d = 1$ case, we define an action of $(\mathrm{GL}(2;\mathbb{C}))^d$ on Y', given by multiplication on the left for the matrices, and by the trivial action on the scalars and on X. As in the $d = 1$ case, we define Y to be the subvariety of Y' given by points of the form $\widehat{U}_T \cdot ([M_1{:}v_1], \ldots, [M_d{:}v_d], x)$ such that each M_i lies in the image of $\widetilde{\sigma}$ (see Set-up 2.2.2). We again have that the variety Y is a geometric quotient and that it is smooth at those points $\widehat{U}_T \cdot ([M_1{:}v_1], \ldots, [M_d{:}v_d], x)$ such that $v_i \neq 0$ for all $i = 1, \ldots, d$ and such that X is smooth at x.

We now fix a point $x_0 \in X^U \cap Z_{\min}$ such that X is smooth at x_0; we wish to show that $X^U \cap Z_{\min}$ is smooth at x_0. As in the $d = 1$ case, we let $M_0 = \begin{pmatrix} a^2 & 2ab & b^2 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$ for some $a, b \neq 0$. Given any $x \in Z_{\min}$, we set $y = \widehat{U}_T \cdot ([M_0{:}1], \ldots, [M_0{:}1], x)$. The first step is to show that y is a well-defined point in Y, for any choice of $x \in Z_{\min}$. Since M_0 is in the image of $\widetilde{\sigma}$, if y is a well-defined point of Y' then it will automatically lie in Y. Thus it suffices to show that the tuple y is a well-defined point of Y', in other words that $([M_0{:}1], \ldots, [M_0{:}1], x)$ lies in the $\widehat{U}_T$-semistable locus for the action of $\widehat{U}_T$ on X' (in the $d = 1$ case it was enough to show that $([M_0{:}1], x)$ lies in the semistable locus for the action of $\mathbb{G}_m$). To do so, we can use the existence of a Hilbert-Mumford criterion for the action of groups of the form $\widehat{U}_T$, as seen in Theorem 1.1.8. That is, the semistable locus $X'^{ss,\widehat{U}_T}_{\min +}$ for the action of $\widehat{U}_T$ on X' is given by

$$X'^{ss,\widehat{U}_T}_{\min +} = \bigcap_{u \in U} u X'^{ss,T}_{\min +}. \tag{2.12}$$

We recall that the notation 'min +' indicates that we have twisted the linearisation by a suitable character and taken a tensor power so that the minimal weight for the action of the grading $\mathbb{G}_m$ lies just to the left of the origin.

To describe the semistable locus for this torus action, we let p_i denote the projection from $W^d \times X$ to the product of the i-th factor of W^d with X, which has an action of $\widehat{U}$. We recall that the semistable locus for the action of the grading $\mathbb{G}_m$ in $\widehat{U}$ on $W \times X$ is given by $(W^0_{\min} \times X^0_{\min}) \setminus Z(W \times X)_{\min}$, using the Hilbert-Mumford criterion. By the combinatorial description of semistability for torus actions (see Section 1.1.1), we obtain that

$$X'^{ss,T}_{\min+} = X'^0_{\min} \setminus \bigcup_{i=1}^{d} p_i^{-1}(Z(W \times X)_{\min}).$$

Thus by (2.12) we obtain that

$$X'^{ss,\widehat{U}_T} = X'^0_{\min} \setminus \bigcup_{i=1}^{d} p_i^{-1}(\mathbb{G}_a Z(W \times X)_{\min}).$$

Since $x_0 \in Z_{\min}$ and $[M_0{:}1] \notin Z(W)_{\min}$, we know that $([M_0{:}1], x_0) \notin \mathbb{G}_a Z(W \times X)_{\min} = \mathbb{G}_a(Z(W)_{\min} \times Z_{\min})$. Therefore the point $([M_0{:}1], \ldots, [M_0{:}1], x_0)$ does indeed lie in $X'^{ss,\widehat{U}_T}_{\min+}$, and so it follows that $y_0 = \widehat{U}_T \cdot ([M_0{:}1], \ldots, [M_0{:}1], x_0) \in Y$ as desired.

Applying the same method from the proof in the $d = 1$ case to each factor, we obtain that $x \in Z_{\min}$ is fixed by $\widehat{\mathbb{G}}_a \times \cdots \times \widehat{\mathbb{G}}_a$ (or equivalently by $\widehat{U}$) if and only if $y := \widehat{U}_T \cdot ([M_0{:}1], \ldots, [M_0{:}1], x)$ is fixed by the product T^d in $\mathrm{GL}(2;\mathbb{C})^d$, where T denotes the maximal torus in $\mathrm{GL}(2;\mathbb{C})$. Thus if we set $y_0 = \widehat{U}_T \cdot ([M_0{:}1], \ldots, [M_0{:}1], x_0)$ for any choice of M_0 in the image of $\widetilde{\sigma}$, then $y_0 \in Y^{T^d}$. Since Y is smooth at y_0 (using the assumption that X is smooth at x_0), it is smooth in a small neighbourhood of y_0, and by the reductivity of T^d we can conclude that Y^{T^d} is smooth in a small neighbourhood of y_0. By considering the smooth neighbourhood in $X^U \cap Z_{\min}$ obtained by pulling back under the geometric quotient map the smooth neighbourhood of y_0 in Y^{T^d}, it follows that $X^U \cap Z_{\min}$ is smooth at x_0. □

Remark 2.2.8 (Smoothness in the Projective Completion algorithm)**.** As depicted in Figure 1.11, the Projective Completion algorithm has three components: the $\widehat{U}$-blow-ups (to go from the red box to the orange box), the partial desingularisation construction (to go from the orange box to the yellow box), and the replacements (to go from a grey rectangle to a red or orange rectangle).

If we start with an irreducible projective variety X which is smooth, the result of this section establishes that the projective completion obtained after the $\widehat{U}$-blow-ups has at worst finite quotient singularities (with the caveat of Remark 2.2.6). As we have seen in Section 2.2.1, the partial desingularisation also preserves smoothness. We note that this result does indeed generalise to smooth Deligne-Mumford stacks as established in [26], so that if we start with a variety with finite quotient singularities (namely the $\widehat{U}$-quotient obtained as a result of applying Theorem 1.2.3 to the linear action on X), the resulting projective completion will still only have finite quotient singularities. Finally, the Replacement algorithm will also preserve smoothness (and more generally it should preserve the property of having only finite quotient singularities, by working with smooth Deligne-Mumford stacks). Indeed, at any stage it involves replacing X by $\overline{Y_\beta^{ss}}$ which is the closure in X of Y_β^{ss} where $S_\beta = GY_\beta^{ss}$ is an unstable stratum for a GIT-instability stratification associated to the linear action of G on X. We know from the results of Section 1.1.3 of Chapter 1 that the resulting stable locus inside $\overline{Y_\beta^{ss}}$ will be contained in $(\overline{Y_\beta^{ss}})^0_{\min} = Y_\beta^{ss}$, which is smooth if X is smooth.

As a result, if we start with an irreducible smooth projective variety X, then the projective completion obtained as the output of the Projective Completion algorithm will be the Non-Reductive GIT quotient of a variety, which will be smooth (or with at worst finite quotient singularities), by a group with internally graded unipotent radical such that semistability coincides with stability both in the classical and non-reductive sense. Consequently, the projective completion can have at worst finite quotient singularities. For this reason, we intend in future work to use the results of Section 2.3 to obtain a formula for the Poincaré series of the projective completion obtained as an output of the Projective Completion algorithm when the initial variety is smooth (see Remark 2.3.4).

2.3 When semistability does not coincide with stability: the ‘desingularised’ GIT quotient

In the previous Section 2.2 we have seen that the blow-up constructions of classical and non-reductive GIT preserve smoothness when the initial variety is smooth (see Remark 2.2.6 for a discussion on the issue of finite quotient singularities). This result is important because it ensures that if we start with a smooth variety, then the GIT quotient in the blow-up has at worst finite quotient singularities, which makes it possible to compute its Poincaré polynomial in terms of cohomological information about the initial variety. The aim of this section is to show how such formulae can be obtained.

In Section 2.3.1 we introduce the existing formula in the case of classical GIT, obtained in [66]. In Section 2.3.2 we establish a formula in the case of externally graded unipotent groups. In Section 2.3.3 we consider the general non-reductive case.

2.3.1 The reductive case

In this section we summarise the results of [66] which provide a formula for computing the Poincaré polynomial of the desingularised GIT quotient in classical GIT, in terms of cohomological information about the initial variety. To this end, let G be a reductive group acting linearly on an irreducible smooth projective variety X, and suppose that $\emptyset \neq X^s \subsetneq X^{ss}$. This ensures that the resulting variety $\widetilde{X}$ has a non-empty stable locus, coinciding with the stable locus.

Strategy. As seen in Section 2.2.1, if X is smooth then the blow-up $\widetilde{X}$ is also smooth. Moreover, since the semistable and stable loci coincide for the induced linear action of G on $\widetilde{X}$, there is an isomorphism

$$H^*(\widetilde{X}/\!/G, \mathbb{Q}) \cong H^*_G(\widetilde{X}^{ss}).$$

Thus to compute the Poincaré series of $\widetilde{X}/\!/G$ it suffices to compute the equivariant Poincaré series of $\widetilde{X}^{ss}$. The strategy is then to relate the Poincaré series $P_t^G(\widetilde{X}^{ss})$ to the Poincaré series $P_t^G(X^{ss})$ by using the explicit description of the blow-ups and the additivity property of Poincaré series. Finally, the formula (2.2) introduced in Section 2.1.1 can be used to inductively compute the Poincaré series $P_t^G(X^{ss})$ from the Poincaré series $P_t^G(X)$. Figure 2.1 illustrates this strategy, which provides a way of computing $P_t(\widetilde{X}/\!/G)$ from $P_t^G(X)$.

Poincaré series after one blow-up. As seen in Section 1.2.1, the first stage of the partial desingularisation construction is to successively blow X^{ss} up along the G-sweep of the locus of semistable points fixed by a connected reductive subgroup of G of maximal dimension amongst those subgroups arising as the stabiliser groups of semistable points in X^{ss}.

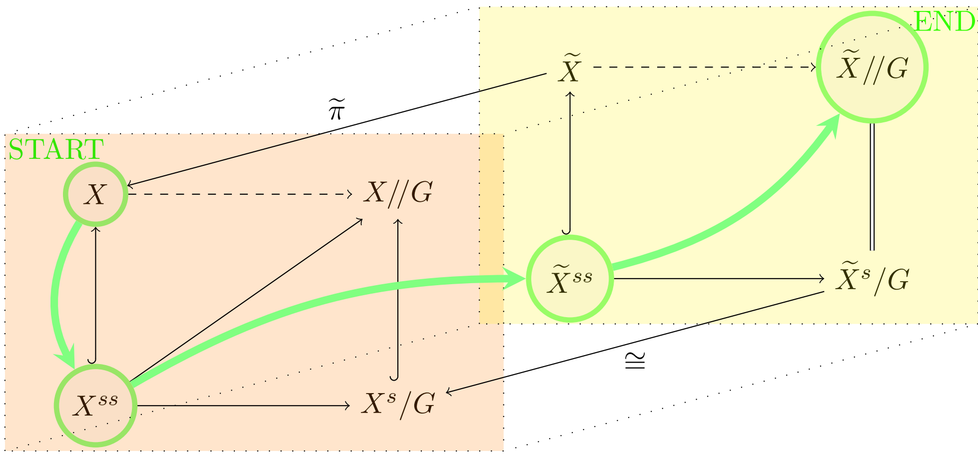

Figure 2.1: Strategy for computing the Poincaré series of the partial desingularisation $\widetilde{X}/\!/G$ in terms of the equivariant Poincaré series of X, for a reductive group G acting linearly on an irreducible smooth projective variety X with non-empty stable locus. Assuming that the G-equivariant Poincaré series of X is known, then the G-equivariant Poincaré series of X^{ss} can be computed inductively using a suitable GIT-instability stratification. In turn, the G-equivariant Poincaré series of $\widetilde{X}^{ss}$ can be computed in terms of that of X^{ss} by explicitly studying the partial desingularisation construction. Finally, the Poincaré series of $\widetilde{X}/\!/G$ coincides with the G-equivariant Poincaré series of $\widetilde{X}^{ss}$ since $\widetilde{X}^{ss} = \widetilde{X}^{s}$ by construction.

Let R be such a subgroup of G. Let $\pi : X_{(1)} \to X^{ss}$ denote the blow-up of X^{ss} along GZ_R^{ss} with exceptional divisor E_R. Consider the linear action of G on $X_{(1)}$ obtained by pulling back the linearisation of the G-action on X and then taking a tensor product with a sufficiently small multiple of the exceptional divisor E_R (which itself has a G-action induced by identifying E_R as the projectivised normal bundle of the centre of the blow-up). For ϵ small enough, the corresponding line bundle is ample.

By the properties of Poincaré series under blow-ups, it follows that

$$P_t^G(X_{(1)}) = P_t(X^{ss}) + P_t(E_R) - P_t(G(X^R \cap X^{ss})). \tag{2.13}$$

To obtain a formula for the equivariant Poincaré series $P_t^G(X_{(1)}^{ss})$, rather than that of $X_{(1)}$, the formula (2.1) which relates the equivariant Poincaré series of the semistable locus to those of the whole variety and of its GIT-unstable strata can be used. In fact, the formula (2.1) can be applied both to the action of G on X and to its induced action on E_R. Through careful analysis and comparison of the GIT-instability stratifications for X and E_R, it can be deduced from (2.13) (see [66, Prop 7.4]) that

$$P_t^G(X_{(1)}^{ss}) = P_t^G(X^{ss}) + P_t^G(E_R^{ss}) - P_t^G(G(X^R \cap X^{ss})). \tag{2.14}$$

The Poincaré series $P_t^G(E_R^{ss})$ can itself be calculated inductively in terms of the Poincaré series $P_t^{N\cap \operatorname{Stab}\beta}(X^R \cap X^{ss})$ (where N is the normaliser of R in G and $\operatorname{Stab}\beta$ is defined as in Section 1.1.1), thus providing an inductive way of computing $P_t^G(X_{(1)}^{ss})$ in terms of cohomological information about X; see [66, Prop 7.4] for the most explicit version of the formula.

Poincaré series of the partial desingularisation. The formula provided by (2.14), which relates the equivariant Poincaré series of the semistable locus after one blow-up to that of the initial variety, can be applied inductively to obtain a formula for the equivariant Poincaré series of the semistable locus of the final variety $\widetilde{X}$ (which coincides with the Poincaré series of $\widetilde{X}//G$) in terms of that of X, as we now explain.

As per Section 2.2.1, let $d_{\max}$ denote the maximal dimension amongst the dimensions of connected reductive subgroups of G arising as the stabiliser groups of points in X^{ss}. Then, let $\mathcal{R}(d_{\max}) = \{R_1, \ldots, R_k\}$ denote a set of conjugacy classes of such reductive subgroups of G. As we have seen in Remark 2.2.1, the result of blowing X^{ss} up successively along $G(X^R \cap X^{ss})$ for $R \in \mathcal{R}(d_{\max})$ is the same as blowing X^{ss} up along $C_{\max}(X^{ss}, G) = \bigsqcup_{R\in\mathcal{R}(d_{\max})} G(X^R \cap X^{ss})$. If we let $Y_{d_{\max}}$ denote the result of this blow-up, then applying (2.14) we obtain:

$$P_t^G(Y_{d_{\max}}^{ss}) = P_t^G(X^{ss}) + \sum_{R\in\mathcal{R}(d_{\max})} \left(P_t^G(E_R^{ss}) - P_t^G(X^R \cap X^{ss})\right). \tag{2.15}$$

To obtain $\widetilde{X}$, we must repeat the above blow-up for each dimension $d_{\max} = d_0 > d_1 > d_2 > d_s = 0$, in decreasing order, arising as the dimension of a connected reductive stabiliser group of a point in the semistable locus. For each d_i, we let $\mathcal{R}(d_i)$ denote a set of representatives of connected reductive subgroups of G arising as the stabiliser group of a semistable point in X. We then inductively define Y_{d_i} to be the result of blowing $Y_{d_{i-1}}$ up successively along $G(Y_{d_{i-1}}^R \cap Y_{d_{i-1}}^{ss})$ for each $R \in \mathcal{R}(d_i)$. By [66, Lem 8.8], for any $R \in \mathcal{R}(d_i)$ the subvariety $G(Y_{d_{i-1}}^R \cap Y_{d_{i-1}}^{ss})$ in $Y_{d_{i-1}}$ is the strict transform of $G(X^R \cap X^{ss})$ for the composition of blow-ups $Y_{d_{i-1}} \to X^{ss}$. To simplify notation, given a connected reductive subgroup R of G arising as the stabiliser group of a point in X^{ss}, we let $G(\widetilde{X^R \cap X^{ss}})$ denote its strict transform in the blow-up Y_{d_i-1} where $d_i = \dim R$. Similarly, we let $\widetilde{E_R^{ss}}$ denote the exceptional divisor for the blow-up of $Y_{d_{i-1}}$ along $G(Y_{d_{i-1}}^R \cap Y_{d_{i-1}}^{ss})$ where $d_i = \dim R$.

By applying (2.15) at each stage, the following formula for the Poincaré series of $\widetilde{X}//G$ is obtained:

$$P_t(\widetilde{X}/\!/G) = P_t^G(X^{ss}) + \sum_{R \in \mathcal{R}} \left(P_t^G(\widetilde{E_R^{ss}}) - P_t^G(G(\widetilde{X^R \cap X^{ss}})) \right) \tag{2.16}$$

where $\mathcal{R}$ is a set of representatives of conjugacy classes of connected reductive subgroups R of G such that there exists $x \in X^R$ such that $G \cdot x$ is closed in X^{ss} and has dimension $\dim G/R$.

2.3.2 The externally graded unipotent case

When semistability does not coincide with stability for the action of $\widehat{U}$ on X, although a projective Non-Reductive GIT quotient for the action of H on X cannot a priori be constructed, by Theorem 1.2.3 a sequence of equivariant blow-ups of X can be performed to obtain a variety $\widehat{X}$ for which semistability does coincide with stability for an induced linear action of $\widehat{U}$ on $\widehat{X}$. This blow-up construction can be viewed as the non-reductive analogue of the partial desingularisation construction of classical GIT. Through this construction, we can obtain a projective Non-Reductive GIT quotient for the action of $\widehat{U}$ on $\widehat{X}$, which can be interpreted as a projective completion of the geometric $\widehat{U}$-quotient of a suitably defined open subset of X (see Section 1.2.2).

As described in the previous Section 2.3.1, in the classical case the partial desingularisation construction can be used to compute the Poincaré series of the desingularised GIT quotient in terms of the G-equivariant Poincaré series of the semistable locus in X; the resulting formula is given in (2.16), and illustrated in Figure 2.1. The aim of this section is to carry out an analogous computation in the non-reductive setting and to obtain a non-reductive analogue of (2.16).

The formula we prove is given in Theorem 2.3.2 below, and illustrated in Figure 2.2. The theorem applies to externally graded unipotent groups $\widehat{U} = U \rtimes \mathbb{G}_m$ satisfying the additional property that U is abelian and that $\mathbb{G}_m$ acts on $\operatorname{Lie} U$ with a single positive weight via the adjoint action. The reason for this assumption is that the blow-up construction of Non-Reductive GIT (Blow-up Construction 2) involves reducing to the case of such groups, by considering a suitable derived series of the unipotent radical and combining the blow-up construction with the procedure of quotienting in stages. Thus it suffices to obtain a formula for such groups. The statement of Theorem 2.3.2 relies on the following notation:

Notation 2.3.1. Suppose that $\widehat{U} = U \rtimes \mathbb{G}_m$ acts linearly on a projective variety X, where U

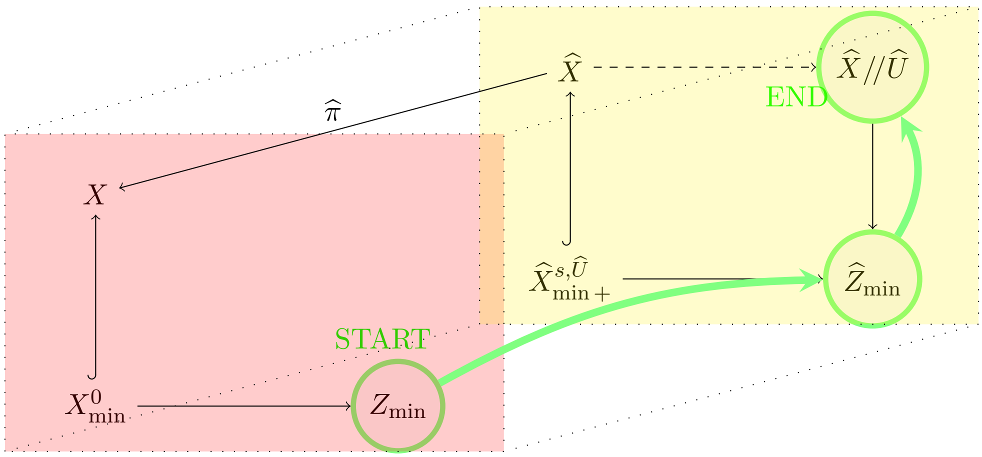

Figure 2.2: Strategy for computing the Poincaré series of $\widehat{X}/\!/G$ in terms of the Poincaré series of $Z_{\min}$, for an externally graded unipotent group $\widehat{U} = U \rtimes \mathbb{G}_m$ acting linearly on an irreducible smooth projective variety X such that there exists a point $z \in Z_{\min}$ with $\mathrm{Stab}_U(z) = \{e\}$. Assuming that the Poincaré series of $Z_{\min}$ is known, then the Poincaré series of $\widehat{Z}_{\min}$ can be computed by explicitly studying the alternative Blow-up Construction 1. The Poincaré series of $\widehat{X}/\!/\widehat{U}$ can then be computed from the Poincaré series of $\widehat{Z}_{\min}$ using (2.5), since by construction ($ss = s \neq \emptyset[\widehat{U}]$) applies for the induced linear action of $\widehat{U}$ on $\widehat{X}$.

is a unipotent algebraic group and $\mathbb{G}_m$ acts on $\operatorname{Lie} U$ with positive weights. We let

$$d_{\min}(X^0_{\min}, \widehat{U}) := d_{(r+1)} < d_{(r)} < \cdots < d_{(1)} < d_{(0)} := d_{\max}(X^0_{\min}, \widehat{U})$$

denote the integers $d \in \mathbb{N}$ arising as the dimensions of stabilisers in $\widehat{U}$ of points in $X^0_{\min}$. In Notation 1.2.6, we defined for any $d \in \mathbb{N}$ the subvariety $C_d(X^0_{\min}, \widehat{U})$ of X consisting of points in $X^0_{\min}$ with d-dimensional stabiliser in $\widehat{U}$. More generally, we can define for any $d \in \mathbb{N}$ the closed subvariety

$$C_{\geq d}(X^0_{\min}, \widehat{U}) := \left\{x \in X^0_{\min} \;\middle|\; \dim \mathrm{Stab}_{\widehat{U}}(x) \geq d\right\}$$

of X. To simplify notation, where it is clear from the context which X and $\widehat{U}$ are being considered, we let $C_{(i)} := C_{\geq d_{(i)}}(X, \widehat{U})$ for $i = 0, \ldots, r+1$ and $C_{\max} := C_{(0)} = C_{d_{\max}(X^0_{\min}, \widehat{U})}(X^0_{\min}, \widehat{U})$.

Theorem 2.3.2 (Poincaré polynomial of the 'partial desingularisation' in the simplest non-reductive case)**.** Let $\widehat{U} = U \rtimes \mathbb{G}_m$ where the unipotent radical U is abelian and $\mathbb{G}_m$ acts with a single positive weight on $\operatorname{Lie} U$ via the adjoint action. Suppose that $\widehat{U}$ acts linearly on an irreducible smooth projective variety X such that there exists a $z \in Z_{\min}$ satisfying $\mathrm{Stab}_U(z) = \{e\}$. Let $\widehat{X}$ denote the variety resulting from applying the alternative Blow-up Construction 1 to the action of $\widehat{U}$ on X. Then we have:

(i) at each stage i of the alternative Blow-up Construction 1, the centre of the blow-up $C_{\max}((X_{(i)})^0_{\min}, \widehat{U})$ and its intersection with $Z(X_{(i)})_{\min}$ are smooth, and correspond to resolutions of singularities of $C_{\geq d_{(i)}}(X, \widehat{U})$ and $C_{\geq d_{(i)}}(X^0_{\min}, \widehat{U}) \cap Z_{\min}$ respectively;

(ii) the Poincaré series of the $\widehat{X}/\!/\widehat{U}$ is given by:

$$P_t(\widehat{X}/\!/\widehat{U}) = \frac{1-t^{2d}}{1-t^2}\left(P_t(Z_{\min}) + \sum_{i=0}^{r} t^2(1-t^{d^{(i)}})P_t\left(C_{\max}((X_{(i)})^0_{\min}, \widehat{U}) \cap Z(X_{(i)})_{\min}\right)\right)$$

where $d := \operatorname{codim}(UZ_{\min}, X)$ and $d^{(i)} := \operatorname{codim}(C_{\geq d_{(i)}}(X^0_{\min}, \widehat{U}) \cap Z_{\min}, Z_{\min})$.

Proof. Since we have assumed that there exists a $z \in Z_{\min}$ such that $\mathrm{Stab}_U(z) = \{e\}$, the alternative Blow-up Construction 1 applies (and we have that $d_{(r+1)} = 0$). We will prove the formula using the same strategy as in the reductive setting, namely by considering a single blow-up first and applying the resulting formula iteratively for each stage of the construction.

The first step of the alternative Blow-up Construction 1 is to blow X up along $C_{\max} = C_{\max}(X^0_{\min}, \widehat{U})$, the locus of points in $X^0_{\min}$ with maximal dimension unipotent stabiliser. We let $X_{(1)}$ denote the blow-up of X along the closure of $C_{\max}$ in X, and let $E_{(1)}$ denote the exceptional divisor. Since $C_{\max}$ is $\widehat{U}$-invariant, there is an induced action of $\widehat{U}$ on $X_{(1)}$. As in the reductive case, we consider the linearisation of this action given by pulling back along the blow-up map the linearisation of the $\widehat{U}$-action on $X^0_{\min}$ and taking a tensor product with a sufficiently small multiple ϵ of the exceptional divisor. By [8, Prop 8.8], the subvariety $Z(X_{(1)})_{\min}$ is the proper transform of $Z_{\min}$ under this blow-up (or equivalently the blow-up of $Z_{\min}$ along $C_{\max} \cap Z_{\min}$).

Since X is smooth, the subvariety $X^0_{\min}$ is also smooth and therefore by Theorem 2.2.4 we have that $C_{\max}$ is smooth. Moreover, by considering the action of $\mathbb{G}_m$ on the closed subvariety $C_{\max}$ of $X^0_{\min}$, by the results of [13] we can conclude that $C_{\max} \cap Z_{\min}$ is also smooth. It follows that both $X_{(1)}$ and $Z(X_{(1)})_{\min}$ are smooth, since they are the blow-ups of smooth varieties along smooth subvarieties.

Using the properties of Poincaré series under blow-ups for smooth subvarieties, we obtain the following formula for the equivariant Poincaré series of $Z(X_{(1)})_{\min} = \mathrm{Bl}_{C_{\max}\cap Z_{\min}} Z_{\min}$:

$$\begin{aligned} P_t(Z(X_{(1)})_{\min}) &= P_t(Z_{\min}) + P_t(E_{(1)}|_{Z(X_{(1)})_{\min}}) - P_t(C_{\max} \cap Z_{\min}) \\ &= P_t(Z_{\min}) + P_t(C_{\max} \cap Z_{\min})\left(t^2 + t^4 + \cdots + t^{2(\operatorname{codim}(C_{\max}\cap Z_{\min}, Z_{\min})-1)}\right) \\ &= P_t(Z_{\min}) + \frac{t^2(1-t^{2d^{(0)}})}{1-t^2} P_t(C_{\max} \cap Z_{\min}). \end{aligned} \tag{2.17}$$

The next step of the alternative Blow-up Construction 1 is to repeat the same procedure but with X replaced by $X_{(1)}$[13]. That is, the smooth variety $(X_{(1)})^0_{\min}$ is blown up along $C_{\max}((X_{(1)})^0_{\min}, \widehat{U})$, which is smooth by Theorem 2.2.4; we let $X_{(2)}$ denote the resulting blow-up, so that $X_{(2)}$ is smooth. Just as in the $i = 0$ case, we also have that $C_{\max}((X_{(1)})^0_{\min}, \widehat{U}) \cap Z(X_{(1)})_{\min}$ is smooth. Thus (2.17) applies with $X_{(1)}$ replaced by $X_{(2)}$ and X by $X_{(1)}$.

By definition we have that $d_{\max}((X_{(1)})^0_{\min}, \widehat{U}) = d_{(1)}$ and so the centre $C_{\max}((X_{(1)})^0_{\min}, \widehat{U})$ of the second blow-up can be identified with the proper transform of $C_{\geq d^{(1)}}(X^0_{\min}, \widehat{U})$ for the first blow-up $X_{(1)} \to X^0_{\min}$, or equivalently with the blow-up of $C_{\geq d^{(1)}}(X^0_{\min}, \widehat{U})$ along $C_{\max}(X^0_{\min}, \widehat{U})$. More generally, we have that for any $i \geq 1$, the subvariety $C_{\geq d_{(i)}}((X_{(1)})^0_{\min}, \widehat{U})$ can be identified with the proper transform of $C_{\geq d_{(i)}}(X^0_{\min}, \widehat{U})$. Since $C_{\max}((X_{(1)})^0_{\min}, \widehat{U})$ is smooth, it follows that this variety can be viewed as a resolution of singularities of the closed subvariety $C_{\geq d_{(1)}}(X^0_{\min}, \widehat{U})$ of $X^0_{\min}$. The same is true for $C_{\max}((X_{(1)})^0_{\min}, \widehat{U}) \cap Z(X_{(1)})_{\min}$, in the sense that $C_{\max}((X_{(1)})^0_{\min}, \widehat{U}) \cap Z(X_{(1)})_{\min}$ is the proper transform of $C_{\geq d^{(1)}}(X^0_{\min}, \widehat{U}) \cap Z_{\min}$ and, since it is smooth, can be viewed as a resolution of singularities of $C_{\geq d^{(1)}}(X^0_{\min}, \widehat{U}) \cap Z_{\min}$. From this description of $C_{\max}((X_{(1)})^0_{\min}, \widehat{U}) \cap Z(X_{(1)})_{\min}$, we have that

$$\operatorname{codim}\left(C_{\max}((X_{(1)})^0_{\min}, \widehat{U}) \cap Z(X_{(1)})_{\min}, Z(X_{(1)})_{\min}\right) = \operatorname{codim}\left(C_{(1)} \cap Z_{\min}, Z_{\min}\right) = d^{(1)}. \tag{2.18}$$

Using (2.18) and applying (2.17), while replacing $X_{(1)}$ by $X_{(2)}$ and X by $X_{(1)}$, yields the following formula for the Poincaré series of $P_t(Z(X_{(2)})_{\min})$:

$$P_t(Z(X_{(2)})_{\min}) = P_t(Z(X_{(1)})_{\min}) + \frac{t^2(1 - t^{2d^{(1)}})}{1 - t^2} P_t\left(C_{\max}((X_{(1)})^0_{\min}, \widehat{U}) \cap Z(X_{(1)})_{\min}\right). \tag{2.19}$$

The above procedure is then repeated for each i until we reach $i = r + 1$. That is, at each stage the smooth variety $X_{(i)}$ is blown up along the smooth subvariety $C_{\max}((X_{(i)})^0_{\min}, \widehat{U})$, which can be identified as the proper transform of $C_{\geq d^{(i)}}((X_{(i-1)})^0_{\min}, \widehat{U})$ for the blow-up $X_{(i)} \to X_{(i-1)}$. In turn, the subvariety $C_{\geq d^{(i)}}((X_{(i-1)})^0_{\min}, \widehat{U})$ of $X_{(i-1)}$ can be identified as the proper

[13] To be exact, the construction requires from the start to blow X up along the closure of $C_{\max}(X^0_{\min}, \widehat{U})$, rather than $X^0_{\min}$ along $C_{\max}(X^0_{\min}, \widehat{U})$, to ensure that a projective variety is obtained at each stage. While $C_{\max}(X^0_{\min}, \widehat{U})$ is smooth by Theorem 2.2.4, its closure may not be. Thus as in the reductive case, we must first resolve the singularities of the closure before blowing up. If we let $\overline{X_{(1)}}$ denote the resulting blow-up, then by choosing ϵ small enough, the open subvariety $(\overline{X_{(1)}})^0_{\min}$ of $\overline{X_{(1)}}$ will be contained in $X_{(1)}$ (see [8, Prop 8.5]). Since it is the open subvariety $(\overline{X_{(1)}})^0_{\min}$ that we are interested in, rather than the whole of $\overline{X_{(1)}}$, to simplify notation we work with $X_{(1)}$ instead of $\overline{X_{(1)}}$. Nevertheless, where the definitions require a projective variety, the variety $X_{(1)}$ should be replaced with $\overline{X_{(1)}}$, and similarly for each $i = 1, \ldots, r + 1$.

transform of $C_{\geq d^{(i)}}((X_{(i-2)})^0_{\min}, \widehat{U})$ for the blow-up $X_{(i-1)} \to X_{(i-2)}$. By iteration we obtain that the subvariety $C_{\max}((X_{(i)})^0_{\min}, \widehat{U})$ of $X_{(i)}$ can be identified with the proper transform of $C_{\geq d_{(i)}}(X^0_{\min}, \widehat{U})$ for the composition of blow-ups $X_{(i)} \to X_{(i-1)} \to \cdots \to X_{(0)} = X$. Similarly, each $Z(X_{(i)})_{\min}$ is the proper transform of $Z(X_{(i-1)})_{\min}$ for the blow-up $X_{(i)} \to X_{(i-1)}$, and iteratively can be viewed as the proper transform of $Z_{\min}$ for the composition of blow-ups $X_{(i)} \to X_{(0)} = X$.

The variety $X_{(r+1)}$ obtained at the $r+1$-th step satisfies the property that

$$d_{\max}((X_{(r+1)})^0_{\min}, \widehat{U}) = d_{(r)} = 0$$

and so $C_{\max}((X_{(r+1)})^0_{\min}, \widehat{U}) = (X_{(r+1)})^0_{\min}$. Thus $X_{(r+1)}$ is the desired smooth variety $\widehat{X}$ which admits a linear $\widehat{U}$-action satisfying ($ss = s \neq \emptyset[\widehat{U}]$)[14].

We can therefore apply the formula (2.5) introduced in Section 2.1.2 to the action of $\widehat{U}$ on $\widehat{X}$ when ($ss = s \neq \emptyset[\widehat{U}]$) is satisfied:

$$P_t(\widehat{X}/\!/\widehat{U}) = P_t(\widehat{Z}_{\min})\frac{1-t^{2\widehat{d}}}{1-t^2}$$

where $\widehat{d} = \operatorname{codim}(U\widehat{Z}_{\min}, \widehat{X})$. Since $\widehat{Z}_{\min}$ is the proper transform of $Z_{\min}$ for the composition of blow-ups $\widehat{X} \to X$, we have that $\operatorname{codim}(U\widehat{Z}_{\min}, \widehat{X}) = \operatorname{codim}(UZ_{\min}, X) = d$ so that $\widehat{d} = d$. Thus

$$P_t(\widehat{X}/\!/\widehat{U}) = P_t(\widehat{Z}_{\min})\frac{1-t^{2d}}{1-t^2}.$$

Finally, the Poincaré series $P_t(\widehat{Z}_{\min})$ can be computed by applying (2.17) iteratively at each stage of the blow-up construction, thus giving the desired formula for $P_t(\widehat{X}/\!/\widehat{U})$:

$$P_t(\widehat{X}/\!/\widehat{U}) = \frac{1-t^{2d}}{1-t^2}\left(P_t(Z_{\min}) + \sum_{i=0}^{r} t^2(1-t^{2d^{(i)}})P_t\left(C_{\max}((X_{(i)})^0_{\min}, \widehat{U}) \cap Z(X_{(i)})_{\min}\right)\right).$$

□

Remark 2.3.3 (A general formula for Blow-up Construction 2)**.** Theorem 2.3.2 is formulated under the assumption that $\widehat{U} = U \rtimes \mathbb{G}_m$ has a unipotent radical U which is abelian, that $\mathbb{G}_m$ acts with a single positive weight on $\operatorname{Lie} U$ via the adjoint action, and also that there exists a point in $Z_{\min}$ with trivial stabiliser group in U. The resulting formula can nevertheless be used to obtain a formula valid in the most general setting, namely when an externally

[14]To be exact, the variety $\widehat{X}$ is in fact $\overline{X_{(r+1)}}$, see footnote 13.

graded unipotent group $\widehat{U}$ acts linearly on an irreducible smooth projective variety. Indeed, in this case Blow-up Construction 2 can be applied, by choosing a filtration of U satisfying the conditions (†) (see Theorem 1.2.3). Applying the formula of Theorem 2.3.2 at each stage of Blow-up Construction 2 (the condition that at each stage there exists a point in $Z_{\min}$ with trivial unipotent stabiliser group can be made to hold thanks to the conditions of (†), see [8, Rk 7.12]), with the caveat of Remark 2.2.6, results in a formula for the Poincaré series of the projective completion $(X \times \mathbb{P}^1 \times \cdots \times \mathbb{P}^1)\widehat{/\!/}(\widehat{U} \times \mathbb{G}_m \times \cdots \times \mathbb{G}_m)$ in terms of cohomological information encoded by $Z_{\min}$, by the varieties corresponding to $Z_{\min}$ for each intermediate quotient and by the centres of the blow-ups required for Blow-up Construction 2.

Remark 2.3.4 (Incorporating Poincaré series into the Projective Completion algorithm)**.** The Projective Completion algorithm introduced in Section 1.3 of Chapter 1 and illustrated in Figure 1.11 consists of a procedure for obtaining, given the linear action of a linear algebraic group H with internally graded unipotent radical on an irreducible projective variety X, an open subset of X admitting a quasi-projective geometric quotient, together with a projective completion of this geometric quotient, obtained itself as a Non-Reductive GIT quotient. The algorithm has two parts: the blow-up constructions (on the right-hand side of Figure 1.11), and the replacement constructions (on the left-hand side of Figure 1.11). The results of Section 2.3.1 above and of the present section can be used to incorporate a calculation of the Poincaré series of the projective completion, obtained as an output of the Projective Completion algorithm, into the part of the algorithm corresponding to the blow-up constructions, provided we assume that the initial variety X is smooth. Thus by describing the behaviour of the Poincaré series of the projective completion under the replacements involved in the algorithm, we expect to be able to generalise the Projective Completion algorithm in the case where X is smooth. That is, in addition to providing an open subset of X admitting a quasi-projective geometric quotient and a projective completion of this geometric quotient, the more general Projective Completion algorithm would provide a way of computing the Poincaré series of the projective completion in terms of the equivariant Poincaré series of X and of various subvarieties of X (more precisely of suitable desingularisations of these subvarieties of X).

2.3.3 The general case

In this section we describe how the results of Sections 2.3.1 and 2.3.2 can be combined to provide a formula for computing the Poincaré series of Non-Reductive GIT quotients in the general non-reductive case, namely for the action of linear algebraic groups with internally graded unipotent radical when semistability does not coincide with stability (in either the classical or non-reductive sense).

A general formula for the Poincaré series of the projective completion. Suppose that a linear algebraic group $H = U \rtimes R$ with internally graded unipotent radical (let $\lambda : \mathbb{G}_m \to Z(R)$ denote the grading one-parameter subgroup and R_λ the quotient $R/\lambda(\mathbb{G}_m)$) acts linearly on an irreducible smooth projective variety X. Suppose moreover that the condition ($ss = s \neq \emptyset[\widehat{U}]$) is not satisfied for the action of $\widehat{U} := U \rtimes \lambda(\mathbb{G}_m) \subseteq H$ on X, and that semistability does not coincide with stability for the action of R_λ on $\widehat{X}/\!/\widehat{U}$; this places us in the most general setting of Non-Reductive GIT (see Theorem 1.2.10). In this case we can obtain a formula for the Poincaré series of $\widetilde{X}/\!/H$ simply by combining the formulae of Section 2.3.1 (when semistability does not coincide with stability for reductive group actions) and Section 2.3.2 (when semistability does not coincide for actions by externally graded unipotent groups). The procedure can be more easily shown through a visual representation, refer to Figure 2.3. Note that Figure 2.3 simply combines Figures 2.1 and 2.2. The two assumptions made to obtain the general formula for the Poincaré series are that there exists a point in $Z_{\min}$ with trivial unipotent stabiliser (so that the results of Section 2.3.2 apply), and that the stable locus for the induced action of R_λ is non-empty (so that the results of Section 2.3.1 apply). Figure 2.3 illustrates a procedure for obtaining the Poincaré series of $\widetilde{X}/\!/H$, starting from the R_λ-equivariant Poincaré series of $Z_{\min}$ and similar equivariant Poincaré series of subvarieties of $Z_{\min}$ which appear in the inductive description of $P_t^{R_\lambda}(Z_{\min}^{ss,R_\lambda})$.

It is important to note that the procedure described in Figure 2.3 requires an understanding of the GIT-instability stratification for the induced action of R_λ on $\widehat{Z}_{\min}$. Unfortunately this may be difficult to achieve in practice, in particular in the context of moduli problems, since the intermediate quotient $\widehat{X}/\!/\widehat{U}$ and the resulting $\widehat{Z}_{\min}$ may not admit an obvious moduli-theoretic interpretation.

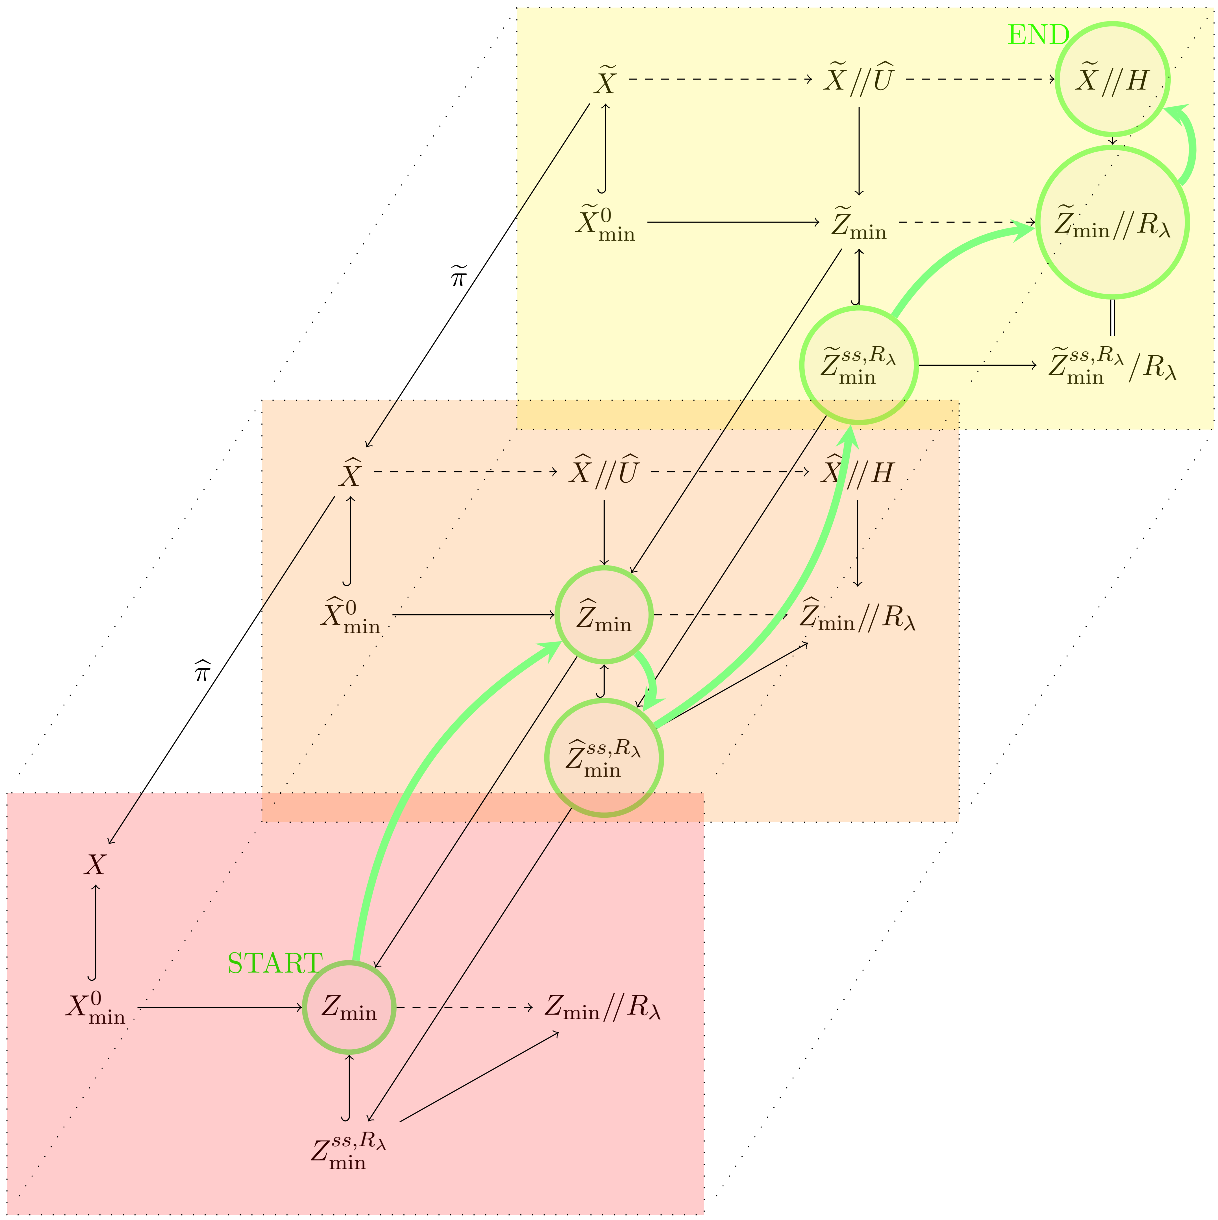

Figure 2.3: Strategy for computing the Poincaré series of $\widetilde{X}/\!/H$, starting from the equivariant Poincaré series of $Z_{\min}$ and equivariant Poincaré series of subvarieties of $Z_{\min}$ which appear in the inductive description of $P_t^{R_\lambda}(Z_{\min}^{ss,R_\lambda})$. We assume that $H = U \rtimes R$ is a linear algebraic group with internally graded unipotent radical U (let $\lambda : \mathbb{G}_m \to Z(R)$ denote the grading one-parameter subgroup, and R_λ the quotient $R/\lambda(\mathbb{G}_m)$) acting linearly on an irreducible smooth projective variety X, such that there exists a point $z \in Z_{\min}$ with $\mathrm{Stab}_U(z) = \{e\}$ and such that the stable locus for the induced action of R_λ on $Z_{\min}/\!/R_\lambda$ is non-empty. Assuming that the R_λ-equivariant Poincaré series of $Z_{\min}$ and of subvarieties of $Z_{\min}$ which appear in the inductive description of $P_t^{R_\lambda}(Z_{\min}^{ss,R_\lambda})$ are known, then the equivariant Poincaré series of $\widehat{Z}_{\min}$ can be computed using (2.5). We can then apply (2.1) to compute the equivariant Poincaré series of $\widehat{Z}_{\min}^{ss,R_\lambda}$. The subvariety $\widetilde{Z}_{\min} := Z(\widetilde{X})_{\min}$ of $\widetilde{X}$ coincides with the variety obtained as a result of applying the partial desingularisation construction to the action of R_λ on $\widehat{Z}_{\min}$. As a result, we can use (2.14) to obtain the equivariant Poincaré series of $\widetilde{Z}_{\min}^{ss,R_\lambda}$, starting from the equivariant Poincaré series of $\widehat{Z}_{\min}^{ss,R_\lambda}$. Since by construction semistability coincides with stability for the action of R_λ on $\widetilde{Z}_{\min}$, we have that $P_t(\widetilde{Z}_{\min}^{ss,R_\lambda}) = P_t(\widetilde{Z}_{\min}/\!/R_\lambda)$. Finally, by (2.7) we obtain a formula for the Poincaré series of $\widetilde{X}/\!/H$ in terms of that of $\widetilde{Z}_{\min}/\!/R_\lambda$.

We therefore conclude this section by describing an alternative strategy for computing $P_t(\widetilde{X}//H)$, which does not rely on the GIT-instability stratification for the induced action of R_λ on $\widehat{Z}_{\min}$.

An alternative strategy. As illustrated by Figure 2.3, the strategy considered above involves computing the equivariant Poincaré series of $\widehat{Z}_{\min}^{ss,R_\lambda}$ from the equivariant Poincaré series of $\widehat{Z}_{\min}$ and of subvarieties appearing in the instability stratification for the induced linear action of R_λ on $\widehat{Z}_{\min}$. Nevertheless, we expect to be able to compute $P_t(\widehat{Z}_{\min}^{ss,R_\lambda})$ in a different way, by exploiting the commutativity of the square

$$\begin{array}{ccc} \widehat{Z}_{\min}^{ss,R_\lambda} & \hookrightarrow & \widehat{Z}_{\min} \\ \downarrow & & \downarrow \\ Z_{\min}^{ss,R_\lambda} & \hookrightarrow & Z_{\min} \end{array}$$

where the vertical arrow on the left is the restriction of the blow-down map $\widehat{Z}_{\min} \to Z_{\min}$. This follows from the fact that points in $\widehat{Z}_{\min}$ which lie over unstable points in $Z_{\min}$ must themselves be unstable, as shown in [88]. The strategy illustrated in Figure 2.3 uses the composition of maps $\widehat{Z}_{\min}^{ss,R_\lambda} \hookrightarrow \widehat{Z}_{\min} \to Z_{\min}$. But using the commutativity of the above diagram, we expect that we could instead use the composition of maps $\widehat{Z}_{\min}^{ss,R_\lambda} \to Z_{\min}^{ss,R_\lambda} \hookrightarrow Z_{\min}$ to compute the equivariant Poincaré series of $\widehat{Z}_{\min}^{ss,R_\lambda}$. That is, instead of requiring the instability stratification for the action of R_λ on $\widehat{Z}_{\min}$, we would only require the instability stratification for the action of R_λ on $Z_{\min}$, which should be easier to describe from a moduli-theoretic perspective.

While obtaining an explicit formula for this alternative strategy in the general case is still work in progress, we can provide a formula under the assumption that the stable locus for the action of R_λ on $Z_{\min}$ is non-empty and coincides with the semistable locus. In this case, we will see that we can obtain a formula for the Poincaré series of $\widehat{X}//\widehat{U}$ by starting with the Poincaré series of $Z_{\min}//R_\lambda$ rather than the equivariant Poincaré series of $Z_{\min}$ (see Corollary 2.3.5 below).

The additional assumption on the action of R_λ on $Z_{\min}$ is helpful for two related reasons. Firstly, it means that the partial desingularisation is not required to obtain a projective completion that is also a geometric quotient, since the assumption ensures that semistability coincides with stability for the induced action of R_λ on $\widehat{X}//\widehat{U}$. Thus we only need consider the quotient $\widehat{X}//H$, rather than $\widetilde{X}//H$ as in the general case of Figure 2.3. Secondly, if all semistable points in $Z_{\min}$ for the action of R_λ are stable, then the stable locus for the action of R_λ on $\widehat{Z}_{\min}$ coincides with the proper transform of the stable locus $Z_{\min}^{s,R_\lambda}$. Indeed, the subvariety $\widehat{Z}_{\min}$ coincides with the proper transform of $Z_{\min}$ in $\widehat{X}$ (thanks to the assumption that there exists a point in $Z_{\min}$ with trivial unipotent stabiliser, see the proof of Theorem 2.3.2) and, by the general arguments of [88] which studies the behaviour of stability under blow-ups, the blow-up procedure producing $\widehat{Z}_{\min}$ from $Z_{\min}$ preserves stability and instability

Thus under this assumption we can obtain a formula for the Poincaré series of $\widehat{X}//H$ in terms of that of $Z_{\min}//R_\lambda$. The strategy for obtaining this formula can be more easily shown through a visual representation, refer to Figure 2.4. Obtaining such a formula is particularly desirable from the point of view of classification problems. Indeed, when applying the theory to specific classification problems, we expect a moduli-theoretic interpretation of $Z_{\min}//R_\lambda$, but not necessarily of $Z_{\min}$. For example, in the Non-Reductive GIT set-up for the moduli problem for Higgs bundles, which we consider in Part II of this thesis, the quotient $Z_{\min}//R_\lambda$ can be interpreted as a product of moduli spaces of semistable Higgs bundles, parametrising Higgs bundles of a fixed Higgs Harder-Narasimhan type which are isomorphic to their Higgs Harder-Narasimhan graded.

We conclude this section by giving an explicit formula for the simplest case. This corresponds to the case where the unipotent radical U is abelian and where $\lambda(\mathbb{G}_m)$ acts with a single positive weight on $\operatorname{Lie} U$ via the adjoint action. Under these assumptions, we obtain from Theorem 2.3.2 and from the discussion above the following formula for the Poincaré series of $\widehat{X}//H$:

Corollary 2.3.5 (Poincaré series of $\widehat{X}//H$)**.** Let $H = U \rtimes R$ be a linear algebraic group with internally graded radical, let $\lambda : \mathbb{G}_m \to Z(R)$ denote the grading one-parameter subgroup and let R_λ denote the quotient $R/\lambda(\mathbb{G}_m)$. Suppose that U is abelian and that $\lambda(\mathbb{G}_m)$ acts with a single positive weight on $\operatorname{Lie} U$ via the adjoint action. Let H act linearly on an irreducible smooth projective variety X such that the following conditions are satisfied:

(i) there exists a point $z \in Z_{\min}$ such that $\operatorname{Stab}_U(z) = \{e\}$;

(ii) the semistable locus $Z_{\min}^{ss,R_\lambda}$ is non-empty and $Z_{\min}^{ss,R_\lambda} = Z_{\min}^{s,R_\lambda}$.

Let $\widehat{X}$ denote the result of applying the alternative Blow-up Construction 1 to the action of

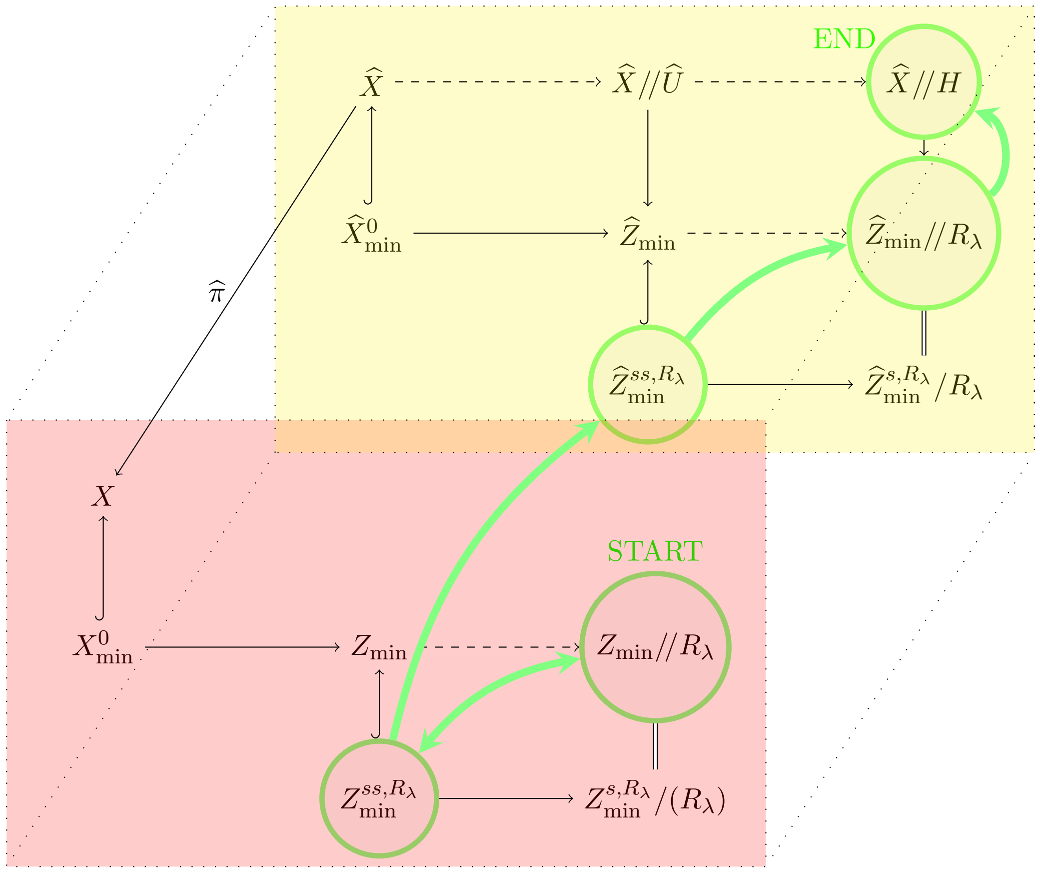

Figure 2.4: Strategy for computing the Poincaré series of $\widehat{X}/\!/H$ in terms of the Poincaré series of $Z_{\min}/\!/R_\lambda$, under the assumption that $Z^{ss,R_\lambda}_{\min} = Z^{s,R_\lambda}_{\min}$. We assume that $H = U \rtimes R$ is a linear algebraic group with internally graded unipotent radical U (let $\lambda : \mathbb{G}_m \to Z(R)$ denote the grading one-parameter subgroup and R_λ the quotient $R/\lambda(\mathbb{G}_m)$) acting linearly on an irreducible smooth projective variety X, such that there exists a point $z \in Z_{\min}$ with $\mathrm{Stab}_U(z) = \{e\}$ and such that $Z^{ss,R_\lambda}_{\min} = Z^{s,R_\lambda}_{\min} \neq \emptyset$ for the induced action of R_λ on $Z_{\min}$. Thanks to this assumption, we have that $P_t(Z_{\min}/\!/R_\lambda) = P_t(Z^{ss,R_\lambda}_{\min})$ and moreover that $\widehat{Z}^{ss,R_\lambda}_{\min}$ corresponds to the proper transform of $Z^{ss,R_\lambda}_{\min}$ under the blow-up $\widehat{X} \to X$. As a result, the formula (2.5) can be applied to obtain the Poincaré series of $\widehat{Z}^{ss,R_\lambda}_{\min}$ in terms of that of $Z^{ss,R_\lambda}_{\min}$. We can then use the fact that $P_t(\widehat{Z}^{ss,R_\lambda}_{\min}) = P_t(\widehat{Z}_{\min}/\!/R_\lambda)$ and apply (2.7) to obtain a formula for the Poincaré series of $\widehat{X}/\!/H$ in terms of that of $\widehat{Z}_{\min}/\!/R_\lambda$.

$\widehat{U} \subseteq H$ on X. Then, using the notation from Theorem 2.3.2, we have:

$$P_t(\widehat{X}/\!/H) = \frac{1-t^{2d}}{1-t^2}\Bigg(P_t\Big(Z_{\min}/\!/R_\lambda\Big) \\ + \sum_{i=0}^{r} t^2(1-t^{2d^{(i)}})P_t\Big((C_{\max}((X_{(i)})^0_{\min},\widehat{U}) \cap Z(X_{(i)})^{ss,R_\lambda}_{\min})/\!/R_\lambda\Big)\Bigg),$$

where

$$d := \operatorname{codim}(UZ_{\min}, X) \text{ and}$$

$$d^{(i)} = \operatorname{codim}\Big((C_{\max}((X_{(i)})^0_{\min},\widehat{U}) \cap Z(X_{(i)})^{ss,R_\lambda}_{\min})/\!/R_\lambda, Z_{\min}/\!/R_\lambda)\Big).$$

Remark 2.3.6 (Application to unstable Higgs bundles of rank 2)**.** We will show in Section 4.3 of Chapter 4 how the formula of Corollary 2.3.5 above can be applied to compute the Poincaré series of a partial compactification of a moduli space involving unstable Higgs bundles of rank 2. In this case, points in the GIT quotient $Z_{\min}/\!/R_\lambda$ can be identified as isomorphism classes of Higgs bundles which are isomorphic to their Higgs Harder-Narasimhan graded (see Definition 3.1.1), and the varieties $(C_{\max}((X_{(i)})^0_{\min},\widehat{U}) \cap Z(X_{(i)})^{ss,R_\lambda}_{\min})/\!/R_\lambda$ can be identified as partial desingularisations of certain rank 1 Brill-Noether loci.

Part II

Applications to Higgs bundles

Overview

The aim of Part I was to describe a systematic way of constructing and studying the geometry of moduli spaces in algebraic geometry, using a combination of classical and Non-Reductive Geometric Invariant Theory (GIT). The aim of Part II is to show how the general results of Part I can be applied in practice, by considering their application to the classification problem for Higgs bundles. For this reason Part II is divided into two chapters which mirror the two chapters of Part 1: the first (Chapter 3) focuses on the construction of moduli spaces for Higgs bundles, while the second (Chapter 4) focuses on studying the geometry of these new moduli spaces.

In Chapter 3 we use Non-Reductive GIT to construct two refined instability stratifications of the stack of (twisted) Higgs bundles. The first is a refinement of the Higgs Harder-Narasimhan stratification, determined by the instability type of the Higgs bundles, the second is a refinement of the Harder-Narasimhan stratification, determined by the instability type of the underlying bundle. This leads to the construction of new quasi-projective moduli spaces for Higgs bundles which are not necessarily semistable. In particular, we obtain a notion of 'stability within instability', analogous to that of stability within semistability (relevant when the rank and degree are coprime). In the rank 2 case, we obtain a complete moduli-theoretic interpretation of this notion, and also of the two refined stratifications in their entirety.

In Chapter 4 we turn to the study of the geometry of these new moduli spaces in the rank 2 case. The new moduli spaces associated to a given Higgs Harder-Narasimhan type μ can be split into two types: the moduli space parametrising 'μ-stable' Higgs bundles of Higgs Harder-Narasimhan type μ, and those parametrising 'μ-unstable' Higgs bundles of Higgs Harder-Narasimhan type μ. Both can be described explicitly when working over $\mathbb{P}^1$. In general, for the second type of moduli spaces, the existence of limits under the Higgs field scaling $\mathbb{C}^*$-action can be used to describe the geometry of the moduli spaces and their Poincaré series. For the first moduli space, which can be considered the more important of the two types since in the context of unstable Higgs bundles it can be viewed as the analogue of the moduli space of stable Higgs bundles, the tools of Chapter 2 can be applied. That is, the moduli space of μ-stable Higgs bundles can be constructed as a Non-Reductive GIT quotient and as such its Poincaré series can be computed using the new cohomological formulae of Chapter 2.

Chapter 3

Stratifications and moduli spaces for the stack of Higgs bundles

3.0 Introduction

The classification problem for Higgs bundles. A classification problem in algebraic geometry has four components: 1. the objects to be classified, 2. equivalence of objects, 3. families of objects, and 4. equivalence of families. For Higgs bundles, these four elements are as follows.

1. Given a compact Riemann surface of genus g and a line bundle $L \to \Sigma$, an L-twisted Higgs bundle of rank r and degree d on Σ is a pair (E, ϕ) where E is a holomorphic vector bundle on Σ of rank r and degree d and $\phi : E \to E \otimes L$ is a holomorphic map. For simplicity we refer to these objects as Higgs bundles, leaving implicit the fixed additional data[1].
2. Two Higgs bundles (E, ϕ) and (E', ϕ') are equivalent if they are isomorphic, that is, if there exists an isomorphism $\psi : E \to E'$ making the following diagram commute:

$$\begin{array}{ccc} E & \xrightarrow{\psi} & E' \\ \downarrow{\scriptstyle \phi} & & \downarrow{\scriptstyle \phi'} \\ E \otimes L & \xrightarrow{\psi \otimes \mathrm{id}_L} & E' \otimes L. \end{array}$$

3. Given a variety S, a family of Higgs bundles parametrised by S is a vector bundle E_S on $\Sigma \times S$ (let $\pi : \Sigma \times S \to \Sigma$ denote the natural projection) together with a holomorphic map $\phi_S : E_S \to E_S \otimes \pi^* L$. Note that any $s \in S$ gives rise to a Higgs bundle (E_s, ϕ_s) on Σ.
4. Two such families (E_S, ϕ_S) and (E'_S, ϕ'_S) are equivalent if, for every $s \in S$, the Higgs bundles

[1]When L is the canonical line bundle of Σ, we recover Higgs bundles in the classical sense as they were first introduced by Hitchin in [51]. Many features of the moduli space of semistable Higgs bundles are preserved when considering the more general case of twisted Higgs bundles (see [31]), which also have links to theoretical physics (see for example [5] and [21]).

(E_s, ϕ_s) and (E'_s, ϕ'_s) are isomorphic[2].

All of these ingredients are encoded in the geometric object that is the moduli stack of Higgs bundles, which we denote by $\mathscr{H}_{r,d}(\Sigma, L)$.

As is common for moduli stacks arising from classification problems in algebraic geometry, the stack of Higgs bundles is an algebraic stack, locally of finite type over $\mathbb{C}$ (see [18, Thm 7.18]), and so the problem of describing its geometry arises. The relevance of this problem is that it can be viewed as a reformulation of the classification problem: to solve a classification problem is to understand the geometry of its associated moduli stack.

Classical GIT for Higgs bundles. The geometry of the stable and semistable substacks of the moduli stack of Higgs bundles is already well understood. Indeed, imposing the condition of (semi)stability for Higgs bundles ensures that quasi-projective moduli spaces for semistable and stable Higgs bundles can be constructed using classical GIT, thanks to the correspondence established by Nitsure in [81] between (semi)stability in the sense of GIT and (semi)stability in the sense of Higgs bundles[3]. We denote the moduli space of semistable Higgs bundles by $\mathcal{M}^{ss}_{r,d}(\Sigma, L)$, which contains as an open subset the moduli space $\mathcal{M}^{s}_{r,d}(\Sigma, L)$ of stable Higgs bundles. When the rank r and degree d are coprime, $\mathcal{M}^{ss}_{r,d}(\Sigma, L)$ is a quasi-projective coarse moduli space for the substack $\mathscr{H}^{ss}_{r,d}(\Sigma, L) \subseteq \mathscr{H}_{r,d}(\Sigma, L)$ of semistable Higgs bundles. If the coprime condition is not satisfied, then restricting to the moduli space of stable Higgs bundles $\mathcal{M}^{s}_{r,d}(\Sigma, L)$ gives a quasi-projective coarse moduli space for the substack of stable Higgs bundles $\mathscr{H}^{s}_{r,d}(\Sigma, L)$, while $\mathcal{M}^{ss}_{r,d}(\Sigma, L)$ is a good moduli space for $\mathscr{H}^{ss}_{r,d}(\Sigma, L)$, in the sense of [2][4]. The moduli space $\mathcal{M}^{ss}_{r,d}(\Sigma, L)$ is widely studied thanks to its rich geometry, and in particular thanks to its structure as a hyperkähler manifold that is a completely integrable Hamiltonian system (see for example [51, 52]).

The question underlying Part II of this thesis is whether there are other 'classes' of Higgs

[2]To identify two families (E_S, ϕ_S) and (E'_S, ϕ'_S) as equivalent if there exists an isomorphism of bundles $E_S \to E'_S$ commuting with the maps ϕ_S and ϕ'_S would perhaps be a more 'obvious' equivalence relation, but it is too rigid and does not lead to the existence of a fine moduli space, even when restricting to stable Higgs bundles.

[3]We note that GIT is not the only way to construct the moduli space of stable Higgs bundles as a complex space. Indeed it was first constructed by Hitchin in the rank 2 case using analytic methods. An analytic construction of the moduli space in the higher rank case appears in the recent paper [27], based on the so-called 'Kuranishi slice method'.

[4]Note that $\mathcal{M}^{ss}_{r,d}(\Sigma, L)$ is a coarse moduli space for the substack $\mathscr{H}^{ps}_{r,d}(\Sigma, L)$ of $\mathscr{H}^{ss}_{r,d}(\Sigma, L)$ parametrising polystable Higgs bundles, see also footnote 7.

bundles (different to the class of stable Higgs bundles) determining substacks of the stack of Higgs bundles which also admit quasi-projective coarse moduli spaces, and if so whether these varieties admit a similarly rich structure to the moduli space of stable Higgs bundles.

Instability stratifications for the stack of Higgs bundles. Instability stratifications represent a useful tool for studying the stack of Higgs bundles beyond the semistable stratum. There are two natural stratifications of the stack of Higgs bundles. The first is the Higgs Harder-Narasimhan stratification, which is determined by the instability type of the Higgs bundle. The second, called the Harder-Narasimhan stratification, is obtained by considering the instability type of the underlying bundle. They are both instances of instability stratifications; an instability stratification is simply a stratification[5] which is indexed by a discrete instability type (for the stratifications mentioned above it is either the Higgs Harder-Narasimhan type of the Higgs bundle or the Harder-Narasimhan type of the underlying vector bundle).

Instability stratifications play an important role in understanding the geometry of moduli stacks. Indeed, it is proposed by Halpern-Leistner in [44] that a solution to a classification problem should be more than just a moduli stack: it should be an algebraic stack equipped with a so-called Θ-stratification: the concept of a Θ-stratification formalises that of an instability stratification, encoding intrinsic and structured notions of filtrations, semistability and instability.

Refined instability stratifications. Another powerful technique for studying the geometry of a stack is to associate to it an algebraic space (or scheme, or variety) which best approximates the stack. The concept of a coarse moduli space makes this association precise [62]. The coarse moduli space of a stack retains the geometry of this stack, and it is easier to study than the stack since it is a space in a more familiar category (for example an algebraic space, a scheme or a variety). Thus by studying the coarse moduli space, it is possible to infer geometric properties about the stack or the classification problem itself. While a moduli stack does not generally admit a coarse moduli space (for example due to the presence of infinite automorphism groups

[5]We use the word ‘stratification’ in the sense of [4]: we define a stratification of a moduli stack to be a decomposition of the stack into a disjoint union of strata, indexed by a partially ordered set, such that the boundary of any stratum lies in the union of strata with strictly higher index.

of objects), we can nevertheless aim to break down the stack in such a way that each ‘piece’ admits a coarse or good moduli space.

The notion of a refined instability stratification combines these two approaches to the study of the geometry of a stack (namely instability stratifications and coarse moduli spaces). As its name suggests, a refined instability stratification is a refinement of an instability stratification, one which satisfies the property that each stratum admits a quasi-projective coarse moduli space with an explicit projective completion.

The aim of this chapter is to construct such refinements of the Higgs Harder-Narasimhan and Harder-Narasimhan stratifications, which we will call refined Higgs Harder-Narasimhan and refined Harder-Narasimhan stratifications. We do so using Non-Reductive GIT, and the results we obtain extend the known result mentioned above regarding the existence of a quasi-projective coarse moduli space for the substack of stable Higgs bundles.

Structure of the chapter. In Section 3.1 we define the Higgs Harder-Narasimhan and Harder-Narasimhan stratifications of the stack of Higgs bundles and compare the two. We also justify the need to refine both these stratifications in order to obtain refined instability stratifications. Sections 3.2 and 3.3 present the main technical content of the chapter. We use Non-Reductive GIT to construct two refined instability stratifications of the stack of Higgs bundles which refine the Higgs Harder-Narasimhan and Harder-Narasimhan stratifications respectively. For the former, we use the spectral correspondence to enable the application of existing results for sheaves regarding the application of Non-Reductive GIT [55, 54, 59]); for the latter we extend Nitsure’s GIT set-up for the construction of the moduli space of semistable Higgs bundles [81]. We then address the problem of the moduli-theoretic interpretation of these stratifications: we provide a complete moduli-theoretic interpretation of the refined Higgs Harder-Narasimhan and Harder-Narasimhan stratifications in the rank 2 case, and obtain a partial moduli-theoretic interpretation of the refined Harder-Narasimhan stratification in the general rank case.

3.1 Stratifications of the stack of Higgs bundles

The aim of this section is to study two instability stratifications of the stack of Higgs bundles: the Higgs Harder-Narasimhan stratification and the Harder-Narasimhan stratification. After introducing the two stratifications and their properties in Section 3.1.1, we describe their intersection in Section 3.1.2 and finally in Section 3.1.3 we justify the need to refine them in order to obtain refined instability stratifications.

3.1.1 Higgs Harder-Narasimhan and Harder-Narasimhan stratifications

In this section we define the Higgs Harder-Narasimhan and the Harder-Narasimhan stratifications of the stack of Higgs bundles and describe some of their properties.

Definition 3.1.1. The *Higgs Harder-Narasimhan filtration* of a Higgs bundle (E, ϕ) is the unique filtration $0 = E^0 \subset E^1 \subset \cdots \subset E^s = E$ satisfying the following conditions[6]:

(i) E^i is a ϕ-invariant subbundle of E for each $i = 1, \dots, s$;

(ii) (E_i, ϕ_i) is semistable for each i, where $E_i = E^i/E^{i-1}$ and ϕ_i is the map $E_i \to E_i \otimes L$ induced by ϕ;

(iii) $\mu(E_1) > \mu(E_2) > \cdots > \mu(E_s)$.

The *Higgs Harder-Narasimhan graded* of (E, ϕ) is the Higgs bundle $\mathrm{gr}(E, \phi) := (E_1, \phi_1) \oplus \cdots \oplus (E_s, \phi_s)$. The vector

$$\mu = \left(\frac{d_1}{r_1}, \dots, \frac{d_1}{r_1}, \frac{d_2}{r_2}, \dots, \frac{d_s}{r_s}\right)$$

where $d_i = \deg(E_i)$, $r_i = \mathrm{rk}(E_i)$ and each entry d_i/r_i appears r_i times is the *Higgs Harder-Narasimhan type* of (E, ϕ), and we call r and d its *rank* and *degree* respectively. The Higgs Harder-Narasimhan type $(d/r, \dots, d/r)$ associated to a semistable Higgs bundles of rank r and degree d is denoted by μ_0, and we call the Higgs Harder-Narasimhan type μ *unstable* if $\mu \neq \mu_0$.

Given a Higgs Harder-Narasimhan type μ of rank r and degree d, let $\mathscr{H}^{\mu}_{r,d}(\Sigma, L) \subseteq \mathscr{H}_{r,d}(\Sigma, L)$ denote the substack of Higgs bundles with Higgs Harder-Narasimhan type μ. The *Higgs Harder-Narasimhan stratification* of $\mathscr{H}_{r,d}(\Sigma, L)$ is given by

$$\mathscr{H}_{r,d}(\Sigma, L) = \mathscr{H}^{ss}_{r,d}(\Sigma, L) \sqcup \bigsqcup_{\mu \neq \mu_0} \mathscr{H}^{\mu}_{r,d}(\Sigma, L). \tag{3.1}$$

[6]See [96, §3] for the existence and uniqueness of this filtration.

Remark 3.1.2. A Higgs Harder-Narasimhan type μ has an associated convex polygon P_μ in the plane, as is the case for Harder-Narasimhan types, determined by connecting the vertices $(r_1, d_1), \ldots, (r_s, d_s)$. This provides a strict partial ordering on the set of Higgs Harder-Narasimhan types, namely $\mu > \mu'$ if P_μ lies strictly above $P_{\mu'}$. The Higgs Harder-Narasimhan stratification is a stratification with respect to this partial ordering, in the sense of [4]. That is, the following properties are satisfied (see [99, §2.1]):

(i) the boundary of each stratum $\mathscr{H}^\mu_{r,d}(\Sigma, L)$ is contained in the union of strata with strictly larger Higgs Harder-Narasimhan type, so that each stratum $\mathscr{H}^\mu_{r,d}(\Sigma, L)$ is closed in

$$\mathscr{H}^{\leq\mu}_{r,d}(\Sigma, L) := \mathscr{H}^\mu_{r,d}(\Sigma, L) \sqcup \bigsqcup_{\mu'<\mu} \mathscr{H}^{\mu'}_{r,d}(\Sigma, L),$$

which is itself open in $\mathscr{H}_{r,d}(\Sigma, L)$;

(ii) the set of Higgs Harder-Narasimhan types μ' of a fixed rank and degree satisfying $\mu' < \mu$ is finite, so that each stratum $\mathscr{H}^\mu_{r,d}(\Sigma, L)$ admits an open neighbourhood

$$\mathscr{H}^\mu_{r,d}(\Sigma, L) \subseteq \mathscr{H}^{\leq\mu}_{r,d}(\Sigma, L)$$

where $\mathscr{H}^{\leq\mu}_{r,d}(\Sigma, L) = \bigsqcup_{\mu'\leq\mu} \mathscr{H}^{\mu'}_{r,d}(\Sigma, L)$ is a finite union.

We note that the structure of $\mathscr{H}^\mu_{r,d}(\Sigma, L)$ as an algebraic stack that is a locally closed substack $\mathscr{H}_{r,d}(\Sigma, L)$ follows from [40, Cor 4.9].

We now define the Harder-Narasimhan stratification.

Definition 3.1.3. Let $\mathscr{V}_{r,d}(\Sigma)$ denote the stack of rank r and degree d vector bundles on Σ, and let $F : \mathscr{H}_{r,d}(\Sigma, L) \to \mathscr{V}_{r,d}(\Sigma)$ denote the forgetful map sending a Higgs bundle to its underlying vector bundle. Consider the stratification

$$\mathscr{V}_{r,d}(\Sigma) = \mathscr{V}^{ss}_{r,d}(\Sigma) \sqcup \bigsqcup_{\tau\neq\tau_0} \mathscr{V}^\tau_{r,d}(\Sigma)$$

of $\mathscr{V}_{r,d}(\Sigma)$ based on Harder-Narasimhan type, where

(i) $\mathscr{V}^{ss}_{r,d}(\Sigma)$ denotes the substack of semistable vector bundles,

(ii) $\mathscr{V}^\tau_{r,d}(\Sigma)$ the substack of unstable vector bundles of Harder-Narasimhan type τ, and

(iii) τ_0 the Harder-Narasimhan type associated to a semistable vector bundle.

For each τ, let $\mathscr{H}_\tau^{r,d}(\Sigma, L) := F^{-1}(\mathscr{V}_{r,d}^\tau(\Sigma))$ and $\mathscr{H}_{ss}^{r,d}(\Sigma, L) = \mathscr{H}_{\tau_0}^{r,d}(\Sigma, L)$. The *Harder-Narasimhan stratification* of $\mathscr{H}_{r,d}(\Sigma, L)$ is given by

$$\mathscr{H}_{r,d}(\Sigma, L) = \mathscr{H}_{ss}^{r,d}(\Sigma, L) \sqcup \bigsqcup_{\tau \neq \tau_0} \mathscr{H}_\tau^{r,d}(\Sigma, L). \tag{3.2}$$

Remark 3.1.4. As with the Higgs Harder-Narasimhan stratification, the Harder-Narasimhan stratification is a stratification in the sense of [4] due to the upper semi-continuity of the Harder-Narasimhan type, proved in [92, Thm 3, Prop 10]. The structure of $\mathscr{V}_{r,d}^\tau(\Sigma)$ as an algebraic stack that is a locally closed substack $\mathscr{V}_{r,d}(\Sigma)$ follows from [82].

Proposition 3.1.5 identifies the fibre of the forgetful map F as a quotient stack, a result which we will use when studying the geometry of moduli spaces of unstable Higgs bundles of rank 2 in Chapter 4.

Proposition 3.1.5 (Fibre of the forgetful map)**.** Given $[E] \in \mathscr{V}_{r,d}(\Sigma)$, there is an isomorphism of stacks

$$F^{-1}([E]) \cong [\operatorname{Hom}(E, E \otimes L)/\operatorname{Aut}(E)]$$

where the action of $\operatorname{Aut}(E)$ on $\operatorname{Hom}(E, E \otimes L)$ is given by $\psi \cdot \phi = (\psi \otimes \mathrm{id}_L) \circ \phi \circ \psi^{-1}$ for $\psi \in \operatorname{Aut}(E)$ and $\phi \in \operatorname{Hom}(E, E \otimes L)$.

Proof. We prove the isomorphism by establishing an equivalence of categories. First, we note that a map $F^{-1}([E]) \to [\operatorname{Hom}(E, E \otimes L)/\operatorname{Aut}(E)]$ is determined by an $\operatorname{Aut}(E)$-torsor over $F^{-1}([E])$ together with an $\operatorname{Aut}(E)$-equivariant map to $\operatorname{Hom}(E, E \otimes L)$. A suitable torsor T can be constructed as follows. The objects of T over a scheme S are pairs $((F_S, \phi_S), \psi_S)$ where (F_S, ϕ_S) is a family of Higgs bundles parametrised by S, and ψ_S is an isomorphism of E_S with F_S. A morphism between such objects is a morphism of the families living over the identity map $E_S \to E_S$.

Then there is a natural surjection $T \to F^{-1}([E])$ mapping an object $((F_S, \phi_S), \psi_S)$ over S to the equivalence class of the family (F_S, ϕ_S) of Higgs bundles parametrised by S. This family lies in $((F_S, \phi_S), \psi_S)$ since by assumption there exists an isomorphism $\psi_S : F_S \to E_S$. Moreover, the group $\operatorname{Aut}(E)$ acts on T via its natural action on the right factor (post-composition), and trivially on the left factor. Note that since this action is free and transitive, in particular it is free and transitive on the fibres of the projection map $T \to F^{-1}([E])$. Thus T is an $\operatorname{Aut}(E)$-torsor over $F^{-1}([E])$. The map taking a pair $((F_S, \phi_S), \psi_S)$ to the element $(\psi_S \otimes \mathrm{id}_{\pi_\Sigma^* L}) \circ \phi_S \circ \psi_S \in \mathrm{Hom}(E_S, E_S \otimes \pi_\Sigma^* L)$, where $\pi_\Sigma : \Sigma \times S \to \Sigma$ is the projection, gives the desired $\mathrm{Aut}(E)$-equivariant map $T \to \mathrm{Hom}(E, E \otimes L)$.

The inverse map $[\mathrm{Hom}(E, E \otimes L)/\,\mathrm{Aut}(E)] \to F^{-1}([E])$ can be defined as follows. An object of $[\mathrm{Hom}(E, E \otimes L)/\,\mathrm{Aut}(E)]$ over S is a principal $\mathrm{Aut}(E)$-bundle $P \to S$ together with an $\mathrm{Aut}(E)$-equivariant map $f : P \to \mathrm{Hom}(E, E \otimes L)$. If $\{U_i, s_i\}$ denotes a trivialisation of P, then since two Higgs fields ϕ and ϕ' for E give rise to isomorphic Higgs bundles if and only if ϕ and ϕ' lie in the same $\mathrm{Aut}(E)$-orbit, it follows that the maps $f \circ s_i$ and $f \circ s_j$ agree on the intersection $U_i \cap U_j$. Thus the maps $f \circ s_i$ can be glued to obtain a morphism $S \to F^{-1}([E])$. This construction also respects morphisms, and so we obtain the desired inverse. □

3.1.2 Comparison of the stratifications and filtrations

In this section we compare the Higgs Harder-Narasimhan and the Harder-Narasimhan stratifications and filtrations. We will repeatedly use the result that if $E \to E'$ is a non-zero homomorphism of semistable vector bundles, then $\mu(E) \leq \mu(E')$.

In general, we have the following result:

Proposition 3.1.6. Let μ be a Higgs Harder-Narasimhan type and τ a Harder-Narasimhan type, both of rank r and degree d. Then we have:

(i) the set T_μ of all possible Harder-Narasimhan types τ for the underlying bundle of a Higgs bundle of type μ is finite;

(ii) the set U_τ of all possible Higgs Harder-Narasimhan types μ for Higgs bundles with an underlying bundle of type τ is finite;

(iii) Given a Higgs bundle (E, ϕ) of Higgs Harder-Narasimhan type τ and with E of Harder-Narasimhan type μ, the equality $\mu = \tau$ holds if and only if the Higgs Harder-Narasimhan and Harder-Narasimhan filtrations of E coincide.

Remark 3.1.7 (When $\deg L = 0$)**.** If $\deg L = 0$, which happens in particular when considering classical Higgs bundles on a curve of genus 1 (since the canonical line bundle then has degree 0), it can be shown by a straightforward generalisation of [29, Prop 4.1] that the Higgs Harder-Narasimhan and Harder-Narasimhan stratifications coincide.

Proof of Proposition 3.1.6. Part (i) is proved in the special case where $\mu = \mu_0$ in [81, Prop. 3.2]. In the proof, it is shown that if (E, ϕ) is a semistable Higgs bundle, and $0 = E^{0'} \subset E^{1'} \subset \cdots \subset E^{t'} = E$ denotes the Harder-Narasimhan filtration of E, then

$$\mu(E^{i'}/E^{i-1'}) \leq \mu(E) + \frac{(r-1)^2}{r} \deg L \tag{3.3}$$

for every $i = 1, \ldots, t'$. Thus we have in particular that $\mu(E^{1'})$ is bounded above. The decreasing condition on the slopes of a Harder-Narasimhan type ensures that bounding the slope of the maximally destabilising subbundle bounds the number of allowable Harder-Narasimhan types of E. Thus there are only a finite number of possible Harder-Narasimhan types (with a fixed rank and degree) for the underlying bundle of a semistable Higgs bundle. The proof generalises to arbitrary μ as follows.

Let (E, ϕ) have Higgs Harder-Narasimhan type $\mu = (d_1/r_1, \ldots, d_s/r_s)$. Let $0 = E^0 \subset E^1 \subset \cdots \subset E^s = E$ denote its Higgs Harder-Narasimhan filtration, and let $0 = E^{0'} \subset E^{1'} \subset \cdots \subset E^{t'} = E$ denote the Harder-Narasimhan filtration of the underlying vector bundle E. It suffices to find an upper bound for the slope of $E^{1'}$, depending only on μ. While each quotient $(E^\gamma/E^{\gamma-1}, \phi_\gamma)$ is semistable as a Higgs bundle, its underlying bundle may have a non-trivial Harder-Narasimhan filtration $0 = E^{\gamma,0'} \subset E^{\gamma,1'} \subset \cdots \subseteq E^{\gamma,s_\gamma'} = E^\gamma/E^{\gamma-1}$. Let $E^{\gamma,j}$ denote the preimage of $E^{\gamma,j'}$ under the quotient map $E^\gamma \to E^\gamma/E^{\gamma-1}$. Note that $E^{\gamma,j'} = E^{\gamma,j}/E^{\gamma-1}$ and that $E^\gamma = E^{\gamma+1,0} = E^{\gamma,s_\gamma}$. The Higgs Harder-Narasimhan filtration of E can be refined to include all of the subbundles $E^{\gamma,j}$: $0 = E^0 \subset E^{0,1} \subset \cdots \subset E^{0,s_0} \subset E^{1,0} \subset E^{1,1} \subset \cdots \subset E^{s,s_s} = E$. Although the slopes of these subbundles may no longer satisfy the decreasing condition, each successive quotient $E^{\gamma,j}/E^{\gamma,j-1} \cong \left(E^{\gamma,j}/E^{\gamma-1}\right) / \left(E^{\gamma,j-1}/E^{\gamma-1}\right) = E^{\gamma,j'}/E^{\gamma,j-1'}$ is semistable.

Suppose that $E^{1'} \subseteq E^{k,l}$ but $E^{1'} \nsubseteq E^{k,l-1}$. Then the restriction of $\phi : E \to E \otimes L$ to $E^{1'}$ induces a non-zero map $E^{1'} \to E^{k,l}/E^{k,l-1} \otimes L$. Since both $E^{1'}$ and $E^{k,l}/E^{k,l-1} \otimes L$ are semistable vector bundles, it follows that $\mu(E^{1'}) \leq \mu\left(E^{k,l}/E^{k,l-1} \otimes L\right) = \mu\left(E^{k,l}/E^{k,l-1}\right) + \deg L$. By applying (3.3) in this setting, we have the inequality

$$\mu\left(E^{k,l'}/E^{k,l-1'}\right) \leq \mu\left(E^k/E^{k-1}\right) + \frac{(r_k-1)^2}{r_k} \deg L.$$

Since $E^{k,l}/E^{k,l-1} \cong E^{k,l'}/E^{k,l-1'}$, the same inequality holds for $E^{k,l}/E^{k,l-1}$. Thus we have:

$$\mu(E^{1'}) \leq \mu\left(E^{k,l}/E^{k,l-1}\right) + \deg L \leq \frac{d_k}{r_k} + \left(\frac{(r_k-1)^2}{r_k} + 1\right) \deg L.$$

We let $r_{\max}$ denote the maximum of the r_γ for $\gamma = 1, \ldots, s$. Then since d_1/r_1 is maximal among all of the d_γ/r_γ for $\gamma = 1, \ldots, s$, we have:

$$\mu(E^{1\prime}) \leq \frac{d_1}{r_1} + \left(\frac{(r_{\max}-1)^2}{r_{\max}} + 1\right) \deg L.$$

The term on the right only depends on μ, and thus can be considered a fixed bound.

To prove part (ii), it suffices to observe that if a Higgs bundle has an underlying bundle of Harder-Narasimhan type $\tau = \left(\frac{d_1'}{r_1'}, \ldots, \frac{d_1'}{r_1'}, \frac{d_2'}{r_2'}, \ldots, \frac{d_t'}{r_t'}\right)$, then its Higgs Harder-Narasimhan type $\mu = \left(\frac{d_1}{r_1}, \ldots, \frac{d_s}{r_s}\right)$ satisfies $\frac{d_1}{r_1} \leq \frac{d_1'}{r_1'}$. There are only finitely many such Higgs Harder-Narasimhan types μ for a fixed τ.

To prove part (iii), we note that if both types coincide, then the successive quotients associated to both filtrations have the same rank. It follows then from the uniqueness of maximally destabilising subbundles that the filtrations coincide. □

In the rank 2 case, we can describe the relationship between Higgs Harder-Narasimhan and Harder-Narasimhan types and filtrations even more explicitly, as per Proposition 3.1.8 below which we will use in Chapter 4 for our description of moduli spaces of unstable Higgs bundles of rank 2.

Proposition 3.1.8 (When $r = 2$)**.** Let (E, ϕ) be a rank 2 Higgs bundle. Then we have:

(i) if (E, ϕ) is unstable, the Higgs Harder-Narasimhan filtration of E coincides with its Harder-Narasimhan filtration, so that the Higgs Harder-Narasimhan type of (E, ϕ) coincides with the Harder-Narasimhan type of E;

(ii) if E is unstable with Harder-Narasimhan type $\tau = (d_1, d_2)$ such that $d_1 - d_2 > \deg L$ (or more generally if $H^0(\Sigma, E^{1\vee} \otimes E/E^1 \otimes L) = \{0\}$ where $0 \subseteq E^1 \subseteq E$ denotes the Harder-Narasimhan filtration of E), then (E, ϕ) is also unstable and has Higgs Harder-Narasimhan type $\mu = \tau$.

Figure 3.1 illustrates the above Proposition 3.1.8.

Proof. We prove part (i) first, and thus suppose that (E, ϕ) is unstable of Higgs Harder-Narasimhan type $\mu = (d_1, d_2)$. We let $0 = E^0 \subseteq E^1 \subseteq E^2 = E$ denote its Higgs Harder-Narasimhan filtration. Since (E, ϕ) is unstable, the underlying bundle E must also be unstable.

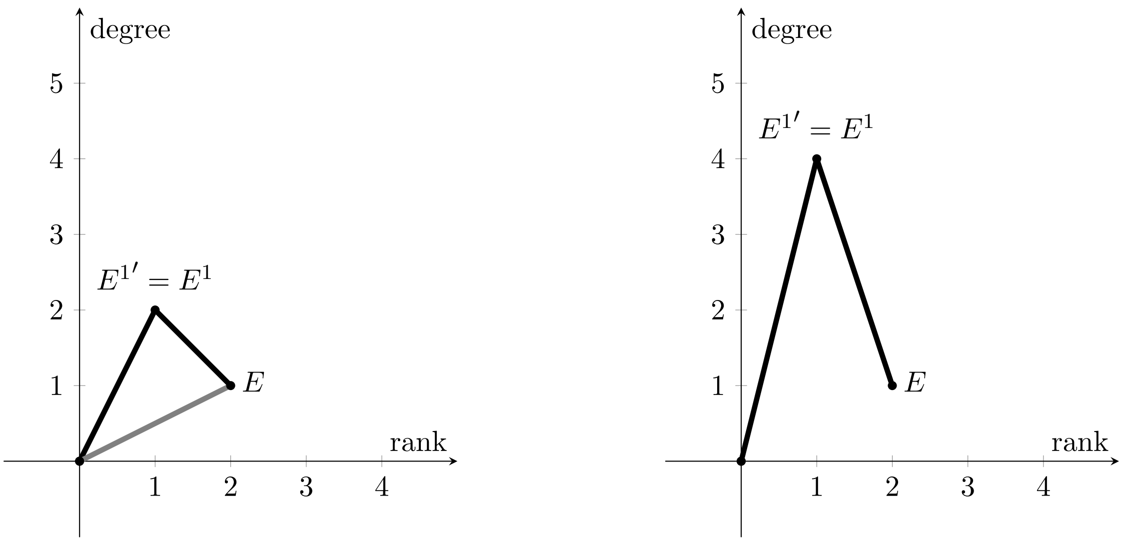

(a) If $d_1 - d_2 \leq \deg L$ (and $H^0(\Sigma, (E^{1'})^\vee \otimes E/E^{1'} \otimes L) \neq \{0\}$), then (E, ϕ) is either semistable (corresponding to the grey polygon), if the map $\phi_{21} : E_1' \to E_2' \otimes L$ induced by the Higgs field is non-zero, or unstable of Higgs Harder-Narasimhan type $\mu = \tau$ (corresponding to the black polygon) if $\phi_{2,1}$ is the zero map. In the latter case, the Higgs Harder-Narasimhan filtration of E coincides with its Harder-Narasimhan filtration.

(b) If $d_1 - d_2 \leq \deg L$, then (E, ϕ) is unstable of Higgs Harder-Narasimhan type $\mu = \tau$ (corresponding to the black polygon), and the Higgs Harder-Narasimhan filtration of E coincides with its Harder-Narasimhan filtration.

Figure 3.1: Relationship between the Higgs Harder-Narasimhan and Harder-Narasimhan filtrations for a rank 2 Higgs bundle (E, ϕ) with unstable underlying bundle. We let $\tau = (d_1, d_2)$ denote the Harder-Narasimhan type of E, and $0 = E^{0'} \subseteq E^{1'} \subseteq E^{2'} = E$ its Harder-Narasimhan filtration, while $0 = E^0 \subseteq E^1 \subseteq E^2 = E$ denotes its Higgs Harder-Narasimhan filtration. For the purpose of the illustration we assume that $d = 1$ and that $\deg L = 5$.

Let $0 = E^{0'} \subseteq E^{1'} \subseteq E^{2'} = E$ denote its Harder-Narasimhan filtration, and $\tau = (d_1', d_2')$ its Harder-Narasimhan type. We must have that $d_1 \leq d_1'$ since otherwise E^1 would be a destabilising subbundle of E of larger degree than $E^{1'}$, contradicting the Harder-Narasimhan type of E.

The composition of the inclusion of $E^{1'}$ into E with the quotient map $E \to E/E^1$ produces a map $E^{1'} \to E/E^1$. Since both $E^{1'}$ and E^1 are line bundles, they are stable. If the map is non-zero, then we must have that $d_1' \leq d_2$, which is not possible since $d_2 < d_1 \leq d_1'$. Thus the map is zero and so $E^{1'} \subseteq E^1$. Therefore both line bundles coincide. It follows that the two filtrations coincide, and in particular that $\mu = \tau$.

For part (ii), suppose that (E, ϕ) is such that E is unstable of Harder-Narasimhan type $\tau = (d_1, d_2)$. Then (E, ϕ) is semistable if and only if the map $\phi_{2,1} : E^1 \to E/E^1 \otimes L$ induced by the Higgs field is non-zero. This map can be identified as a section in $H^0(\Sigma, E_1^\vee \otimes E/E^1 \otimes L)$, and thus if this space of global sections is zero, there can be no semistable Higgs bundle with E as its underlying bundle. Thus (E, ϕ) is unstable and by (i) must have Higgs Harder-Narasimhan type τ. Finally, we note that since $E_1^\vee \otimes E/E^1 \otimes L$ has degree $d_2 - d_1 + \deg L$, if $d_1 - d_2 > \deg L$ then the line bundle has negative degree and so $H^0(\Sigma, E_1^\vee \otimes E/E^1 \otimes L) = \{0\}$. □

3.1.3 Refined instability stratifications

In this section we introduce refined instability stratifications and explain why the instability stratifications given by the Higgs Harder-Narasimhan and Harder-Narasimhan stratifications are not refined instability stratifications; the aim of Sections 3.2 and 3.3 will be to refine them in order to obtain refined instability stratifications.

Remark 3.1.9 (Link with Θ-stratifications)**.** As mentioned in the introduction to this chapter, the concept of instability stratifications for moduli stacks has recently been formalised by Halpern-Leistner (see [44]), and the results we obtain in this chapter could be formulated in this language. Although we intend to do so in future work, for the purpose of this thesis the theory of Θ-stratification is not needed. We note however that while the Higgs Harder-Narasimhan stratification of the stack of Higgs bundles is a Θ-stratification, the Harder-Narasimhan stratification is not. Rather it induces a Θ-stratification of an auxiliary stack: the stack-theoretic quotient of the stack of Higgs bundles by the Higgs field scaling $\mathbb{C}^*$-action, considered in [43].

Definition 3.1.10 (Refined instability stratifications)**.** A refinement of an instability stratification of a moduli stack is a *refined instability stratification* if it is itself a stratification in the sense of [4] (see note 5), and moreover if each of its strata admits a quasi-projective coarse moduli space with an explicit projective completion. The strata of a refined instability stratification are called the *refined strata*.

Both the Higgs Harder-Narasimhan and Harder-Narasimhan stratifications are instability stratifications, but they are not refined instability stratifications, as we now show.

Why the Higgs Harder-Narasimhan stratification is not a refined instability stratification. We consider the open Higgs Harder-Narasimhan stratum first, namely the semistable stratum, since it represents a familiar setting in which we can illustrate the need for refinements in order to obtain quasi-projective coarse moduli spaces. To do so we must assume that the rank and the degree are not coprime. Indeed, under this assumption, the moduli space $\mathcal{M}^{ss}_{r,d}(\Sigma, L)$ is not a coarse moduli space for $\mathscr{H}^{ss}_{r,d}(\Sigma, L)$ in that it does not parametrise isomorphism classes of semistable Higg bundles. Rather it parametrises S-equivalence classes, where two semistable Higgs bundles are S-equivalence if their associated Jordan-Hölder graded are isomorphic. In fact, a coarse moduli space for $\mathscr{H}^{ss}_{r,d}(\Sigma, L)$ does not exist: since any semistable Higgs bundle deforms to its Jordan-Hölder graded (this is the 'jump phenomenon', see [81, Rk 4.6]), any quasi-projective coarse moduli space of semistable Higgs bundles must identify a Higgs bundle with its Jordan-Hölder graded (so that the moduli space is separated), to which it need not be isomorphic.

Nevertheless, by restricting to the substack $\mathscr{H}^{s}_{r,d}(\Sigma, L)$ of stable Higgs bundles we obtain a quasi-projective coarse moduli space in the form of $\mathcal{M}^{s}_{r,d}(\Sigma, L)$[7]. And indeed, the stable stratum $\mathscr{H}^{s}_{r,d}(\Sigma, L)$ is the starting point for obtaining a refined Higgs Harder-Narasimhan stratification, in the sense that it will correspond to the open refined stratum (see Theorem 3.2.2).

The situation for the unstable Higgs Harder-Narasimhan strata is analogous to that of the semistable stratum in the non-coprime case, in so far as a quasi-projective coarse moduli space for an entire Higgs Harder-Narasimhan stratum cannot exist; it must also be further stratified. Indeed, as in the above case of S-equivalence, given an unstable Higgs bundle (E, ϕ), it follows from [81, Prop 4.5] that one can find a family of Higgs bundles parametrised by $\mathbb{A}^1$ such that for each non-zero $t \in \mathbb{A}^1$, the fibre over t is isomorphic to (E, ϕ), while the fibre over 0 is isomorphic to $\mathrm{gr}(E, \phi)$. This 'jump phenomenon' implies that any coarse moduli space for a given stratum would have to identify a Higgs bundle with its associated graded, to which it need not be isomorphic. This precludes the possibility of obtaining a quasi-projective (and hence separated) coarse moduli space for a whole Higgs Harder-Narasimhan stratum (assuming that not all Higgs bundles in the stratum are isomorphic to their graded). We note that the question of whether a notion of S-equivalence exists for unstable Higgs bundles, just as it does for semistable Higgs bundles, remains open.

[7] The moduli space $\mathcal{M}^{ss}_{r,d}(\Sigma, L)$ can itself be viewed as a coarse moduli space provided we restrict to the substack $\mathscr{H}^{ps}_{r,d}(\Sigma, L)$ of polystable Higgs bundles, i.e. semistable Higgs bundles which are isomorphic to their Jordan-Hölder graded. Indeed, restricting to this subclass of semistable Higgs bundles avoids the 'jump phenomenon' mentioned above.

Why the Harder-Narasimhan stratification is not a refined instability stratification. We now illustrate why the Harder-Narasimhan stratification must also be refined to obtain a refined Θ-stratification, by using the Higgs field scaling $\mathbb{C}^*$-action and specialising to the rank 2 case.

An important feature of the moduli space $\mathcal{M}^{ss}_{r,d}(\Sigma, L)$ is the $\mathbb{C}^*$-action it admits, given by scaling the Higgs field: $t \cdot (E, \phi) = (E, t\phi)$ for $t \in \mathbb{C}^*$. This action extends naturally to the whole stack $\mathscr{H}_{r,d}(\Sigma, L)$ of Higgs bundles. It is a non-trivial fact that limits at 0 under the $\mathbb{C}^*$-action always exist inside the semistable stratum $\mathscr{H}^{ss}_{r,d}(\Sigma, L)$ (see [95, Cor 10.3]). Since this stratum admits a quasi-projective coarse moduli space (or a quasi-projective good moduli space if the rank and degree are not coprime), it follows that such limits are always unique. This is no longer true for the whole stack, due to its non-separatedness. Indeed, given $[(E, \phi)] \in \mathscr{H}_{r,d}(\Sigma, L)$, the Higgs bundle $[(E, 0)] \in \mathscr{H}_{r,d}(\Sigma, L)$ is a well-defined limit of $t \cdot [(E, \phi)]$ as $t \to 0$, but this may not be the only limit. A simple example arises when choosing a semistable Higgs bundle (E, ϕ) with unstable underlying bundle. Then there exists a semistable Higgs bundle $[(E_0, \phi_0)] \in \mathscr{H}^{ss}_{r,d}(\Sigma, L)$ such that $t \cdot [(E, \phi)] \to [(E_0, \phi_0)]$ as $t \to 0$. Since E is unstable, we have that $[(E_0, \phi_0)] \neq [(E, 0)]$. Thus $t \cdot [(E, \phi)]$ has at least two distinct limits in $\mathscr{H}_{r,d}(\Sigma, L)$ as $t \to 0$.

In the rank 2 case, both these limits lie in the same Harder-Narasimhan stratum. Indeed, if E has Harder-Narasimhan type $\tau \neq \tau_0$ with associated Harder-Narasimhan graded $\operatorname{gr} E = E_1 \oplus E_2$ and (E, ϕ) is a semistable Higgs bundle, then by [35, Prop 3.1] we have that

$$[(E_0, \phi_0)] = \left[\left(\operatorname{gr} E, \begin{pmatrix} 0 & 0 \\ \phi_{2,1} & 0 \end{pmatrix} \right) \right] \in \mathscr{H}^{2,d}_{\tau}(\Sigma, L) \cap \mathscr{H}^{ss}_{2,d}(\Sigma, L),$$

where $\phi_{2,1}$ is given by the composition $\pi_{E_1} \circ \phi|_{E_1} : E_1 \to E_2$ with $\pi_{E_1} : E \to E/E_1 \cong E_2$ the quotient map (note that $\phi_{2,1}$ is non-zero since (E, ϕ) is semistable by assumption). We also have that $[(E, 0)] \in \mathscr{H}^{2,d}_{\tau}(\Sigma, L)$. Thus any quasi-projective coarse moduli space for $\mathscr{H}^{2,d}_{\tau}(\Sigma, L)$ would have to identify the non-isomorphic Higgs bundles $[(E, 0)]$ and $[(E_0, \phi_0)]$ (the former has Higgs Harder-Narasimhan type τ, the latter is semistable).

We will see in Section 3.3 that to obtain a refined Harder-Narasimhan stratification in general, we should first refine each Harder-Narasimhan stratum based, roughly speaking, on how much the Higgs Harder-Narasimhan filtration of the Higgs bundle differs from the Harder-Narasimhan filtration of its underlying bundle (this is determined by the Higgs field). In the rank 2 case, by Proposition 3.1.8 there are only two possibilities: the Higgs Harder-Narasimhan filtration of a Higgs bundle in a given Harder-Narasimhan stratum either coincides with the filtration of the underlying bundle, or it is trivial (see Remark 3.3.12). In higher rank, there are more possibilities. We will make this refinement precise in Section 3.3.3 (see Definition 3.3.10), where we will define the so-called 'Higgs stratification' of the stack of Higgs bundles.

3.2 Refined Higgs Harder-Narasimhan stratification

As seen in Section 3.1.3, in general a quasi-projective coarse moduli space cannot exist for an entire Higgs Harder-Narasimhan stratum; further stratification is needed to obtain a refined instability stratification. We construct such a stratification in this section using Non-Reductive GIT. Section 3.2.1 shows that each Higgs Harder-Narasimhan stratum can be identified as a quotient stack (see Proposition 3.2.1), a necessary step for the application of GIT (either classical or Non-Reductive) to the construction of moduli spaces. In Section 3.2.2 we use this identification to construct a refinement of the Higgs Harder-Narasimhan stratification which is a refined instability stratification (see Theorem 3.2.2). Finally in Section 3.2.3 we provide a complete moduli-theoretic interpretation of this refined stratification in the case of rank 2 Higgs bundles (see Theorem 3.2.9).

3.2.1 Identification of the strata as quotient stacks

The difficulty in applying GIT (either classical or Non-Reductive) to classification problems lies in achieving a set-up for the classification problems in which the theory can be applied. That is, it is necessary to identify a parameter space for the objects to be classified (this is typically achieved by equipping the objects with some additional structure) with an action of a group such that the orbits coincide with isomorphism classes of the objects to be classified. This procedure is most easily formulated in the language of stacks: given a moduli stack, the application of GIT (either classical or Non-Reductive GIT) requires the identification of the moduli stack with a quotient stack.

Since our goal in this Section 3.2 is to apply Non-Reductive GIT to construct moduli spaces for Higgs bundles of a fixed Higgs Harder-Narasimhan type, the first step is therefore to identify a Higgs Harder-Narasimhan stratum as a quotient stack. We do so in this section (see

Proposition 3.2.1 below).

One approach to identifying a Higgs Harder-Narasimhan stratum as a quotient stack is to generalise the set-up from [81], in which Nitsure constructs a moduli space for semistable Higgs bundles using classical GIT. To use this general set-up to construct moduli spaces using Non-Reductive GIT requires showing that, just as semistability for Higgs bundles can be made to coincide with GIT-semistability (see [81, Prop 5.7]), instability for Higgs bundles can be made to coincide with GIT-instability (as introduced in Section 1.1.1). This strategy is indeed the one employed in [55, 54] in the context of coherent sheaves on projective varieties, where instability for sheaves is related to GIT-instability for the linear action of a reductive group on a projective variety, by using the set-up from Simpson's construction of a moduli space for semistable sheaves [96].

Nevertheless, extending this strategy, even just to the case of Higgs bundles on a smooth projective curve, is lengthy and technical (see [45, Prop 3.20]). An alternative approach[8], and the one which we have chosen to take in this section, involves using the spectral correspondence so that the results for sheaves referred to above can be applied directly, thus bypassing the need to generalise them to Higgs bundles. Indeed, the spectral correspondence establishes an equivalence of categories between L-twisted Higgs bundles on a curve and a subcategory of coherent sheaves on the total space of the line bundle $L \to \Sigma$. Thus results for sheaves can be applied to Higgs bundles by using this identification. We describe the correspondence in more detail below.

The spectral correspondence. The key to applying the results for sheaves to the case of Higgs bundles is the spectral correspondence, described by Simpson in [94]. We summarise the correspondence below in the special case of L-twisted[9] Higgs bundles on a smooth projective curve (rather than in the more general case of Higgs sheaves on an arbitrary dimension smooth projective variety as considered in [94]).

The correspondence relies on the following set-up. We fix an ample line bundle $\mathcal{O}_\Sigma(1)$ on

[8] I am grateful to Yuuji Tanaka for pointing out the spectral correspondence to me, and for suggesting that it might be used to apply results about sheaves to Higgs bundles and vice-versa.

[9] The correspondence is stated in [94] for Higgs sheaves over a projective scheme with Higgs field valued in the cotangent bundle, but is valid for Higgs fields valued in any vector bundle over the projective scheme (and in the case of curves which we are considering, for any line bundle on the curve [84]) [61, §2.2].

Σ and let $\pi : \operatorname{Tot} L \to \Sigma$ denote the projection map from the total space $\operatorname{Tot} L$ of the line bundle L to Σ. We let Y denote the projective completion of $\operatorname{Tot} L$ given by the projective bundle of L, and let $D = Y \setminus \operatorname{Tot} L$ denote the divisor at infinity. Then π extends to Y, and for simplicity we also denote the resulting map by $\pi : Y \to \Sigma$. Finally, we choose an integer $k \gg 1$ such that the line bundle $\mathcal{O}_Y(1) := \pi^*\mathcal{O}_\Sigma(k) \otimes_{\mathcal{O}_Y} \mathcal{O}_Y(D)$ is ample on Y. We then have that $\mathcal{O}_{\operatorname{Tot} L}(1) := \mathcal{O}_Y(1)|_{\operatorname{Tot} L} = \pi^*\mathcal{O}_\Sigma(k)$.

The spectral correspondence establishes an equivalence of categories between:

(i) the category of L-twisted Higgs bundles on Σ;

(ii) the category of coherent sheaves on Y of pure dimension one and with support not intersecting D.

The correspondence is proved in [94, Lem 6.8] and can be summarised as follows. The map $\pi : \operatorname{Tot} L \to \Sigma$ is an affine morphism, and so $\operatorname{Tot} L$ can be identified as the sheafified spectrum of the sheaf of rings $\pi_*\mathcal{O}_{\operatorname{Tot} L}$ on Σ. Moreover, there is a natural identification of this sheaf with the symmetric algebra $\operatorname{Sym}^\bullet(L^\vee)$ on the line bundle L. As a result, the push-forward π_* maps a quasi-coherent sheaf $\mathcal{E}$ on $\operatorname{Tot} L$ to a quasi-coherent sheaf $E = \pi_*\mathcal{E}$ on Σ with an action of $\operatorname{Sym}^\bullet(L^\vee)$. This action is equivalent to the data of a Higgs field $\phi : E \to E \otimes L$. Conversely, a quasi-coherent Higgs sheaf (E, ϕ) on Σ can be interpreted as a quasi-coherent sheaf E with an action of $L^\vee$, which induces an action of $\operatorname{Sym}^\bullet(L^\vee)$. It can then be identified under π as a quasi-coherent sheaf $\mathcal{E}$ on $\operatorname{Tot} L$.

The condition that $E = \pi_*\mathcal{E}$ is a coherent sheaf is equivalent to the condition that $\mathcal{E}$ is a coherent sheaf and that the closure of the support of $\mathcal{E}$ in Y does not intersect D. Finally, the condition that E is torsion-free is equivalent to the condition that $\mathcal{E}$ has pure dimension given by $\dim \Sigma = 1$.

Under the spectral correspondence described above, we have the following further correspondences:

(i) for any coherent sheaf $\mathcal{E}$ on Y with support not meeting D and of pure dimension one, the Hilbert polynomials of $\mathcal{E}$ and of $E = \pi_*\mathcal{E}$ differ by scaling by k. That is, if $\mathcal{P}(E, x)$ and $\mathcal{P}(\mathcal{E}, x)$ denote the Hilbert polynomials of E and $\mathcal{E}$ respectively, then $\mathcal{P}(\mathcal{E}, x) = \mathcal{P}(E, kx)$;

(ii) if (E, ϕ) is the Higgs bundle on Σ corresponding to $\mathcal{E}$ under the spectral correspondence, then $\mathcal{E}$ is semistable[10] if and only if (E, ϕ) is semistable;

(iii) given (E, ϕ) as above, then (E, ϕ) has Higgs Harder-Narasimhan type $\mu = \left(\frac{d_1}{r_1}, \ldots, \frac{d_s}{r_s}\right)$ if and only if $\mathcal{E}$ has Harder-Narasimhan type $k^*\mu = (\mathcal{P}_1, \ldots, \mathcal{P}_s)$ where

$$\mathcal{P}_i(x) = r_i \deg(\mathcal{O}_\Sigma(1))x + d_i + r_i(1-g)$$

is the Hilbert polynomial of the vector bundle $E_i = E^i/E^{i-1}$ determined by the Higgs Harder-Narasimhan filtration of E.

These statements all follow from [94]: the first is established in the text preceding [94, Lem 6.9] while the second point corresponds to [94, Lem 6.9]. The third follows in a straightforward way from the previous statement: the spectral correspondence provides a correspondence between the maximally destabilising subbundle of $\mathcal{E}$ and the maximally destabilising Higgs subbundle of E, and so by induction we obtain a correspondence between the Harder-Narasimhan filtration of $\mathcal{E}$ and the Higgs Harder-Narasimhan filtration of E.

The spectral correspondence described above can be used to identify a Higgs Harder-Narasimhan stratum as a quotient stack, and this is the key to enabling the application of Non-Reductive GIT to the construction of moduli spaces for unstable Higgs bundles. Proposition 3.2.1 establishes this identification.

Proposition 3.2.1 (Identification of a Higgs Harder-Narasimhan stratum as a quotient stack)**.** Let μ be a Higgs Harder-Narasimhan type of rank r and degree d. Then there exists a quasi-projective variety $\mathring{Q}_n^\mu$ acted upon by a reductive group G_n such that for n sufficiently large,

$$\mathscr{H}^\mu_{r,d}(\Sigma, L) \cong \left[\mathring{Q}_n^\mu / G_n\right]. \tag{3.4}$$

Proof. We consider again the set-up and notation introduced above to describe the spectral correspondence. Let $\mathcal{P}$ denote the Hilbert polynomial given by $\mathcal{P}(x) = k\deg(\mathcal{O}_\Sigma(1))x + d + r(1-g)$ (in other words $\mathcal{P}$ is the Hilbert polynomial of a vector bundle E on Σ of rank r and degree d). We let $\mathscr{C}_{\mathcal{P},1}(Y)$ denote the stack of coherent sheaves of pure dimension one on the projective completion Y of $\operatorname{Tot} L$. Moreover, we let $\mathscr{C}_{\mathcal{P},1}(\operatorname{Tot} L)$ denote the open substack of $\mathscr{C}_{\mathcal{P},1}(Y)$ consisting of coherent sheaves $\mathcal{E}$ with a support that does not intersect the divisor $D = Y \setminus \operatorname{Tot} L$. It follows from the spectral correspondence that there is an isomorphism of

[10] In the sense of Giesecker-stability, see [55, Def 5.2].

stacks

$$\mathscr{C}_{\mathcal{P},1}(\operatorname{Tot} L) \cong \mathscr{H}_{r,d}(\Sigma, L).$$

Moreover, as we have seen in our review of the spectral correspondence above, a Higgs bundle (E, ϕ) of Higgs Harder-Narasimhan type μ corresponds to a sheaf $\mathcal{E}$ of Harder-Narasimhan type $k^*\mu$. Thus the isomorphism above restricts to an isomorphism

$$\mathscr{C}_{\mathcal{P},1}^{k^*\mu}(\operatorname{Tot} L) \cong \mathscr{H}_{r,d}^{\mu}(\Sigma, L), \tag{3.5}$$

where $\mathscr{C}_{\mathcal{P},1}^{k^*\mu}(\operatorname{Tot} L)$ is defined to be the intersection of $\mathscr{C}_{\mathcal{P},1}(\operatorname{Tot} L)$ with the locally closed substack $\mathscr{C}_{\mathcal{P},1}^{k^*\mu}(Y)$ of $\mathscr{C}_{\mathcal{P},1}(Y)$ parametrising sheaves of Harder-Narasimhan type $k^*\mu$.

The substack $\mathscr{C}_{\mathcal{P},1}^{k^*\mu}(Y)$ can be identified as a quotient stack by the results of [55, 54] which build on Simpson's construction of a moduli space for semistable sheaves on any projective scheme [96]. The identification is as follows.

Simpson's construction uses the quot scheme $\operatorname{Quot}(V_n \otimes \mathcal{O}_Y(-n), \mathcal{P})$ where V_n is a vector space of dimension $\mathcal{P}(n)$. This scheme parametrises quotients of the sheaf $V_n \otimes \mathcal{O}_Y(-n)$ on Y with Hilbert polynomial $\mathcal{P}$. There is a natural action of $\mathrm{GL}(V_n)$ on this quot scheme such that orbits for the action coincide with isomorphism classes of such quotient sheaves (see [96, Thm 1.21] or [56, §4.3.2]). We let $Q_n \subseteq \operatorname{Quot}(V_n \otimes \mathcal{O}_Y(-n), \mathcal{P})$ denote the $\mathrm{GL}(V_n)$-invariant open subscheme consisting of quotients $\rho : V_n \otimes \mathcal{O}_Y(-n) \to \mathcal{E}$ such that $\mathcal{E}$ has pure dimension one and such that the induced map of sections $H^0(Y, \rho(n)) : V_n \to H^0(Y, \mathcal{E}(n))$ is an isomorphism where $\mathcal{E}(n) := \mathcal{E} \otimes \mathcal{O}_Y(n)$. It is shown in [96, Thm 1.1] that the set of semistable sheaves with Hilbert polynomial $\mathcal{P}$ is bounded, from which it follows that for n sufficiently large, any coherent sheaf $\mathcal{E}$ of pure dimension one can be identified as a point in Q_n. As a result, the open substack $\mathscr{C}_{\mathcal{P},1}^{ss}(Y) \subseteq \mathscr{C}_{\mathcal{P},1}(Y)$ of semistable sheaves can be identified as a quotient stack, for n sufficiently large (see [54, Rk 5.14]):

$$\mathscr{C}_{\mathcal{P},1}^{ss}(Y) \cong \left[Q_n^{ss} / \mathrm{GL}(V_n)\right],$$

where $Q_n^{ss} \subseteq Q_n$ denotes the open subscheme of Q_n parametrising semistable sheaves. This result is extended to sheaves with Hilbert polynomial $\mathcal{P}$ of a fixed Harder-Narasimhan type μ in [54, Prop 5.21]. That is, it is shown in [55, Lem 6.3] that the set of sheaves with Hilbert polynomial $\mathcal{P}$ and with a fixed Harder-Narasimhan type μ is bounded. As a result, for n

sufficiently large, any such sheaf can be identified as a point in Q_n. Thus for any Harder-Narasimhan type μ we obtain an isomorphism for n sufficiently large (see [54, Prop 5.21]):

$$\mathscr{C}^{\mu}_{\mathcal{P},1}(Y) \cong [Q^{\mu}_n / \operatorname{GL}(V_n)] ,$$

where $Q^{\mu}_n \subseteq Q_n$ denotes the locally closed subscheme of Q_n parametrising sheaves of Harder-Narasimhan type μ.

Hence by defining $\mathring{Q}^{\mu}_n \subseteq Q^{k^*\mu}_n$ to be the open subscheme of $Q^{k^*\mu}_n$ consisting of quotient sheaves on Y with a support that does not intersect the divisor $D = Y \setminus \operatorname{Tot} L$, and by setting $G_n := \operatorname{GL}(V_n)$, we obtain the desired isomorphism for n sufficiently large:

$$\mathscr{H}^{\mu}_{r,d}(\Sigma, L) \cong \mathscr{C}^{k^*\mu}_{\mathcal{P},1}(\operatorname{Tot} L) \cong \left[\mathring{Q}^{\mu}_n / G_n\right] .$$

□

3.2.2 Construction of the refined stratification

In this section we use Proposition 3.2.1 to construct a refinement of the Higgs Harder-Narasimhan stratification which is a refined instability stratification. The theorem we prove is the following:

Theorem 3.2.2 (Refined Higgs Harder-Narasimhan stratification)**.** There exists a refinement of the Higgs Harder-Narasimhan stratification of $\mathscr{H}_{r,d}(\Sigma, L)$, called the *refined Higgs Harder-Narasimhan stratification*, which is a refined instability stratification. Moreover, the open stratum of this stratification coincides with the stable stratum $\mathscr{H}^{s}_{r,d}(\Sigma, L)$, provided the latter is non-empty.

Remark 3.2.3 (Moduli-theoretic interpretation of the refined Higgs Harder-Narasimhan stratification)**.** The Non-Reductive GIT construction we will give of the refined Higgs Harder-Narasimhan stratification depends a priori on some choices, since it relies on the Projective Completion algorithm which, as we have noted in Remark 1.3.4, is not necessarily deterministic. Nevertheless, just as we expect that the Projective Completion algorithm can be made canonical (see Remark 1.3.4), we similarly expect that the refined Higgs Harder-Narasimhan stratification is independent of the choices made in its construction, and that it can be interpreted in a moduli-theoretic way. Indeed, we show this in the rank 2 case in Section 3.2.3

below, by providing a complete moduli-theoretic interpretation of the refined Higgs Harder-Narasimhan stratification. This is why we have referred to the stratification of Theorem 3.2.2 as *the* refined Higgs Harder-Narasimhan stratification. We hope in future work, beyond this thesis, to obtain a moduli-theoretic interpretation in the general rank case. We intend to do so using an alternative strategy, described in Remark 3.2.6 below) because the Non-Reductive GIT construction of this refined stratification in the higher rank case is much more difficult to describe explicitly than in the rank 2 case.

Since the stratification of Theorem 3.2.2 above is a refinement of the Higgs Harder-Narasimhan stratification, it produces for each Higgs Harder-Narasimhan stratum a distinguished open stratum. Provided that the refined stratification can be described in a moduli-theoretic way, then the refined Higgs Harder-Narasimhan stratification can be used to define a notion of 'stability' within a given Higgs Harder-Narasimhan stratum.

Definition 3.2.4 (μ-stability)**.** Given a Higgs Harder-Narasimhan type μ of rank r and degree d, we let $\mathscr{H}^{\mu\text{-}s}_{r,d}(\Sigma, L)$ denote the refined stratum of the stratification given in Theorem 3.2.2 which is open and in general non-empty[11]. A Higgs bundle (E, ϕ) is *μ-stable* if $[(E, \phi)] \in \mathscr{H}^{\mu\text{-}s}_{r,d}(\Sigma, L)$. The corresponding quasi-projective coarse moduli space is denoted $\mathcal{M}^{\mu\text{-}s}_{r,d}(\Sigma, L)$, and called the *moduli space of μ-stable Higgs bundles.*

Remark 3.2.5 (Analogy of μ-stability with stability and moduli-theoretic interpretation)**.** If $\mu = \mu_0$, then as stated in Theorem 3.2.2 we have that μ_0-stability coincides with stability and so $\mathcal{M}^{\mu_0\text{-}s}_{r,d}(\Sigma, L) = \mathcal{M}^{s}_{r,d}(\Sigma, L)$. In this way, provided we can obtain an explicit moduli-theoretic interpretation of the refined Higgs Harder-Narasimhan stratification, then we can view μ-stability for unstable Higgs bundles of Higgs Harder-Narasimhan type μ as the analogue of stability for semistable Higgs bundles.

In the rank 2 case for which we are able to obtain a complete moduli-theoretic interpretation of the refined Higgs Harder-Narasimhan stratification (see Theorem 3.2.9 below), we can provide a moduli-theoretic interpretation of μ-stability (see Proposition 3.2.14). In the general rank case, while we do not at present have a complete description of μ-stability in terms of intrinsic

[11] We use the qualifier 'in general' since we expect that in some cases the μ-stable stratum may be empty (just as in some cases the stable stratum in $\mathscr{H}_{r,d}(\Sigma, L)$ can be empty), yet we do not want the notion of μ-stability to change in these special cases, just as the notion of stability does not change if the stable stratum happens to be empty.

properties of Higgs bundles, it will follow from the proof of Theorem 3.2.2 that μ-stable Higgs bundles are contained in the substack of Higgs bundles (E,ϕ) satisfying the following properties:

(i) $(E,\phi) \ncong \mathrm{gr}(E,\phi)$;

(ii) $\dim \mathrm{End}_{-1}(E,\phi) = 0$ where $\mathrm{End}_{-1}(E,\phi)$ consists of endomorphisms of (E,ϕ) such that $\phi(E^i) \subseteq E^{i-1} \otimes L$ where $0 = E^0 \subseteq E^1 \subseteq \cdots \subseteq E^s$ is the Higgs Harder-Narasimhan filtration of E.

This provides a first step towards a moduli-theoretic interpretation of μ-stability in the general rank case.

Proof of Theorem 3.2.2. The construction we give uses Non-Reductive GIT and we start by giving its outline. The strategy is to extend the action of G_n on $\mathring{Q}_n^\mu$ to a linear action of G_n on a projective variety $\overline{\mathring{Q}_n}$ containing $\mathring{Q}_n^\mu$, such that $\mathring{Q}_n^\mu$ is contained in a GIT-unstable stratum $S_{\beta_n(\mu)} \subseteq \overline{\mathring{Q}_n}$ associated to the linear action of G_n on $\overline{\mathring{Q}_n}$ (and to a choice of invariant inner product on $\mathrm{Lie}\, G_n$). Non-Reductive GIT (more precisely Theorem 1.3.6 (iv)) can then be used to stratify this unstable stratum in such a way that each stratum admits a quasi-projective geometric quotient with an explicit projective completion. The restriction of this stratification to $\mathring{Q}_n^\mu \subseteq S_{\beta_n(\mu)}$ induces a stratification of the quotient stack $\left[\mathring{Q}_n^\mu / G_n\right]$, and consequently of $\mathscr{H}_{r,d}^\mu(\Sigma, L)$ by using the identification of Proposition 3.2.1. By its construction, this stratification will satisfy the property that each stratum admits a quasi-projective coarse moduli space with an explicit projective completion.

Having outlined the construction above, we now describe how it can be achieved. It uses the results of [55, 54] concerning the stack of coherent sheaves: thanks to the spectral correspondence, by restricting to those sheaves with support not intersecting the divisor $D \subseteq Y$ we can obtain corresponding results for the stack of Higgs bundles on Σ.

We use the notation introduced in the proof of Proposition 3.2.1. To apply GIT to the action of $\mathrm{SL}(V_n)$ on $\mathring{Q}_n$, it is necessary to obtain a linear action of $\mathrm{SL}(V_n)$ on a projective completion of $\mathring{Q}_n$. A suitable projective completion and linearisation of the $\mathrm{SL}(V_n)$-action is obtained by considering Grothendieck's embedding of the quot scheme $\mathrm{Quot}(V_n \otimes \mathcal{O}_Y(-n), \mathcal{P})$ into a Grassmannian, and then using the Plücker embedding of the Grassmannian into a projective space (see [55, §5]). We note that this embedding depends on an additional parameter m, and we must choose $m \gg n \gg 1$. The action of $\mathrm{SL}(V_n)$ can be shown to extend to this projective space, which we denote by $X_{n,m}$, and to lift to its ample line bundle. We let $\overline{\mathring{Q}_n}$ denote the closure of $\mathring{Q}_n$ inside $X_{n,m}$. Restricting the linearisation of its $\mathrm{SL}(V_n)$ action to $\overline{\mathring{Q}_n}$ provides a linear action of $\mathrm{SL}(V_n)$ on the closure $\overline{\mathring{Q}_n}$ of $\mathring{Q}_n$ inside $X_{n,m}$.

The linear action of $\mathrm{SL}(V_n)$ on $\overline{\mathring{Q}_n}$ induces a GIT-instability stratification (determined by a choice of invariant inner product on the Lie algebra of $\mathrm{SL}(V_n)$) of $\overline{\mathring{Q}_n}$, which we denote by

$$\overline{\mathring{Q}_n} = \bigsqcup_{\beta \in \mathcal{B}} S_\beta.$$

A correspondence is established in [55] between Harder-Narasimhan types μ and GIT-instability types β, given by $\mu \mapsto \beta_{n,m}(\mu)$, such that any sheaf in $\overline{\mathring{Q}_n}$ of Harder-Narasimhan type μ lies in the stratum indexed by $\beta_{n,m}(\mu)$ for $m \gg n \gg 1$ (see [55, Lem 6.6]). It follows from [54, Thm 5.19] that for $m \gg n \gg 1$, the subscheme $\mathring{Q}_n^\mu \subseteq \overline{\mathring{Q}_n}$ is a closed subscheme of $S_{\beta_{n,m}(k^*\mu)} \cap \mathring{Q}_n$.

Applying Theorem 1.3.6 to the action of $\mathrm{SL}(V_n)$ on $\overline{\mathring{Q}_n}$ results in a stratification of $\overline{\mathring{Q}_n}$, which by Theorem 1.3.6 (iv) refines the GIT-instability stratification of $\overline{\mathring{Q}_n}$. Thus we obtain an $\mathrm{SL}(V_n)$-equivariant stratification of $S_{\beta_{n,m}(k^*\mu)}$ satisfying the property that each stratum admits a quasi-projective geometric quotient with an explicit projective completion, and we can restrict this stratification to the intersection $\mathring{Q}_n \cap S_{\beta_{n,m}(k^*\mu)}$ to obtain a stratification of this intersection satisfying the same properties.

The stratification of $\mathring{Q}_n \cap S_{\beta_{n,m}(k^*\mu)}$ can be further restricted to its closed subscheme $\mathring{Q}_n^\mu$. This $\mathrm{SL}(V_n)$-stratification of $\mathring{Q}_n^\mu$ induces a stratification of the quotient stack $\left[\mathring{Q}_n^\mu / G_n\right]$, and thus by Proposition 3.2.1 we obtain a stratification of $\mathscr{H}_{r,d}^\mu(\Sigma, L)$. By construction, it satisfies the property that each stratum has a quasi-projective coarse moduli space with an explicit projective completion. By combining these stratifications for each Higgs Harder-Narasimhan type μ, we obtain as required a refinement of the Higgs Harder-Narasimhan stratification of $\mathscr{H}_{r,d}(\Sigma, L)$ which is a refined instability stratification.

Finally, we note that the stratification of the semistable stratum S_0 of $\overline{\mathring{Q}_n}$ has as its open stratum the stable locus $\overline{\mathring{Q}_n}^s \subseteq \overline{\mathring{Q}_n}$, by Remark 1.3.2 (assuming this stable locus is non-empty). The intersection of this open stratum with $\mathring{Q}_n$ coincides with $\mathring{Q}_n^s$, namely points in the quot scheme parametrising stable sheaves, by [96, Thm 1.19]. Thus under the identification of $\mathscr{H}_{r,d}^{ss}(\Sigma, L)$ with $\left[\mathring{Q}_n^{ss} / G_n\right]$ for n sufficiently large, the resulting stratification has as its open stratum the stratum $\mathscr{H}^{s}_{r,d}(\Sigma, L)$ of stable Higgs bundles (provided the stable stratum is non-empty). □

Remark 3.2.6 (An alternative approach to stratifications)**.** As noted in Remark 3.2.3, at present we do not have an explicit moduli-theoretic interpretation of the refined Higgs Harder-Narasimhan stratification in the general rank case, nor of μ-stability, although we expect one would exist. The difficulty in obtaining it lies in tracking each step of the Non-Reductive GIT construction of the stratification and interpreting it in terms of intrinsic properties of Higgs bundles, which is difficult to do in general due to the numerous blow-ups involved (we have achieved this so far only in the rank 2 case, see Section 3.2.3 below).

Nevertheless, considering the problem from a different perspective might lead more readily to a moduli-theoretic interpretation, starting with the perspective on stratifications developed in [72]. Indeed, stratifications of strictly semistable loci are considered in [72] where it is shown that given the linear action of a reductive group G on a smooth projective variety, the GIT-instability stratification (see Section 1.1.1 of Chapter 1) can be refined to obtain a stratification by locally closed non-singular G-invariant subvarieties such that the open stratum is given by the stable locus (in contrast to the GIT-instability stratification which has the semistable locus as its open stratum). The remaining strata are defined inductively in terms of sets of stable points for the linear action of G on closed non-singular subvarieties of X and on their projectivised normal bundles.

In [70], this refinement is applied to the moduli space of semistable vector bundles on a smooth projective curve. The motivation for this application was to determine a complete set of relations for the generators of the cohomology ring of the moduli space, although this was obtained via different means in [25]. The problem addressed in [70] is how to interpret the stratification of the strictly semistable locus in terms of intrinsic invariants for vector bundles, that is, in terms of natural refinements of the notion of the Harder-Narasimhan type of a vector bundle. An answer is provided through the definition of a notion of *balanced δ-filtrations* (see [70, §6]).

Refinements of Harder-Narasimhan type filtrations are also considered, albeit from a different perspective, in the recent paper [42], where a suggested application is indeed the stratification of stacks of semistable objects by type of the weight-filtration, as per [72]. These refinements are defined iteratively (leading to the appearance of 'iterated logarithms'), and can be interpreted in terms of gradient flows on quiver representations.

We hope in future work to investigate whether this perspective and its formalism might be applied and generalised to the unstable strata in order to obtain a moduli-theoretic interpretation of the refined stratification of Theorem 3.2.2 in a more straightforward manner, namely without having to rely on the Projective Completion algorithm.

3.2.3 The rank 2 case

In this section we consider the problem of interpreting the refined Higgs Harder-Narasimhan stratification constructed in Section 3.2.2 above in terms of intrinsic properties of Higgs bundles. We examine the case of (L-twisted) Higgs bundles of rank 2 and with odd degree, and prove Theorem 3.2.9 which provides a complete moduli-theoretic interpretation of the refined Higgs Harder-Narasimhan stratification.

We will prove Theorem 3.2.9 using the same approach as adopted in the previous two sections, namely using the spectral correspondence to enable the application of an existing result for sheaves. In this case the relevant result is [59, Thm 5.3.2.5], which establishes using Non-Reductive GIT the existence of moduli spaces for unstable sheaves with a Harder-Narasimhan type $\mu = \left(\frac{d_1}{r_1}, \dots, \frac{d_s}{r_s}\right)$ of length[12] two.

The statement of Theorem 3.2.9 relies on the following notation and definitions.

Notation 3.2.7 (Nilpotent endomorphisms)**.** Given a sheaf $\mathcal{E}$ with Harder-Narasimhan filtration[13] $0 = \mathcal{E}^0 \subseteq \mathcal{E}^1 \subseteq \cdots \subseteq \mathcal{E}^s = \mathcal{E}$, we let $\mathrm{End}_{-1}(\mathcal{E})$ denote the subset of $\mathrm{End}(\mathcal{E})$ given by

$$\mathrm{End}_{-1}(\mathcal{E}) = \{\psi \in \mathrm{End}(\mathcal{E}) \mid \psi(\mathcal{E}^i) \subseteq \mathcal{E}^{i-1} \text{ for every } i = 1, \dots, s\}.$$

Given a Higgs bundle (E, ϕ) with Higgs Harder-Narasimhan filtration $0 = E^0 \subseteq E^1 \subseteq \cdots \subseteq E^s = E$, we let $\mathrm{End}_{-1}(E, \phi)$ denote the subset of $\mathrm{End}(E, \phi)$ given by

$$\mathrm{End}_{-1}(E, \phi) = \{\psi \in \mathrm{End}(E, \phi) \mid \psi(E^i) \subseteq E^{i-1} \text{ for every } i = 1, \dots, s\}.$$

We call elements of $\mathrm{End}_{-1}(E, \phi)$ or $\mathrm{End}_{-1}(\mathcal{E})$ *nilpotent endomorphisms.*

[12] The length of a Harder-Narasimhan type $\mu = \left(\frac{d_1}{r_1}, \dots, \frac{d_s}{r_s}\right)$ is the integer s.

[13] Though we have only defined Harder-Narasimhan filtrations for Higgs bundles and vector bundles in Section 3.1, an analogous definition can be made for sheaves, using the notion of Giesecker-stability (see [55, Def 6.1] for example).

Definition 3.2.8 (Refined Higgs Harder-Narasimhan strata)**.** Let $\mu \neq \mu_0$ be a Higgs Harder-Narasimhan type of rank r and degree d. We define the following two substacks of $\mathscr{H}^{\mu}_{2,d}(\Sigma, L)$:

(i) the stack $\mathscr{H}^{\mu,\mathrm{dec}}_{r,d}(\Sigma, L)$ consisting of Higgs bundles (E, ϕ) satisfying $(E, \phi) \cong \mathrm{gr}(E, \phi)$;

(ii) the stack $\mathscr{H}^{\mu,\mathrm{indec}}_{r,d}(\Sigma, L)$ consisting of Higgs bundles (E, ϕ) satisfying $(E, \phi) \ncong \mathrm{gr}(E, \phi)$.

Moreover, given $\delta \in \mathbb{N}$, we let $\mathscr{H}^{\mu,\delta,\mathrm{indec}}_{r,d}(\Sigma, L)$ denote the substack of $\mathscr{H}^{\mu,\mathrm{indec}}_{r,d}(\Sigma, L)$ consisting of Higgs bundles (E, ϕ) satisfying $\dim \mathrm{End}_{-1}(E, \phi) = \delta$.

We can now describe the refined Higgs Harder-Narasimhan stratification constructed in the proof of Theorem 3.2.2 for L-twisted Higgs bundles of rank 2 and odd degree.

Theorem 3.2.9 (When $r = 2$ and d is odd)**.** Given an odd degree d, the refined Higgs Harder-Narasimhan stratification of $\mathscr{H}_{2,d}(\Sigma, L)$ constructed in the proof of Theorem 3.2.2 is given by

$$\mathscr{H}_{2,d}(\Sigma, L) = \mathscr{H}^{ss}_{2,d}(\Sigma, L) \sqcup \bigsqcup_{\mu \neq \mu_0} \bigsqcup_{\delta \in \mathbb{N}} \mathscr{H}^{\mu,\delta,\mathrm{indec}}_{2,d}(\Sigma, L) \sqcup \mathscr{H}^{\mu,\mathrm{dec}}_{2,d}(\Sigma, L). \tag{3.6}$$

Remark 3.2.10 (The refined Higgs Harder-Narasimhan stratification is a stratification)**.** As per Definition 3.1.10, a refined instability stratification must in particular be a stratification. To show that (3.6) is a stratification, we make explicit its indexing set. If we let $H_{2,d}$ denote the set of Higgs Harder-Narasimhan types of rank 2 and degree d, then the indexing set can be identified with the subset of $H_{2,d} \times \mathbb{N} \times \{0, 1\}$ consisting of triples (μ, δ, i) such that $i = \delta = 0$ if $\mu = \mu_0$, and $\delta = 0$ if $i = 1$. Given a triple (μ, δ, i) in this subset, the corresponding stratum $S_{(\mu,\delta,i)}$ is given by

$$S_{(\mu,\delta,i)} := \begin{cases} \mathscr{H}^{ss}_{2,d}(\Sigma, L) & \text{if } \mu = \mu_0; \\ \mathscr{H}^{\mu,\delta,\mathrm{indec}}_{2,d}(\Sigma, L) & \text{if } i \neq 1; \\ \mathscr{H}^{\mu,\mathrm{dec}}_{2,d}(\Sigma, L) & \text{if } i = 1. \end{cases}$$

The partial order is given by the condition that $(\mu, \delta, i) < (\mu', \delta', i')$ if either of the following three statements hold:

(i) $\mu < \mu'$ (with respect to the partial order for Higgs Harder-Narasimhan types introduced in Remark 3.1.2);

(ii) $\mu = \mu'$, $i = i'$ and $\delta < \delta'$;

(iii) $\mu = \mu'$, $\delta' = 0$ and $i' > i$.

To see that (3.6) is a stratification with respect to this partial order, it suffices to note the following topological properties of the refined strata:

(i) the substack $\mathscr{H}^{\mu,\mathrm{dec}}_{r,d}(\Sigma, L)$ is closed in $\mathscr{H}^{\mu}_{r,d}(\Sigma, L)$, while the substack $\mathscr{H}^{\mu,\mathrm{indec}}_{r,d}(\Sigma, L)$ is open in $\mathscr{H}^{\mu}_{r,d}(\Sigma, L)$ (this follows from the openness of indecomposability, see [15, §4.1][14]);

(ii) the decomposition $\mathscr{H}^{\mu,\mathrm{indec}}_{2,d}(\Sigma, L) = \bigsqcup_{\delta\in\mathbb{N}} \mathscr{H}^{\mu,\delta,\mathrm{indec}}_{2,d}(\Sigma, L)$ is a stratification in the sense that for any $\delta \in \mathbb{N}$, there is an inclusion

$$\overline{\mathscr{H}^{\mu,\delta,\mathrm{indec}}_{2,d}(\Sigma, L)} \subseteq \mathscr{H}^{\mu,\delta,\mathrm{indec}}_{2,d}(\Sigma, L) \sqcup \bigsqcup_{\delta'>\delta} \mathscr{H}^{\mu,\delta',\mathrm{indec}}_{2,d}(\Sigma, L).$$

This follows from the upper semi-continuity of fibre dimension.

Remark 3.2.11 (Assumption that d is odd)**.** In the above Theorem 3.2.9, we have assumed that the degree d is odd. This assumption is only required for the semistable stratum, to ensure that it has a quasi-projective coarse moduli space (if the degree is even, then the semistable stratum must be further stratified, with refined open stratum given by the stable locus). By contrast, the refined strata for the unstable locus admit quasi-projective coarse moduli spaces with explicit projective completions even if the degree is even.

Proof of Theorem 3.2.9. The assumption that the degree d is odd ensures that $\mathscr{H}^{ss}_{2,d}(\Sigma, L)$ coincides with $\mathscr{H}^{s}_{2,d}(\Sigma, L)$. By Theorem 3.2.2, it follows that it is a stratum of the refined Higgs Harder-Narasimhan stratification constructed in the proof of Theorem 3.2.2. Thus we need only consider the unstable Higgs Harder-Narasimhan strata.

Let $\mu = (d_1, d_2)$ denote an unstable Higgs Harder-Narasimhan type of rank 2. Applying the isomorphism of (3.5), which follows from the spectral correspondence, we have that

$$\mathscr{H}^{\mu}_{2,d}(\Sigma, L) \cong \mathscr{C}^{k^*\mu}_{\mathcal{P},1}(\mathrm{Tot}\, L) \subseteq \mathscr{C}^{k^*\mu}_{\mathcal{P},1}(Y)$$

where we recall (as introduced in Section 3.2.1) that:

(i) Y is a projective completion of $\mathrm{Tot}\, L$ and k an integer such that the line bundle $\pi^*\mathcal{O}_\Sigma(1)\otimes_{\mathcal{O}_Y} \mathcal{O}_Y(D)$ on Y is ample, where $D = Y \setminus \mathrm{Tot}\, L$ and $\pi : Y \to \Sigma$ is the extension of the projection map $\pi : \mathrm{Tot}\, L \to \Sigma$;

(ii) $\mathcal{P}$ is the Hilbert polynomial defined by $\mathcal{P}(x) = 2\deg(\mathcal{O}_\Sigma(1))x + d + 2(1 - g)$;

(iii) $k^*\mu = (\mathcal{P}_1, \mathcal{P}_2(x))$ where $\mathcal{P}_i(x) = \deg(\mathcal{O}_\Sigma(1))x + d_i + 1 - g$ for $i = 1, 2$.

[14] This is proved in [15, §4.1] for vector bundles (and holds more generally for sheaves) rather than for Higgs bundles. We expect that the proof could be generalised to Higgs bundles. Alternatively, by identifying L-twisted Higgs bundles as sheaves on the total space of L via the spectral correspondence, we can directly apply the result in the case of sheaves.

In [59, Ch 5], Jackson uses Non-Reductive GIT to construct moduli spaces for unstable coherent sheaves with a Harder-Narasimhan type τ of length 2 on a projective variety. The strategy is precisely that described in the proof of Proposition 3.2.1 above, applied in the special case where the Harder-Narasimhan type has length two[15]. That is, the strategy consists in explicitly describing the output of the Projective Completion algorithm applied to the linear action of $\mathrm{SL}(V_n)$ on the closure of Q_n^τ in the projective space in which Q_n is embedded (the linearisation of the action of $\mathrm{SL}(V_n)$ on Q_n is obtained via an $\mathrm{SL}(V_n)$-embedding of Q_n into a projective space, as seen in the proof of Proposition 3.2.1). By reapplying this procedure inductively to the complement of the open stratum obtained, a stratification of the Harder-Narasimhan stratum $\mathscr{C}^\tau_{\mathcal{P},1}(Y)$ is achieved, satisfying the property that each stratum admits a quasi-projective coarse moduli space with an explicit projective completion[16]. It is given by (see [59, Thm 5.3.2.5]):

$$\mathscr{C}^\tau_{\mathcal{P},1}(Y) = \bigsqcup_{\delta\in\mathbb{N}} \mathscr{C}^{\tau,\delta,\mathrm{indec}}_{\mathcal{P},1}(Y) \sqcup \mathscr{C}^{\tau,\mathrm{dec}}_{\mathcal{P},1}(Y) \tag{3.7}$$

where the substacks $\mathscr{C}^{\tau,\delta,\mathrm{indec}}_{\mathcal{P},1}(Y)$ and $\mathscr{C}^{\tau,\mathrm{dec}}_{\mathcal{P},1}(Y)$ are defined analogously to the substacks $\mathscr{H}^{\mu,\delta,\mathrm{indec}}_{r,d}(\Sigma, L)$ and $\mathscr{H}^{\mu,\mathrm{dec}}_{r,d}(\Sigma, L)$ of $\mathscr{H}^{\mu}_{r,d}(\Sigma, L)$ (see Definition 3.2.8).

Setting $\tau = k^*\mu$ (noting that since μ is a Higgs Harder-Narasimhan type of rank 2, τ is a Harder-Narasimhan type of length 2), we obtain a stratification of $\mathscr{C}^{k^*\mu}_{\mathcal{P},1}(Y)$ using (3.7), which restricts to a stratification of the open substack $\mathscr{C}^{k^*\mu}_{\mathcal{P},1}(\mathrm{Tot}\, L) \subseteq \mathscr{C}^{k^*\mu}_{\mathcal{P},1}(Y)$ consisting of coherent sheaves of pure dimension one on Y with a support which does not intersect the divisor $D = Y \setminus \mathrm{Tot}\, L$. By (3.5) there is an isomorphism $\mathscr{C}^{k^*\mu}_{\mathcal{P},1}(\mathrm{Tot}\, L) \cong \mathscr{H}^{\mu}_{2,d}(\Sigma, L)$. Under this isomorphism, which in particular preserves the endomorphisms of the objects (therefore the dimension of nilpotent endomorphims) and the property of being decomposable or indecomposable (as a sheaf or as a Higgs bundle respectively), we obtain a stratification of $\mathscr{H}^{\mu}_{2,d}(\Sigma, L)$ given by:

$$\mathscr{H}^{\mu}_{2,d}(\Sigma, L) = \bigsqcup_{\delta\in\mathbb{N}} \mathscr{H}^{\mu,\delta,\mathrm{indec}}_{2,d}(\Sigma, L) \sqcup \mathscr{H}^{\mu,\mathrm{dec}}_{2,d}(\Sigma, L).$$

By the construction of the stratification using Non-Reductive GIT (more precisely the Projective

[15] In fact [59] considers the more general case of coherent sheaves, rather than coherent sheaves of pure dimension, but the exact same construction can be used for the substack of coherent sheaves of pure dimension. Since we are interested in Higgs bundles, which correspond via the spectral correspondence to coherent sheaves of pure dimension, we do not need to consider the more general case of coherent sheaves.

[16] We note that the results of [59] are not formulated in the language of stacks, and that we are simply restating the results in this language.

Completion algorithm), it follows that each stratum admits a quasi-projective coarse moduli space with an explicit projective completion. $\square$

The above Theorem 3.2.9 provides a complete moduli-theoretic interpretation in the rank 2 case of the refined Higgs Harder-Narasimhan stratification constructed in the proof of Theorem 3.2.2. Moreover, provided $\mathscr{H}_{2,d}^{\mu,\mathrm{indec}}(\Sigma, L)$ is non-empty, then in general we expect the (non-empty) open stratum of the refined stratification to be given by $\mathscr{H}_{2,d}^{\mu,0,\mathrm{indec}}(\Sigma, L)$, based on Proposition 3.2.12 and Remark 3.2.13 below. We note that Proposition 3.2.12 is valid for Higgs bundles of any rank (not just rank 2) and we therefore state it in this generality.

Proposition 3.2.12 (Constructing Higgs bundles with no nilpotent endomorphisms)**.** Let $\mu \neq \mu_0$ denote a Higgs Harder-Narasimhan type and let E denote a vector bundle of Harder-Narasimhan type μ such that $E \cong \operatorname{gr} E = E_1 \oplus \cdots \oplus E_s$. Then there exists a Higgs field $\phi : E \to E \otimes L$ such that

$$\mathrm{End}_{-1}(E, \phi) = \{0\}.$$

Proof. We fix an isomorphism $E \cong \operatorname{gr} E$ and write $E = E_1 \oplus \cdots \oplus E_s$. To define a suitable Higgs field ϕ, choose s pairwise distinct sections $\sigma_i \in H^0(\Sigma, L)$ (recall that we have assumed from the start that $h^0(\Sigma, L) \neq 0$). Then each section σ_i defines a map $\phi_i : E_i \to E_i \otimes L$ and we obtain a Higgs field $\phi : E = E_1 \oplus \cdots \oplus E_s \to E \otimes L$ by setting $\phi = \phi_1 \oplus \cdots \oplus \phi_s$. We aim to show that $\mathrm{End}_{-1}(E, \phi) = \{0\}$. To this end, suppose that $\psi \in \mathrm{End}_{-1}(E, \phi)$, so that the following diagram commutes:

$$\begin{array}{ccc} E & \xrightarrow{\psi} & E \\ \downarrow{\scriptstyle \phi} & & \downarrow{\scriptstyle \phi} \\ E \otimes L & \xrightarrow{\psi \otimes \mathrm{id}_L} & E \otimes L. \end{array}$$

Restricting to $E^i \subseteq E$, post-composing with the projection map $E \to E_j$ (we assume that $i \neq j$, and letting $\psi_{ij} : E_i \to E_j$ denote the corresponding morphism leads to the following commutative diagram:

$$\begin{array}{ccc} E_i & \xrightarrow{\psi_{ij}} & E_j \\ \downarrow{\scriptstyle \phi_i} & & \downarrow{\scriptstyle \phi_j} \\ E_i \otimes L & \xrightarrow{\psi_{ij} \otimes \mathrm{id}_L} & E_j \otimes L. \end{array}$$

The commutativity of (3.2.3) implies that $\psi_{ij} \otimes (\sigma_i - \sigma_j) = 0 \in H^0(\Sigma, E_i^\vee \otimes E_j \otimes L)$. Thus in particular the zero locus of ψ_{ij}, which is closed in Σ, must contain all points $x \in \Sigma$ where

$\sigma_i(x) \neq \sigma_j(x)$. The locus of such points is open and in Σ since by assumption $\sigma_i \neq \sigma_j$, and thus it is dense in Σ which by assumption is a smooth, hence irreducible, projective curve. It follows that ψ_{ij} must vanish on all of Σ, and as this holds true for all $i, j \in \{1, \ldots, s\}$, we have that ψ is identically zero. Hence $\mathrm{End}_{-1}(E, \phi) = \{0\}$ as required. □

Remark 3.2.13 (μ-stability for unstable Higgs bundles of rank 2)**.** Proposition 3.2.12 above applied in the rank 2 case suggests that, given an unstable Higgs Harder-Narasimhan type μ of rank 2, the minimal value δ such that $\mathscr{H}_{2,d}^{\mu,\delta,\mathrm{indec}}(\Sigma, L)$ is non-empty should in general be $\delta = 0$ (assuming that $\mathscr{H}_{2,d}^{\mu,\mathrm{indec}}(\Sigma, L)$ is non-empty).

Indeed, applying Proposition 3.2.12 in the rank 2 case shows that, given any vector bundle $E_1 \oplus E_2$ of Harder-Narasimhan type μ, then by choosing distinct sections σ_1 and σ_2 of L we can construct a Higgs bundle

$$\left(E_1 \oplus E_1, \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix}\right) \tag{3.8}$$

of Higgs Harder-Narasimhan type μ, which is isomorphic to its graded and which has no non-trivial nilpotent endomorphisms. Since the dimension of nilpotent endomorphisms can only increase upon taking the limit of a Higgs bundle to its associated Higgs Harder-Narasimhan graded, provided there exists an indecomposable Higgs bundle of Harder-Narasimhan type μ with Higgs Harder-Narasimhan graded isomorphic to a Higgs bundle of the form (3.8) (a condition which would not be met in special cases only), a Higgs bundle (E, ϕ) satisfying $\mathrm{End}_{-1}(E, \phi) = \{0\}$ will exist.

For this reason, in general the open and non-empty stratum of the restriction of the refined Higgs Harder-Narasimhan stratification to a Higgs Harder-Narasimhan stratum $\mathscr{H}_{2,d}^{\mu}(\Sigma, L)$ will coincide with $\mathscr{H}_{2,d}^{\mu,0,\mathrm{indec}}(\Sigma, L)$. Thus we obtain a moduli-theoretic characterisation of μ-stability for unstable Higgs bundles (introduced in Definition 3.2.4) in the rank 2 case, as per Proposition 3.2.14 below.

Proposition 3.2.14 (μ-stable Higgs bundles of rank 2)**.** Given a Higgs Harder-Narasimhan type μ of rank 2 and degree d, a Higgs bundle (E, ϕ) of Higgs Harder-Narasimhan type μ is μ-stable if and only if it satisfies the following conditions:

(i) $(E, \phi) \ncong \mathrm{gr}(E, \phi)$;

(ii) $\dim \mathrm{End}_{-1}(E, \phi) = 0$.

We also define the following notion of μ-instability for unstable Higgs bundles of rank 2 and Higgs Harder-Narasimhan type μ. We will study the geometry of the moduli spaces of μ-stable and μ-unstable Higgs bundles of rank 2 and of Higgs Harder-Narasimhan type μ in Section 4.3 of Chapter 4.

Definition 3.2.15 (μ-unstable Higgs bundles)**.** Given a Higgs Harder-Narasimhan type μ of rank 2 and degree d, a Higgs bundle (E, ϕ) of Higgs Harder-Narasimhan type μ is *μ-unstable* if it satisfies the following conditions:

(i) $(E, \phi) \ncong \mathrm{gr}(E, \phi)$;

(ii) $\dim \mathrm{End}_{-1}(E, \phi) > 0$.

3.3 Refined Harder-Narasimhan stratification

In this section we show that the Harder-Narasimhan stratification can be refined using Non-Reductive GIT to produce a refined instability stratification. In Section 3.3.1 we identify each Harder-Narasimhan stratum as a quotient stack (see Proposition 3.3.3) using Nitsure's parametrisation of Higgs bundles from [81]. This identification allows us in Section 3.3.2 to construct a refinement of the Harder-Narasimhan stratification which is a refined instability stratification using Non-Reductive GIT (see Theorem 3.3.5). We do so by extending to unstable Higgs bundles Nitsure's GIT set-up for the construction of the moduli space of semistable Higgs bundles. Finally, in Section 3.3.3 we compare this stratification to the refined Higgs Harder-Narasimhan stratification constructed in Section 3.2 above (including in the rank 2 case). In particular we show that the refined Harder-Narasimhan stratification is a refinement of the so-called 'Higgs stratification' (see Definition 3.3.10) of the stack of Higgs bundles (see Proposition 3.3.11).

3.3.1 Identification of the strata as quotient stacks

In order to construct the refined Harder-Narasimhan stratification using Non-Reductive GIT, it is necessary to achieve a set-up in which results from Non-Reductive GIT can be applied. For the Higgs Harder-Narasimhan stratification considered in Section 3.2, we were able to use an existing set-up for sheaves, via the spectral correspondence which identifies Higgs bundles as a special class of sheaves on a higher dimensional variety. The reason this approach worked is that, under this correspondence, the Higgs Harder-Narasimhan type of a Higgs bundle coincides (up to scaling) with the Harder-Narasimhan of the corresponding sheaf.

Unfortunately this approach does not work for the Harder-Narasimhan stratification, since it is not obvious how we could interpret the Harder-Narasimhan type (rather than the Higgs Harder-Narasimhan type) of a Higgs bundle in terms of a discrete invariant of the corresponding sheaf. Instead, to construct a refinement of the Harder-Narasimhan stratification we extend Nitsure's GIT construction of the moduli space of semistable Higgs bundles, given in [81], starting with an identification of each Harder-Narasimhan stratum as a quotient stack. The construction involves parametrising Higgs bundles in such a way that isomorphism classes correspond to orbits for a group action, and we describe this construction below.

Nitsure's parametrisation of Higgs bundles ([81, §3]). Let $m = d + r(1-g)$ and let $Q_{r,d}$ denote the Quot scheme parametrising quotient sheaves $\mathcal{O}_\Sigma^{\oplus m} \to \mathcal{F}$ on Σ of rank r and degree d. We let $\mathcal{O}_{\Sigma \times Q_{r,d}}^{\oplus m} \to \mathcal{U}_{r,d}$ denote the universal quotient sheaf on $\Sigma \times Q_{r,d}$. There is a natural $\mathrm{GL}(m)$ action on the scheme $Q_{r,d}$ given by

$$A \cdot (q : \mathcal{O}_\Sigma^{\oplus m} \to \mathcal{U}_q) = q \circ A : \mathcal{O}_\Sigma^{\oplus m} \xrightarrow{A} \mathcal{O}_\Sigma^{\oplus m} \xrightarrow{q} \mathcal{U}_q$$

for any $A \in \mathrm{GL}(m)$. Define $R_{r,d} \subseteq Q_{r,d}$ to be the subset of all $q \in Q_{r,d}$ satisfying the properties that:

(i) the sheaf $(\mathcal{U}_{r,d})_q$ is locally free, where $(\mathcal{U}_{r,d})_q = \mathcal{U}_{r,d}|_{\{q\}\times\Sigma}$ is a quotient sheaf of $\mathcal{O}_\Sigma^{\oplus m}$ on Σ of rank r and degree d;

(ii) the canonical map $H^0(\Sigma, \mathcal{O}_\Sigma^{\oplus m}) \to H^0(\Sigma, (\mathcal{U}_{r,d})_q)$ is an isomorphism.

By [80, Thm 5.3], the set $R_{r,d}$ is a $\mathrm{GL}(m)$-invariant open and reduced subset of $\mathcal{Q}_{r,d}$ and $\mathcal{U}_{r,d}|_{R_{r,d}\times\Sigma}$ is locally free, hence a vector bundle over $R_{r,d} \times \Sigma$, denoted $(\mathcal{U}_{r,d})_{R_{r,d}}$. Moreover, by [80, Lem 5.2 & Thm 5.3], the following conditions are satisfied, provided $d > r(2g-1)$:

(i) the family $(\mathcal{U}_{r,d})_R$ of vector bundles over Σ parametrised by $R_{r,d}$ has the local universal property[17] for families of vector bundles of rank r and degree d;

(ii) two vector bundles in $R_{r,d}$ are isomorphic if and only if they lie in the same orbit under the $\mathrm{GL}(m)$ action on $R_{r,d}$.

[17] Given a moduli problem, a family $\mathcal{U}$ parametrised by a scheme B has the local universal property if given any family $\mathcal{F}$ parametrised by a scheme B', and any point $b \in B'$, there exists an open neighbourhood U of b' and a map $U \to B$ such that the restriction of $\mathcal{F}$ to U coincides with the pull-back of $\mathcal{U}$ along the map $U \to B$.

Let $\mathrm{HF}_{r,d}$ denote[18] the functor from the category of $R_{r,d}$-schemes to the category of groups defined by $(f : S \to R_{r,d}) \mapsto H^0(S, \pi_{S*}(\mathrm{id}_\Sigma \times f)^*(\pi_\Sigma^* L \otimes \mathrm{End}((\mathcal{U}_{r,d})_{R_{r,d}}))$, where $\pi_\Sigma : \Sigma \times R \to \Sigma$ denotes the projection onto Σ.

Remark 3.3.1 (Understanding the functor $HF_{r,d}$)**.** By the universal property of $R_{r,d}$ for vector bundles, an $R_{r,d}$-scheme $f : S \to R_{r,d}$ corresponds to a family E_S of vector bundles on Σ parametrised by S. Then $\mathrm{HF}_{r,d}(S \to R_{r,d})$ can be interpreted as the group of all Higgs fields ϕ_S such that (E_S, ϕ_S) is a family of Higgs bundles parametrised by S. In particular, if S is a point, then $S \to R_{r,d}$ is a vector bundle E on Σ and its image under $\mathrm{HF}_{r,d}$ is the group $H^0(\Sigma, L \otimes \mathrm{End}\, E)$ of Higgs fields for E.

By [81, §3], the functor $\mathrm{HF}_{r,d}$ is representable by a linear[19] scheme $f : F_{r,d} \to R_{r,d}$ which satisfies the following properties:

(i) the family $(E_{F_{r,d}}, \phi_{F_{r,d}})$ of Higgs bundles over Σ parametrised by $F_{r,d}$, determined by the identity map $F_{r,d} \to F_{r,d}$, has the local universal property for families comprising Higgs bundles (E, ϕ) of rank r and degree d satisfying:

  (a) E is generated by its sections, and

  (b) $H^1(\Sigma, E) = 0$;

(ii) the $\mathrm{GL}(m)$-action on $R_{r,d}$ lifts to an action on $F_{r,d}$, given by $A \cdot (\mathcal{U}_q, \phi) = (\mathcal{U}_{A \cdot q}, \phi)$;

(iii) two Higgs bundles in $F_{r,d}$ are isomorphic if and only if they lie in the same orbit under the $\mathrm{GL}(m)$-action on $F_{r,d}$.

(iv) Given $(E, \phi) \in F_{r,d}$, there is an isomorphism $\mathrm{Aut}(E, \phi) \cong \mathrm{Stab}_{GL(m)}(E, \phi)$.

We note that the centre of $\mathrm{GL}(m)$ acts trivially on $F_{r,d}$, thus for the purpose of using GIT the action of $\mathrm{SL}(m)$ can be considered instead. Nevertheless at the level of stacks we must work with $\mathrm{GL}(m)$ to record all the automorphisms.

Using the above set-up, we can prove the analogue of Proposition 3.2.1 for the Harder-Narasimhan strata, namely that each Harder-Narasimhan stratum can be identified as a quotient stack. The statement relies on the following definition:

Definition 3.3.2. Two Harder-Narasimhan types $\tau = \left(\frac{d_1}{r_1}, \ldots, \frac{d_1}{r_1}, \frac{d_2}{r_2}, \ldots, \frac{d_s}{r_s}\right)$ and

[18] The notation $\mathrm{HF}_{r,d}$ stands for 'Higgs field'. See Remark 3.3.1.

[19] By linear we mean that the fibres of the map can be identified with vector spaces.

$\tau' = \left(\frac{d'_1}{r'_1}, \ldots, \frac{d'_1}{r'_1}, \frac{d'_2}{r'_2}, \ldots, \frac{d'_t}{r'_t}\right)$ are *equivalent* if $s = t$, if $r_i = r'_i$ for each i, and if there exists some $e \in \mathbb{Z}$ such that $d_i = d'_i + r_i e$ for each i. We let $[\tau]$ denote the equivalence class of a Harder-Narasimhan type under this equivalence relation. We note that if E has Harder-Narasimhan type τ, and M is any line bundle on Σ, then the Harder-Narasimhan type of $E \otimes M$ is equivalent to τ (by taking e to be the degree of M).

Proposition 3.3.3 (Identification of a Harder-Narasimhan stratum as a quotient stack)**.** Let τ be a Harder-Narasimhan type of rank r and degree d. Then there exists a quasi-projective variety $F_\tau^{r,d}$ acted upon by a reductive group G_d such that for d sufficiently large (depending only on $[\tau]$), there is an isomorphism of stacks

$$\mathscr{H}_\tau^{r,d}(\Sigma, L) \cong \left[F_\tau^{r,d}/G_d\right].$$

Remark 3.3.4. In the above Proposition 3.3.3, the assumption must be made that the degree d is sufficiently large, although it only depends on $[\tau]$ rather than on τ. This assumption can be made without loss of generality, since by choosing a line bundle on Σ of positive degree and tensoring the underlying vector bundle of a Higgs bundle by any tensor power of this line bundle, we obtain for any d' congruent to d modulo r an isomorphism of stacks

$$\mathscr{H}_{r,d}(\Sigma, L) \cong \mathscr{H}_{r,d'}(\Sigma, L),$$

noting that the equivalence class of Harder-Narasimhan types is preserved by this isomorphism.

Proof of Proposition 3.3.3. Consider the map $f : F_{r,d} \to R_{r,d}$ from the set-up described above, and fix a Harder-Narasimhan type τ. Let $R_{r,d}^\tau \subseteq R_{r,d}$ denote the subvariety parametrising vector bundles with Harder-Narasimhan type τ, and let $F_\tau^{r,d} := f^{-1}(R_{r,d}^\tau)$. It follows from the definition of $F_{r,d}$ that $F_\tau^{r,d}$ parametrises Higgs bundles (E, ϕ) such that E has Harder-Narasimhan type τ and satisfies:

(a) E is generated by its sections;

(b) $H^1(\Sigma, E) = 0$.

We can then apply [54, Thm 5.19 & Lem 5.2] to conclude that for d sufficiently large (depending on $[\tau]$), any vector bundle of Harder-Narasimhan type τ satisfies conditions (a) and (b) above. Moreover, since two Higgs bundles in $F_{r,d}$ are isomorphic if and only if they lie in the same orbit under the $\mathrm{GL}(m)$-action on $F_{r,d}$, the same is true when we restrict to $F_\tau^{r,d}$.

To identify the Harder-Narasimhan stratum $\mathscr{H}_\tau^{r,d}(\Sigma, L)$ as the quotient stack $\left[F_\tau^{r,d}/\operatorname{GL}(m)\right]$ for d sufficiently large, we proceed as per the proof of [54, Prop 5.21], which establishes an isomorphism of stacks in the case of sheaves of a fixed Harder-Narasimhan type. The restriction of the universal family $(E_{F_{r,d}}, \phi_{F_{r,d}})$ to $F_\tau^{r,d} \times \Sigma \subseteq F_{r,d} \times \Sigma$ has the local universal property for families of Higgs bundles of Harder-Narasimhan type τ, provided d is sufficiently large. Thus we obtain an atlas $F_\tau^{r,d} \to \mathscr{H}_\tau^{r,d}(\Sigma, L)$ for the Harder-Narasimhan stratum $\mathscr{H}_\tau^{r,d}(\Sigma, L)$. Moreover, two morphisms $f_i : B \to F_\tau^{r,d}$ define equivalent families of Higgs bundles of Harder-Narasimhan type τ if and only if there exists a morphism $\varphi : B \to \operatorname{GL}(m)$ such that $f_1(b) = \varphi(b) \cdot f_2(b)$ for all $b \in B$. Thus for d sufficiently large, the map $F_\tau^{r,d} \to \mathscr{H}_\tau^{r,d}(\Sigma, L)$ descends to an isomorphism of stacks

$$\mathscr{H}_\tau^{r,d}(\Sigma, L) \cong \left[F_\tau^{r,d}/G_d\right],$$

where $G_d = \operatorname{GL}(m)$ (we recall that $m = d + r(1-g)$). □

3.3.2 Construction of the refined stratification

The identification of a Harder-Narasimhan stratum as a quotient stack established in Proposition 3.3.3 above enables the application of Non-Reductive GIT to construct a refined instability stratification which is a refinement of the Harder-Narasimhan stratification. We do so in this section. The result we prove is the following theorem:

Theorem 3.3.5 (Refined Harder-Narasimhan stratification)**.** There exists a refinement of the Harder-Narasimhan stratification of $\mathscr{H}_{r,d}(\Sigma, L)$, called the *refined Harder-Narasimhan stratification*, which is a refined instability stratification. Moreover, the open stratum of this stratification coincides with the stratum $\mathscr{H}_s^{r,d}(\Sigma, L) := F^{-1}(\mathscr{V}_{r,d}^s(\Sigma))$, provided the latter is non-empty.

The refined Harder-Narasimhan stratification produces for each Harder-Narasimhan stratum a distinguished open stratum. Thus, provided that the refined stratification can be described in a moduli-theoretic way, then the refined Harder-Narasimhan stratification can be used to define a notion of 'stability' within a given Harder-Narasimhan stratum, just as we defined a notion of 'stability' within a given Higgs Harder-Narasimhan stratum in Definition 3.2.4 above.

Definition 3.3.6 (τ-stability)**.** Given a Harder-Narasimhan type τ, we let $\mathscr{H}_{\tau\text{-}s}^{r,d}(\Sigma, L)$ denote the refined stratum of the stratification given in Theorem 3.3.5 which is open and in general non-empty[20]. A Higgs bundle (E, ϕ) is *τ-stable* if $[(E, \phi)] \in \mathscr{H}^{r,d}_{\tau\text{-}s}(\Sigma, L)$. The corresponding quasi-projective coarse moduli space is denoted $\mathcal{M}^{r,d}_{\tau\text{-}s}(\Sigma, L)$, and called the *moduli space of τ-stable Higgs bundles*.

Remark 3.3.7 (Moduli theoretic interpretation of the refined Harder-Narasimhan stratification and of τ-stability)**.** The question of the moduli-theoretic interpretation of the refined Harder-Narasimhan stratification and of τ-stability arises, just as it did in the case of the refined Higgs Harder-Narasimhan stratification and of μ-stability (see Remarks 3.2.3 and 3.2.5). We will give a partial answer to this question in Section 3.3.3, by showing that the refined Harder-Narasimhan stratification constructed in the proof of Theorem 3.3.5 is a refinement of another stratification (finer than the Harder-Narasimhan stratification) which can be described in a moduli-theoretic way (see Proposition 3.3.11). We will also provide in Remark 3.3.12 a complete moduli-theoretic description of the refined Harder-Narasimhan stratification in the rank 2 case.

Finally, we note that it will follow from the proof of Theorem 3.2.2 that τ-stable Higgs bundles are contained in the substack of $\mathscr{H}^{r,d}_{\tau}(\Sigma, L)$ consisting of Higgs bundles (E, ϕ) such that the map $\phi_{s,1} : E^1 \to E/E^{s-1} \otimes L$ is non-zero, where $0 = E^0 \subseteq E^1 \subseteq \cdots \subseteq E^s = E$ denotes the Harder-Narasimhan filtration of E and $\phi_{s,1}$ is the map induced from the Higgs field ϕ by restricting to E^1 and post-composing with the quotient map $E \to E/E^{s-1}$. This provides a first step towards a moduli-theoretic interpretation of τ-stability in the general rank case.

We will construct the refined Harder-Narasimhan stratification of Theorem 3.3.5 using the isomorphism established in Proposition 3.3.3, by constructing a stratification of the quotient stacks $\left[F^{r,d}_{\tau}/G_d\right]$. We do so using Non-Reductive GIT, as we did in the case of the refined Higgs Harder-Narasimhan stratification which we constructed in Section 3.2 above. Once again, the application of Non-Reductive GIT to the action of G_d on $F^{r,d}_{\tau}$ requires a linearisation of the action on a projective completion of $F^{r,d}_{\tau}$. Such a linearisation and projective completion can be obtained using Nitsure's GIT set-up for the construction of the moduli space of semistable Higgs bundles, summarised below.

Nitsure's GIT set-up for Higgs bundles ([81, §5]). Let $G(r, m)$ denote the Grassmannian of r-dimensional quotients of k^m, and U^m_r the tautological bundle on $G(r, m)$. There is a natural

[20] We use the qualifier 'in general' in the same sense as described in footnote 11.

action of $\mathrm{SL}(m)$ on $G(r,m)$ given by multiplication on the right. For any $x \in \Sigma$, there is a corresponding morphism $\iota_x : R_{r,d} \to G(r,m)$ sending a point $q \in R$ to the fibre at x of the vector bundle U_q over Σ. The fibre in $F_{r,d}$ over a point $q \in R_{r,d}$ can be identified with $H^0(\Sigma, \mathrm{End}(\mathcal{U}_{r,d})_q \otimes L)$, as noted in Remark 3.3.1. Thus a point ϕ in this fibre determines a morphism $\phi : (\mathcal{U}_{r,d})_q \to (\mathcal{U}_{r,d})_q \otimes L$. Fixing a basis for L_x induces an endomorphism $\phi_x : (\mathcal{U}_{r,d})_{q,x} \to (\mathcal{U}_{r,d})_{q,x}$, that is, a point in $\mathrm{End}(\mathcal{U}_{r,d})_{q,x}$.

If we let $H(r,m) = \mathrm{Tot}(\mathrm{End}\, U_r^m)$, then the action of $\mathrm{SL}(m)$ on $G(r,m)$ lifts to an action on $H(r,m)$, by defining $A \cdot \sigma(y) = \sigma(Ay)$ for any section $\sigma \in \mathrm{End}\, U_r^m$ and $A \in \mathrm{GL}(m)$. Since the fibre in $H(r,m)$ over a point $(\mathcal{U}_{r,d})_{q,x} \in G(r,m)$ can be identified with $\mathrm{End}(\mathcal{U}_{r,d})_{q,x}$, we obtain a morphism $\widehat{\iota}_x : F_{r,d} \to H(r,m)$ given by $((\mathcal{U}_{r,d})_q, \phi) \mapsto ((\mathcal{U}_{r,d})_{q,x}, \phi_x)$, and this morphism lies over the morphism ι_x. By choosing N points $x_1, \dots, x_N \in \Sigma$, we obtain $\mathrm{SL}(m)$-equivariant maps ι and $\hat{\iota}$ making the following diagram commute:

$$\begin{array}{ccc} F_{r,d} & \xrightarrow{\hat{\iota}} & H(r,m)^N \\ \big\downarrow f & & \big\downarrow \pi \\ R_{r,d} & \xrightarrow{\iota} & G(r,m)^N, \end{array} \tag{3.9}$$

where $\pi : H(r,m)^N \to G(r,m)^N$ denote the natural projection. Consider the projective variety

$$\widehat{G}(r,m) := \mathbb{P}(\mathcal{O}_{G(r,m)} \oplus \mathrm{End}\, U_r^m) \cong \mathbb{P}((\mathcal{O}_{G(r,m)} \oplus \mathrm{End}\, U_r^m) \otimes (\det U_r^m)^{-1}).$$

Then there is an embedding of $H(r,m)$ inside $\widehat{G}(r,m)$ (corresponding to the points which do not lie in the divisor at infinity), and more generally of $H(r,m)^N$ inside $\widehat{G}(r,m)^N$. The action of $\mathrm{GL}(m)$ on $H(r,m)$ lifts to an action on $\widehat{G}(r,m)$, which can be linearised with respect to $\mathcal{O}_{\widehat{G}(r,m)}(1)$ such that the inclusion $H(r,m) \hookrightarrow \widehat{G}(r,m)$ is $\mathrm{GL}(m)$-equivariant. We also let $\pi : \widehat{G}(r,m)^N \to G(r,m)^N$ denote the natural projection. By [81, §2] there is an identification

$$H^0(\widehat{G}, \mathcal{O}_{\widehat{G}(r,m)}(1)) \cong H^0\left(G(r,m), \det U_r^m \oplus \det U_r^m \otimes \mathrm{End}\, U_r^m\right). \tag{3.10}$$

The line bundle $\mathcal{O}_{\widehat{G}(r,m)}(1)$ is very ample by [81, Prop 2.2] and so we obtain a linearisation of the action of $\mathrm{SL}(m)$ on $\widehat{G}(r,m)^N$ with respect to a very ample line bundle.

While the resulting map $\widehat{\iota} : F_{r,d} \to \widehat{G}(r,m)^N$ is not an embedding, Nitsure shows that it is injective and proper when restricted to suitable open subsets of $F_{r,d}$, provided d and N are sufficiently large. These open subsets are defined as follows:

Notation 3.3.8. Given $A \geq 0$, let $T_{A,r,d}$ denote the set of Harder-Narasimhan types of those vector bundles E of rank r and degree d on Σ which satisfy the inequality $\mu(E') \leq \mu(E) + A$ for any non-zero subbundle $E' \subseteq E$. Let $F_{A,r,d} \subseteq F_{r,d}$ denote the open subvariety of Higgs bundles for which the underlying bundle has Harder-Narasimhan type in $T_{A,r,d}$.

By [81, Prop 5.3], given any $A \geq 0$, for d (depending on A, r, g) and N (depending on d) sufficiently large, there exists a sequence of N points $x_1, \ldots, x_N \in \Sigma$ such that the map $\widehat{\iota}: F_{A,r,d} \to \widehat{G}(r,m)^N$ is injective and proper. In particular, there exists an integer A such that $F^{ss}_{r,d} \subseteq F_{A,r,d}$. GIT can then be applied to the linear action of $\mathrm{SL}(m)$ on the closure of $F^{ss}_{r,d}$ inside $\widehat{G}(r,m)^N$.

We now prove Theorem 3.3.5, using the above set-up.

Proof of Theorem 3.3.5. Let τ denote a Harder-Narasimhan type μ of rank r and degree d. Our aim is to construct a stratification of the quotient stack $\left[F^{r,d}_{\tau}/G_d\right]$ such that each stratum admits a quasi-projective coarse moduli space with an explicit projective completion, and we do so using Non-Reductive GIT.

We can choose an integer $A \geq 0$ such that $\tau \in T_{A,r,d}$, where we recall from Notation 3.3.8 that $T_{A,r,d}$ consists of the set of Harder-Narasimhan types of those vector bundles E of rank r and degree d on Σ which satisfy the inequality $\mu(E') \leq \mu(E) + A$ for any proper subbundle $E' \subseteq E$ (for example $A = \frac{d_1'}{r_1'} - \mu(E)$ is suitable, where $\frac{d_1'}{r_1'}$ is the first entry of τ). Then by [81, Prop 5.3], for d and N sufficiently large, there exists a sequence of N points $x_1, \ldots, x_N \in \Sigma$ such that the map $\widehat{\iota}: F_{A,r,d} \to \widehat{G}(r,m)^N$ is injective and proper. By our choice of A, we have that $F^{r,d}_{\tau} \subseteq F_{A,r,d}$ and thus for d and N sufficiently large, there exists a $\mathrm{GL}(m)$-equivariant map $\widehat{\iota}: F^{r,d}_{\tau} \hookrightarrow \widehat{G}(r,m)^N$ which is injective and proper.

By [65], there exists a correspondence $\tau \mapsto \beta(\tau)$ between Harder-Narasimhan types τ and GIT-instability types associated to the GIT-instability stratification $G(r,m)^N = \bigsqcup_{\beta \in \mathcal{B}} S_\beta$ associated to the linear action of $\mathrm{SL}(m)$ on $G(r,m)^N$ (and to a choice of invariant inner product on the Lie algebra of $\mathrm{SL}(m)$), such that for d sufficiently large (depending only on $[\tau]$) and N sufficiently large (depending on d and on $[\tau]$), there is an inclusion $\iota(R^{\tau}_{r,d}) \subseteq S_{\beta(\tau)}$. Thus we have a commutative diagram, for d and N sufficiently large (as qualified in the previous sentence, and assumed from here on):

$$\begin{array}{ccccc} F^{r,d}_{\tau} & \xhookrightarrow{\widehat{\iota}} & \pi^{-1}(S_{\beta(\tau)}) & \subseteq & \widehat{G}(r,m)^N \\ \downarrow f & & \downarrow \pi & & \downarrow \pi \\ R^{\tau}_{r,d} & \xhookrightarrow{\iota} & S_{\beta(\tau)} & \subseteq & G(r,m)^N. \end{array}$$

For $\tau = \tau_0$, we construct a stratification of $F^{r,d}_{\tau} = F^{r,d}_{ss}$ (and consequently of $\mathscr{H}^{r,d}_{ss}(\Sigma, L)$) by considering the linear action of $\mathrm{SL}(m)$ on the closure $\overline{\widehat{\iota}(F^{r,d}_{ss})}$ of $\widehat{\iota}(F^{r,d}_{ss})$ inside $\widehat{G}(r,m)^N$. The application of the stratification theorem of Non-Reductive GIT (Theorem 1.3.6) results in a stratification of $\overline{\widehat{\iota}(F^{r,d}_{ss})}$ such that each stratum has a quasi-projective geometric quotient with an explicit projective completion. Since $\widehat{\iota}$ restricts to an injective and proper map on $F^{ss}_{r,d}$ provided d and N are sufficiently large, the stratification of $\overline{\widehat{\iota}(F^{r,d}_{ss})}$ can be pulled back via $\widehat{\iota}$ to a stratification of $F^{r,d}_{ss}$. Then, using the isomorphism $\mathscr{H}^{r,d}_{ss}(\Sigma, L) \cong \left[F^{r,d}_{ss}/G_d\right]$ for d sufficiently large, we obtain a stratification of $\mathscr{H}^{r,d}_{ss}(\Sigma, L)$ satisfying the required properties. In particular, the open stratum coincides with the stratum $\mathscr{H}^{r,d}_{s}(\Sigma, L)$, provided it is non-empty. This follows from the fact that the open stratum for the non-reductive GIT stratification induced by the action of $\mathrm{SL}(m)$ on $\overline{\widehat{\iota}(F^{r,d}_{ss})}$ coincides with the intersection of $\overline{\widehat{\iota}(F^{r,d}_{ss})}$ with $(\widehat{G}(r,m)^N)^s$, the stable locus for the linear action of $\mathrm{SL}(m)$ on $\widehat{G}(r,m)^N$ (see Remark 1.3.2). By [81, Prop 5.7], which relates stability for Higgs bundles to GIT-stability, we obtain that this open stratum of $\overline{\widehat{\iota}(F^{r,d}_{ss})}$, once interpreted as an open stratum of $\mathscr{H}^{r,d}_{ss}(\Sigma, L)$ using the procedure described above, coincides with the stable stratum $\mathscr{H}^{r,d}_{s}(\Sigma, L)$. This provides the desired refinement of the semistable Harder-Narasimhan stratum.

We now turn to the unstable Harder-Narasimhan strata, and fix a Harder-Narasimhan type $\tau \neq \tau_0$. Returning to the diagram (3.3.2) above, we note that $S_{\beta(\tau)} \cong \mathrm{SL}(m) \times_{P_{\beta(\tau)}} Y^{ss}_{\beta(\tau)}$ (see Section 1.1.1), and therefore that $\pi^{-1}(S_{\beta(\tau)}) \cong \mathrm{SL}(m) \times_{P_{\beta(\mu)}} \pi^{-1}(\overline{Y^{ss}_{\beta(\tau)}})$ since π is $\mathrm{SL}(m)$-equivariant. We can apply the stratification theorem of Non-Reductive GIT (namely Theorem 1.3.6) to the linear action of $P_{\beta(\tau)}$ on the closed subvariety $\pi^{-1}(\overline{Y^{ss}_{\beta(\tau)}})$ of $\widehat{G}(r,m)^N$. This application of Theorem 1.3.6 results in a $P_{\beta(\mu)}$-invariant stratification of this closed subvariety such that each stratum admits a quasi-projective geometric quotient with an explicit projective completion.

This stratification induces an $\mathrm{SL}(m)$-invariant stratification of $\pi^{-1}(S_{\beta(\tau)})$, which can be pulled back via $\widehat{\iota}$ to a stratification of $F_\tau^{r,d}$ satisfying the same properties. Under the isomorphism

$$\mathscr{H}_\tau^{r,d}(\Sigma, L) \cong \left[F_\tau^{r,d}/G_d\right],$$

valid for d sufficiently large, we obtain a stratification of $\mathscr{H}_\tau^{r,d}(\Sigma, L)$ satisfying the required properties. By Remark 3.3.4, we can obtain a stratification for a Harder-Narasimhan stratum of any degree d. □

3.3.3 Comparison with the refined Higgs Harder-Narasimhan stratification

Having constructed refinements of the Higgs Harder-Narasimhan and Harder-Narasimhan stratifications which are refined instability stratifications, we can ask the question of how they relate to each other. In particular, we can ask whether the refinement of the Harder-Narasimhan stratification encodes information about the Higgs Harder-Narasimhan type of the Higgs bundle, and conversely whether the refinement of the Higgs Harder-Narasimhan stratification encodes information about the Harder-Narasimhan type of the underlying bundle.

As we do not at present have a complete moduli-theoretic interpretation of either of the two refined stratifications, we cannot fully answer these questions. Nevertheless, we obtain a partial answer in this section. More precisely, we show that the refined Harder-Narasimhan stratification is a refinement of the so-called 'Higgs stratification' of the stack of Higgs bundles (itself a refinement of the Harder-Narasimhan stratification), which encodes information about how the Higgs Harder-Narasimhan and Harder-Narasimhan filtrations of a given Higgs bundle interact. The result we prove is Proposition 3.3.11 below, the statement of which requires first defining the so-called 'Higgs stratification' of the stack of Higgs bundles.

The Higgs stratification as a refinement of the Harder-Narasimhan stratification. Let (E, ϕ) be a Higgs bundle with underlying bundle E of Harder-Narasimhan type $\tau = \left(\frac{d_1}{r_1}, \ldots, \frac{d_s}{r_s}\right)$. Let $0 = E^0 \subseteq E^1 \subseteq \cdots E^s = E$ denote its Harder-Narasimhan filtration and $E_1 \oplus \cdots \oplus E_s$ its associated Harder-Narasimhan graded. Given $i \in \{1, \ldots, s\}$, we let $\pi_i : E \to E/E^i$ denote the quotient map. By taking a tensor product with $\mathrm{id}_L : L \to L$, we obtain a map $E \otimes L \to E/E^i \otimes L$ which we also denote by π_i for simplicity. Given $i, j \in \{1, \ldots, s\}$, consider the map $\pi_{i-1} \circ \phi|_{E^j} : E^j \to E/E^{i-1} \otimes L$. Then if $\pi_{i-1} \circ \phi|_{E^j}$ restricts to the zero map on $E^{j-1} \subseteq E^j$, the map descends to a map $E^j/E^{j-1} \to E/E^{i-1} \otimes L$. If moreover we have that $\pi_i \circ \phi|_{E^j} : E^j \to E/E^i \otimes L$ is the zero map, then the image of $\pi_{i-1} \circ \phi|_{E^j}$ is contained in $E^i/E^{i-1} \otimes L$. Under these two assumptions, we obtain a well-defined map $E^j/E^{j-1} \to E^i/E^{i-1} \otimes L$, denoted $\phi_{i,j}$, such that $\phi_{i,j} \circ \pi_j = \pi_{i-1} \circ \phi|_{E^j}$.

Note that $\phi_{s,1}$ is always well-defined since $\pi_{s-1} \circ \phi|_{E^1} : E^1 \to E/E^{s-1} \otimes L$ trivially induces a map $E_1 \to E_s \otimes L$, given that $E^0 = 0$ and $E = E^s$. Given $i, j \in \{1, \ldots, s\}$, we now show that if $\phi_{j-1,i}$ and $\phi_{j,i+1}$ are well-defined and equal to 0, then $\phi_{i,j}$ is also well-defined. Since $\phi_{j-1,i} = 0$, it follows that $\pi_{i-1} \circ \phi|_{E^j}$ is zero on $E^{j-1} \subseteq E^j$ and so $\pi_{i-1} \circ \phi|_{E^j}$ descends to a map $E^j/E^{j-1} \to E/E^{i-1} \otimes L$. To see that the image is contained in $E^i/E^{i-1} \otimes L$, we use the fact that $\phi_{i,j+1} = 0$. Indeed this implies that $\pi_i \circ \phi|_{E^j} = 0$ and hence $\phi(E^j) \subseteq E^i \otimes L$. Thus $\operatorname{Im} \pi_{i-1} \circ \phi|_{E^j} \subseteq E^i/E^{i-1} \otimes L$ and so the map induces a map $\phi_{i,j} : E_j \to E_i \otimes L$ as required.

We can define a refinement of the Harder-Narasimhan stratification as follows. Given a Harder-Narasimhan type $\tau = \left(\frac{d_1}{r_1}, \ldots, \frac{d_s}{r_s}\right)$ and $i, j \in \{1, \ldots, s\}$, we define the substack

$$\mathscr{H}^{r,d}_{\tau,i,j}(\Sigma, L) \subseteq \mathscr{H}^{r,d}_{\tau}(\Sigma, L)$$

to be the substack of Higgs bundles (E, ϕ) such that

(i) $\phi_{i',j'}$ is equal to 0 if either $i' > i$ and $j' \leq j$ or $i' \geq i$ and $j' < j$;

(ii) $\phi_{i,j} \neq 0$.

Since $\phi_{s,1}$ is well-defined, the condition that $\phi_{s,1} \neq 0$ is valid. From there, by descending induction on i and ascending induction on j we can see that the above two conditions are valid.

Remark 3.3.9. There is an ordering on the pairs (i, j) given by $(i, j) \leq (i', j')$ if $\frac{d_i}{r_i} - \frac{d_j}{r_j} \leq \frac{d_{i'}}{r_{i'}} - \frac{d_{j'}}{r_{j'}}$.

We note that if $[(E, \phi)] \in \mathscr{H}^{r,d}_{\tau,i,j}(\Sigma, L)$ for some (i, j) satisfying $i > j$, then (E, ϕ) cannot have the same Higgs Harder-Narasimhan filtration as the Harder-Narasimhan filtration of E. Indeed, if $\phi_{i,j} : E_j \to E_i \otimes L$ is well-defined and not equal to 0, it follows that $\phi(E^j) \nsubseteq E^{i-1} \otimes L$, and thus E^j is not a ϕ-invariant subbundle since $E^j \subseteq E^{i-1}$ based on the assumption that $i > j$. If this is not the case, namely if (E, ϕ) with E of Harder-Narasimhan type τ does not lie in $\mathscr{H}^{r,d}_{\tau,i,j}(\Sigma, L)$ for any pair (i, j) satisfying $i > j$, then $\phi_{i,j}$ is well-defined and equal to 0 for every $i > j$. This implies that $\phi(E^i) \subseteq E^i \otimes L$ for every i, and thus ϕ preserves the Harder-Narasimhan filtration of E. By Proposition 3.1.6 (iii), we know that the Higgs Harder-Narasimhan and Harder-Narasimhan filtrations of a Higgs bundle coincide if and only if its Higgs Harder-Narasimhan type and Harder-Narasimhan types are equal. Therefore the substack of $\mathscr{H}_\tau^{r,d}(\Sigma, L)$ consisting of Higgs bundles (E, ϕ) for which the Higgs Harder-Narasimhan and Harder-Narasimhan filtrations coincide corresponds to the intersection $\mathscr{H}_\tau^{r,d}(\Sigma, L) \cap \mathscr{H}_{r,d}^\tau(\Sigma, L)$.

We now define the Higgs stratification of $\mathscr{H}_{r,d}(\Sigma, L)$.

Definition 3.3.10. Let $\tau \neq \tau_0$ denote a Harder-Narasimhan type of rank r and degree d. The *Higgs stratification of the Harder-Narasimhan stratum* $\mathscr{H}_{r,d}^\tau(\Sigma, L)$ is given by

$$\mathscr{H}_{r,d}^\tau(\Sigma, L) = \bigsqcup_{1 \leq j < i \leq s} \mathscr{H}_{\tau,i,j}^{r,d}(\Sigma, L) \sqcup \left(\mathscr{H}_\tau^{r,d}(\Sigma, L) \cap \mathscr{H}_{r,d}^\tau(\Sigma, L)\right). \tag{3.11}$$

The *Higgs stratification of* $\mathscr{H}_{ss}^{r,d}(\Sigma, L)$ is the trivial one. The induced stratification of $\mathscr{H}_{r,d}(\Sigma, L)$ is the *Higgs stratification of* $\mathscr{H}_{r,d}(\Sigma, L)$.

By construction, the Higgs stratification of $\mathscr{H}_{r,d}(\Sigma, L)$ is a refinement of the Harder-Narasimhan stratification. In Proposition 3.3.11 below, we show that the refined Harder-Narasimhan stratification constructed to prove Theorem 3.3.5 is a refinement of the Higgs stratification.

Proposition 3.3.11. The refined Harder-Narasimhan stratification of $\mathscr{H}_{r,d}(\Sigma, L)$ constructed in the proof of Theorem 3.3.5 is a refinement of the Higgs stratification of $\mathscr{H}_{r,d}(\Sigma, L)$.

Proof. Since the Higgs stratification does not stratify the open stratum $\mathscr{H}_{ss}^{r,d}(\Sigma, L)$, it suffices to show that the result is true for each Harder-Narasimhan stratum $\mathscr{H}_\tau^{r,d}(\Sigma, L)$ where $\tau \neq \tau_0$. By definition, given such a Harder-Narasimhan stratum, its Higgs stratification is given by

$$\mathscr{H}_\tau^{r,d}(\Sigma, L) = \bigsqcup_{1 \leq j < i \leq s} \mathscr{H}_{\tau,i,j}^{r,d}(\Sigma, L) \sqcup \left(\mathscr{H}_\tau^{r,d}(\Sigma, L) \cap \mathscr{H}_{r,d}^\tau(\Sigma, L)\right). \tag{3.12}$$

We recall from the proof of Theorem 3.3.5 (and using the notation introduced there) that the refinement of the stratum $\mathscr{H}_\tau^{r,d}(\Sigma, L)$ given by the refined Harder-Narasimhan stratification is constructed from a non-reductive GIT stratification associated to the linear action of $P_{\beta(\tau)}$ on the closed subvariety $\pi^{-1}(\overline{Y_{\beta(\tau)}^{ss}})$ of $\widehat{G}(r, m)^N$. By Theorem 1.3.6 (iii), we know that any such non-reductive GIT stratification refines the Bialynicki-Birula stratification associated to the action of $\lambda_{\beta(\tau)}(\mathbb{G}_m)$ on $\pi^{-1}(\overline{Y_{\beta(\tau)}^{ss}})$. Just as we showed that a non-reductive GIT stratification of $\pi^{-1}(\overline{Y_{\beta(\tau)}^{ss}})$ induces a stratification of $\mathscr{H}_\tau^{r,d}(\Sigma, L)$, the coarse Bialynicki-Birula stratification of $\pi^{-1}(\overline{Y_{\beta(\tau)}^{ss}})$ similarly induces a stratification of $\mathscr{H}_\tau^{r,d}(\Sigma, L)$. We claim that this stratification coincides with the Higgs stratification. It will then follow immediately that the refined Harder-Narasimhan stratification constructed in the proof of Theorem 3.3.5 is a refinement of the Higgs stratification. Thus it suffices to prove this claim.

To simplify notation, we let $\beta = \beta(\tau)$. We write $\tau = \left(\frac{d_1}{r_1}, \ldots, \frac{d_1}{r_1}, \frac{d_2}{r_2}, \ldots, \frac{d_s}{r_s}\right)$, where each entry $\frac{d_i}{r_i}$ is repeated r_i times. Then under the correspondence $\tau \mapsto \beta(\tau)$ of [65, Def 11.1], the vector $\beta = \beta(\tau)$ is given by $\left(\frac{k_1}{m_1}, \ldots, \frac{k_1}{m_1}, \frac{k_2}{m_2}, \ldots, \frac{k_s}{m_s}\right) - \left(\frac{k}{m}, \ldots, \frac{k}{m}\right)$, where each $\frac{k_i}{m_i}$ appears m_i times, for $k_i = N(d_i - r_i g)$ and $m_i = d_i + r_i(1-g)$. The one-parameter subgroup $\lambda_{\beta(\tau)}(\mathbb{G}_m) \subseteq \mathrm{SL}(m)$ consists of diagonal matrices with the vector $\beta(\tau)$ on the diagonal. We note that the condition that $\frac{d_1}{r_1} > \frac{d_2}{r_2} > \cdots > \frac{d_s}{r_s}$ implies that $\frac{k_1}{m_1} > \frac{k_2}{m_2} > \cdots > \frac{k_s}{m_s}$.

We start by describing the minimal weight space (namely the analogue of $Z_{\min}$ for $\pi^{-1}(\overline{Y_\beta^{ss}})$) for the action of $\lambda_\beta(\mathbb{G}_m)$ on $\pi^{-1}(\overline{Y_\beta^{ss}})$. We recall from Nitsure's set-up that $\pi^{-1}(\overline{Y_\beta^{ss}}) \subseteq \widehat{G}(r,m)^N$, where $\pi : \widehat{G}(r,m)^N \to G(r,m)^N$ is a projective bundle over $G(r,m)^N$. Following [81, §2], a point in $\widehat{G}(r,m)$ can be written as an equivalence class $\langle y, [c:\phi]\rangle$ where y is an $r \times m$ matrix, $c \in k$, $\phi \in k^{r\times r}$ and c and ϕ are not both zero, under the equivalences $\langle y, [c:\phi]\rangle = \langle y, [\beta c : \beta\phi]\rangle$ for any $\beta \in k^*$ and $\langle y, [c:\phi]\rangle = \langle \alpha y, [(\det\alpha)^{-1}c : (\det\alpha)^{-1}\alpha\phi\alpha^{-1}]\rangle$ for any $\alpha \in \mathrm{GL}(r)$.

For a point $\widehat{y} = (\langle y_1, [c_1:\phi_1]\rangle, \ldots, y_N, [c_N:\phi_N]\rangle) \in \pi^{-1}(\overline{Y_{\beta(\tau)}^{ss}})$ to be in the minimal weight space for $\pi^{-1}(\overline{Y_{\beta(\tau)}^{ss}})$, it must be fixed by $\lambda(\mathbb{G}_m)$. Thus in particular its projection to $G(r,m)^N$ under π must be fixed by $\lambda(\mathbb{G}_m)$, since π is $\mathrm{SL}(m)$-equivariant. But the image of $\pi^{-1}(\overline{Y_\beta^{ss}})$ in $G(r,m)^N$ coincides with $\overline{Y_\beta}^{ss}$, and the only points in $\overline{Y_\beta}^{ss}$ fixed by $\lambda_\beta(\mathbb{G}_m)$ are those in Z_β (this follows from the results of Section 1.1.3 of Chapter 1 regarding the application of Non-Reductive GIT to GIT-unstable strata). We note that by [64, §16.9], the intersection of Z_β with the image of $R^\tau_{r,d}$ under ι coincides with vector bundles of Harder-Narasimhan type τ which are isomorphic to their Harder-Narasimhan graded.

Thus to determine the minimal weight space in $\pi^{-1}(\overline{Y_\beta^{ss}})$, it suffices to determine the minimal weight space for the action of $\lambda_\beta(\mathbb{G}_m)$ on the fibre $\pi^{-1}(z)$ for any $z \in Z_\beta$. We fix such a point $z = (\langle z_1\rangle, \ldots, \langle z_N\rangle)$. Then any point $\widehat{z} \in \pi^{-1}(z)$ is of the form $\widehat{z} = (\langle z_1, [c_1:\phi_1]\rangle, \cdots, \langle z_N, [c_N:\phi_N]\rangle)$. By definition of the action of $\mathrm{SL}(m)$ on $\widehat{G}(r,m)^N$ (see [81, §2]), the one-parameter subgroup $\lambda(\mathbb{G}_m)$ acts trivially on the coordinates c_i, and via conjugation on the matrices ϕ_i. The weights of this action are therefore given by

$$\{0\} \cup \left\{ \frac{k_j}{m_j} - \frac{k_i}{m_i} \;\middle|\; i,j \in \{1,\dots,s\}, i \neq j \right\},$$

and since $\frac{k_1}{m_1} > \frac{k_2}{m_2} > \cdots > \frac{k_s}{m_s}$, it follows that the minimal weight is $\frac{k_s}{m_s} - \frac{k_1}{m_1}$. Given $\phi \in \mathrm{Mat}_{r\times r}(k)$, we can consider for each $i,j \in \{1,\dots,s\}$ the block $\phi_{i,j} \in \mathrm{Mat}_{r_i \times r_j}(k)$. Matrices ϕ which are zero everywhere except for $\phi_{s,1}$ are weight vectors for the conjugation action of $\lambda(\mathbb{G}_m)$ on $\mathrm{Mat}_{r\times r}(k)$ with weight $\frac{k_s}{m_s} - \frac{k_1}{m_1}$. Thus the minimal weight space for the action of $\lambda_\beta(\mathbb{G}_m)$ on $\pi^{-1}(\overline{Y_\beta^{ss}})$ consists of points $(\langle z_1, [0{:}\phi_1]\rangle, \cdots, \langle z_N, [0{:}\phi_N]\rangle) \in \pi^{-1}(\overline{Y_\beta^{ss}})$ such that $(\langle z_1\rangle, \dots, \langle z_N\rangle) \in Z_\beta$ and such that for each $k = 1, \dots, N$ the equality $(\phi_k)_{ij} = 0$ holds for $i,j \in \{1,\dots,s\}$ except when $i = s$ and $j = 1$ (note that this forces $(\phi_k)_{s1} \neq 0$). It follows that the open stratum of the Bialynicki-Birula stratification corresponding to this minimal weight space consists of points $(\langle y_1, [0{:}\phi_1]\rangle, \cdots, \langle y_N, [0{:}\phi_N]\rangle) \in \pi^{-1}(\overline{Y_\beta^{ss}})$ such that $(\langle y_1\rangle, \dots, \langle y_N\rangle) \in Y_\beta^{ss}$ and such that $(\phi_k)_{s1} \neq 0$ for each $k = 1, \dots, N$.

The intersection of this open stratum with the image of $F_\tau^{r,d}$ under $\widehat{\iota}$ can be interpreted in the following moduli-theoretic way: it consists of Higgs bundles (E, ϕ) with Harder-Narasimhan type τ such that $\phi_{1,s} \neq 0$. Thus the resulting open stratum of $\mathscr{H}_\tau^{r,d}(\Sigma, L)$ coincides with the open stratum $\mathscr{H}_{\tau,s,1}^{r,d}(\Sigma, L)$ of the Higgs stratification of $\mathscr{H}_\tau^{r,d}(\Sigma, L)$.

We can similarly describe the Bialynicki-Birula strata for all weights $\frac{k_j}{m_j} - \frac{k_i}{m_i} < 0$, and we obtain that they coincide with the strata $\mathscr{H}_{\tau,i,j}(\Sigma, L)$). It remains only to relate the stratum $\mathscr{H}_\tau^{2,d}(\Sigma, L) \cap \mathscr{H}_{2,d}^{\tau}(\Sigma, L)$ of the Higgs stratification to a Bialynicki-Birula stratum, and we now show that it is the stratum corresponding to weight 0. Indeed, in the 0-weight space for the action of $\lambda_\beta(\mathbb{G}_m)$ on $\pi^{-1}(\overline{Y_\beta^{ss}})$, points $\widehat{y} = (\langle y_1, [c_1{:}\phi_1]\rangle, \cdots, \langle y_N, [c_N{:}\phi_N]\rangle) \in \pi^{-1}(\overline{Y_\beta^{ss}})$ in the corresponding Bialynicki-Birula stratum satisfy the property that each ϕ_i is a block upper triangular matrix (to ensure that there is no non-zero coordinate with smaller weight). Thus if a Higgs bundle lies in this Bialynicki-Birula stratum, then its Higgs field must preserve the Harder-Narasimhan filtration (this follows from the fact that each ϕ_i for $i = 1, \dots, N$ is block upper triangular), so that its Higgs Harder-Narasimhan and Harder-Narasimhan types coincide. Conversely, if a Higgs bundle satisfies this property, then its image in $\pi^{-1}(\overline{Y_\beta^{ss}})$ will have zero weight simply based on the fact that it lies in $F_{r,d}$. Indeed, there are inclusions $F_\tau^{r,d} \subseteq H(r,m)^N \subseteq \widehat{G}(r,m)^N$ where $H(r,m)^N$ consists of those points $(\langle y_1, [c_1{:}\phi_1]\rangle, \cdots, \langle y_N, [c_N{:}\phi_N]\rangle) \in \widehat{G}(r,m)^N$

such that all coordinates c_i are non-zero, and these coordinates have weight 0 for the action of $\lambda_\beta(\mathbb{G}_m)$. Thus the intersection with $F_\tau^{r,d}$ of the Bialnyicki-Birula stratum associated to the 0-weight space for the action of λ_β on $\pi^{-1}(\overline{Y_\beta^{ss}})$ corresponds inside the Harder-Narasimhan stratum $\mathscr{H}_\tau^{2,d}(\Sigma, L)$ to the intersection $\mathscr{H}_\tau^{2,d}(\Sigma, L) \cap \mathscr{H}_{2,d}^\tau(\Sigma, L)$. □

Remark 3.3.12 (The rank 2 and odd degree case)**.** In the rank 2 case, the Higgs stratification is simple to describe. Given an unstable Harder-Narasimhan τ of rank 2 and odd degree d, the Higgs stratification of $\mathscr{H}_\tau^{2,d}(\Sigma, L)$ is given by

$$\mathscr{H}_\tau^{2,d}(\Sigma, L) = \mathscr{H}_{\tau,2,1}^{2,d}(\Sigma, L) \sqcup (\mathscr{H}_\tau^{2,d}(\Sigma, L) \cap \mathscr{H}_{2,d}^\tau(\Sigma, L)) = \mathscr{H}_{\tau,2,1}^{2,d}(\Sigma, L) \sqcup \mathscr{H}_{2,d}^\tau(\Sigma, L). \quad (3.13)$$

The latter equality follows from Proposition 3.1.8, which implies that when $\tau \neq \tau_0$ there is an inclusion $\mathscr{H}_{2,d}^\tau(\Sigma, L) \subseteq \mathscr{H}_\tau^{2,d}(\Sigma, L)$ (since an unstable Higgs bundle of rank 2 must have an underlying vector bundle of the same instability type). By Proposition 3.3.11, we know that the restriction of the refined Harder-Narasimhan stratification to $\mathscr{H}_\tau^{2,d}(\Sigma, L)$ is a refinement of (3.13). Thus in particular we obtain a stratification of $\mathscr{H}_\tau^{2,d}(\Sigma, L) \cap \mathscr{H}_{2,d}^\tau(\Sigma, L)$. We have shown in [45, Thm 4.1 (ii)][21] that this stratification coincides with the intersection of $\mathscr{H}_\tau^{2,d}(\Sigma, L)$ with the refined Higgs Harder-Narasimhan stratification, described explicitly in Theorem 3.2.9.

Thus it remains only to consider the restriction of the refined Harder-Narasimhan stratification to the stratum $\mathscr{H}_{\tau,2,1}^{2,d}(\Sigma, L)$. We recall that this stratum parametrises Higgs bundles $[(E, \phi)]$ of Higgs Harder-Narasimhan type τ, such that the induced Higgs field $\phi_{2,1} : E^1 \to E/E^1 \otimes L$ is non-zero, where E^1 is the canonically destabilising subbundle in the Harder-Narasimhan filtration of E. Thus in particular E^1 is not ϕ-invariant, and so (E, ϕ) cannot have Higgs Harder-Narasimhan type τ. It follows from Theorem 3.2.9 that (E, ϕ) must be semistable. Therefore $\mathscr{H}_{\tau,2,1}^{2,d}(\Sigma, L)$ coincides with the intersection $\mathscr{H}_{2,d}^{ss}(\Sigma, L) \cap \mathscr{H}_\tau^{2,d}(\Sigma, L)$ of the Higgs semistable stratum with the Harder-Narasimhan stratum $\mathscr{H}_\tau^{2,d}(\Sigma, L)$.

We note that the locally closed subset of the moduli space $\mathcal{M}_{2,d}^{ss}(\Sigma, L)$ given by Higgs bundles with underlying bundle of Harder-Narasimhan type τ provides a quasi-projective coarse moduli

[21] We were able to obtain this result in [45] because we used the same set-up for constructing both the refined Higgs Harder-Narasimhan and Harder-Narasimhan stratifications. In this thesis we have used a different set-up to that used in [45] to construct the refined Higgs Harder-Narasimhan stratification (namely we have used in this thesis the spectral correspondence rather than Nitsure's construction), and for this reason it is not immediately obvious that the restrictions of the refined Higgs Harder-Narasimhan and Harder-Narasimhan stratifications to an intersection of the form $\mathscr{H}_\tau^{r,d}(\Sigma, L) \cap \mathscr{H}_{r,d}^\tau(\Sigma, L)$ coincide. Nevertheless we expect this to be the case, and that it can be shown through a more detailed analysis of the Non-Reductive GIT constructions of both refinements.

space for this stratum. In fact, it can be shown by explicitly studying the linear action of $P_{\beta(\tau)}$ on $\pi^{-1}(\overline{Y^{ss}_{\beta(\tau)}})$ (which is the set-up used in the proof of Theorem 3.3.5 to construct the refinement of the Harder-Narasimhan stratum $\mathscr{H}^{2,d}_{\tau}(\Sigma, L)$) that the stratum $\mathscr{H}_{\tau,2,1}(\Sigma, L)$ does not need to be further refined: under its identification as points in the parameter space $\pi^{-1}(\overline{Y^{ss}_{\beta(\tau)}})$, it is contained in the 'stable' locus for the linear action of $P_{\beta(\tau)}$ on $\pi^{-1}(\overline{Y^{ss}_{\beta(\tau)}})$ obtained as the output of the Projective Completion algorithm (which in this situation does not involve any choices)[22].

As a result we can obtain, just as in the case of the refined Higgs Harder-Narasimhan stratification, an explicit description of the refined Harder-Narasimhan stratification constructed in the proof of Theorem 3.3.5:

$$\mathscr{H}_{2,d}(\Sigma, L) = \mathscr{H}^{2,d}_{ss}(\Sigma, L) \sqcup \bigsqcup_{\tau} \left(\left(\mathscr{H}^{ss}_{2,d}(\Sigma, L) \cap \mathscr{H}^{2,d}_{\tau}(\Sigma, L) \right) \sqcup \bigsqcup_{\delta \in \mathbb{N}} \mathscr{H}^{\tau,\delta,\mathrm{indec}}_{2,d}(\Sigma, L) \sqcup \mathscr{H}^{\tau,\mathrm{dec}}_{2,d}(\Sigma, L) \right).$$

In particular, we can give an explicit characterisation of τ-stability (introduced in Definition 3.3.6), as per Proposition 3.3.13 below.

Proposition 3.3.13 (τ-stable Higgs bundles of rank 2)**.** Let τ denote a Harder-Narasimhan type of rank 2 and degree d. A Higgs bundle (E, ϕ) with an underlying bundle E of Harder-Narasimhan type τ is *τ-stable* if and only if it is stable as a Higgs bundle.

[22] We have chosen not to include the proof of this result in this thesis because it is quite technical (see [45, Thm 4.11]), and because the existence of a quasi-projective coarse moduli space for the stratum $\mathscr{H}^{2,d}_{\tau,2,1}(\Sigma, L) = \mathscr{H}^{ss}_{2,d}(\Sigma, L) \cap \mathscr{H}^{2,d}_{\tau}(\Sigma, L)$ which it establishes can be deduced without using Non-Reductive GIT.

Chapter 4

Geometry of moduli spaces for unstable Higgs bundles of rank 2

4.0 Introduction

Refined instability stratifications for the stack of Higgs bundles. In Chapter 3 we constructed two refined instability stratifications of the stack of Higgs bundles. The first is a refinement of the Higgs Harder-Narasimhan stratification, determined by the instability type of the Higgs bundles, while the second is a refinement of the Harder-Narasimhan stratification, determined by the instability type of the underlying vector bundle. The refined stratifications are constructed using Non-Reductive GIT, and satisfy the property that each stratum admits a quasi-projective coarse moduli space. In this way their construction extends the well-known result regarding the existence of a quasi-projective coarse moduli space for the substack of stable Higgs bundles. We are thus naturally led to the next step, namely the study of the geometry of these new moduli spaces. In order to do this, it is helpful to be able to identify the objects which these moduli spaces parametrise. In other words, a complete moduli-theoretic interpretation of the refined instability stratifications is desired.

The rank 2 case. Although obtaining such an interpretation of the refined instability stratifications in general is still work in progress, we have done so in the rank 2 case in Chapter 3. As we have seen in Chapter 3, the refined Harder-Narasimhan stratification is simply the intersection of the Harder-Narasimhan stratification with the refined Higgs Harder-Narasimhan stratification. The coarse moduli spaces for the refined Higgs Harder-Narasimhan strata associated to a given Higgs Harder-Narasimhan stratum $\mathscr{H}^{\mu}_{2,d}(\Sigma, L)$, where $\mu = (d_1, d_2)$, can be split into three classes:

(1) The first class consists of a single moduli space: the moduli space $\mathcal{M}^{\mu\text{-s}}_{2,d}(\Sigma, L)$ parametrising isomorphism classes of μ-stable Higgs bundles of Higgs Harder-Narasimhan type μ. As per Proposition 3.2.14, these are Higgs bundles (E, ϕ) of Higgs Harder-Narasimhan type μ which are not isomorphic to their graded, and which satisfy $\dim \mathrm{End}_{-1}(E, \phi) = 0$. We note that $\mathcal{M}^{\mu\text{-s}}_{2,d}(\Sigma, L) = \mathcal{M}^{\mu,0,\mathrm{indec}}_{2,d}(\Sigma, L)$, and that it is (provided it is non-empty) the coarse moduli space for the open stratum of the refinement of the Higgs Harder-Narasimhan stratum. This is indeed the reason behind the terminology 'μ-stable', chosen by analogy with the semistable case in which the notion of stability determines an open substack of the semistable stratum admitting a quasi-projective coarse moduli space.

(2) The second class consists of a finite set of moduli spaces of the form $\mathcal{M}^{\mu,\delta,\mathrm{indec}}_{2,d}(\Sigma, L)$ for $\delta > 0$, which parametrise isomorphism classes of μ-unstable Higgs bundles. As per Definition 3.2.15, these are Higgs bundles (E, ϕ) of Higgs Harder-Narasimhan type μ which are not isomorphic to their graded, and which satisfy $\dim \mathrm{End}_{-1}(E, \phi) = \delta$. These are coarse moduli spaces for the locally closed strata of the refinement of the Higgs Harder-Narasimhan stratum.

(3) The third class consists of a single moduli space: the moduli space $\mathcal{M}^{\mu,\mathrm{dec}}_{2,d}(\Sigma, L)$ of Higgs bundles of Higgs Harder-Narasimhan type μ and which are isomorphic to their graded. It is the coarse moduli space for the closed stratum of the refinement of the Higgs Harder-Narasimhan stratum. This moduli space is isomorphic to the product $\mathcal{M}^{ss}_{1,d_1}(\Sigma, L) \times \mathcal{M}^{ss}_{1,d_2}(\Sigma, L)$ of moduli spaces of rank 1 Higgs bundles of degree d_1 and d_2 respectively, which are well-understood (see [34] for example).

Thus the new moduli spaces are those of the form $\mathcal{M}^{\mu,\delta,\mathrm{indec}}_{2,d}(\Sigma, L)$ for $\delta \geq 0$, and we can ask whether their geometry has a similarly rich structure to that of the moduli space of stable Higgs bundles. This is the guiding question of this chapter.

When $\Sigma = \mathbb{P}^1$. Although classical Higgs bundles (namely when $L = K_\Sigma$) are generally considered over curves of genus greater than one to obtain meaningful results[1], when considering the more general case of twisted Higgs bundles, even the genus zero case is interesting. For example, although the moduli space of stable L-twisted Higgs bundles over a curve of genus zero is no longer hyperkähler, it still has the rich structure of an integrable system (see for example [5]), just as in the classical case.

The appeal of working over $\mathbb{P}^1$ in the rank 2 case is that there are no moduli for the underlying bundle, and as a result the moduli spaces can be described explicitly. This was achieved in the semistable case by Rayan in [85], which provides an explicit global description of the moduli space of rank 2 semistable Higgs bundles twisted by the dual of the canonical line bundle on $\mathbb{P}^1$ (these objects are called co-Higgs bundles). Interestingly, many features of the moduli space of stable Higgs bundles in the classical case on curves of genus greater than one are also exhibited by the moduli space of stable L-twisted Higgs bundles on $\mathbb{P}^1$ (for example the degree independence of the Betti numbers of the moduli space, see [87] and the references therein). As a result, considering unstable L-twisted Higgs bundles on $\mathbb{P}^1$ may be viewed as a useful 'test case' for understanding unstable Higgs bundles in the classical sense on curves of higher genus.

For this reason we will start our study of the moduli spaces of unstable Higgs bundles of rank 2 by specialising to the case of Higgs bundles on $\mathbb{P}^1$, for which we can obtain an explicit description of the moduli spaces. Even in this simple case, it is interesting to compare the geometry of the moduli spaces to that of the moduli space of stable Higgs bundles. Moreover, the descriptions of the moduli spaces give an indication of the features we might expect to see exhibited by the moduli spaces of unstable Higgs bundles of rank 2 on curves of higher genus.

The $\mathbb{C}^*$-action. An important feature of the moduli space of stable L-twisted Higgs bundles, on curves of any genus, is the Higgs field scaling $\mathbb{C}^*$-action. Although the moduli space of stable L-twisted Higgs bundles is only quasi-projective, it is semiprojective in the sense that limits at

[1] Indeed, there are no stable Higgs bundles, in the classical sense, over $\mathbb{P}^1$. Nevertheless, the genus zero case for classical Higgs bundles still represents an active area of research, through the study of parabolic Higgs bundles. These are Higgs bundles equipped with an additional so-called parabolic structure resulting from puncturing $\mathbb{P}^1$ and imposing conditions on the behaviour of the Higgs field at the punctures. The resulting moduli space of parabolic Higgs bundles has the advantage of being non-empty and hyperkähler [14].

0 under the $\mathbb{C}^*$-action always exist, and that the fixed point set for this action is projective (see [81, Thm. 6.1]). These two conditions are exactly those required for a quasi-projective variety admitting a $\mathbb{C}^*$-action to be semiprojective, as defined by Hausel in [47]. The semiprojectivity of the moduli space of semistable Higgs bundles can be exploited to describe the topology and geometry of the moduli space. For example, the resulting Bialynicki-Birula stratification allows the computation of the cohomology of $\mathcal{M}^{ss}_{r,d}(\Sigma, L)$ from the cohomology of the fixed point set, and moreover an orbifold compactification of the moduli space can be obtained from the $\mathbb{C}^*$-action via an algebraic analogue of Lerman's symplectic cutting (see [47] for more details).

The Higgs field scaling $\mathbb{C}^*$-action extends naturally to the whole stack $\mathscr{H}_{r,d}(\Sigma, L)$ of Higgs bundles, and therefore the moduli spaces of unstable Higgs bundles which we have constructed also admit a $\mathbb{C}^*$-action. A follow-up question then is whether this action can similarly be used to study the geometry of the moduli spaces. We will answer this question in the rank 2 case.

Structure of this chapter. The aim of this chapter is to study the geometry, and in particular the Poincaré series, of the moduli spaces $\mathcal{M}^{\mu,\delta,\mathrm{indec}}_{2,d}(\Sigma, L)$ of unstable Higgs bundles of rank 2 which we have constructed in Chapter 3. In Section 4.1.1 we consider the case where $\Sigma = \mathbb{P}^1$. Doing so allows us to explicitly describe the geometry of the moduli spaces thanks to two maps: the map taking a Higgs bundle to its associated graded and the Hitchin morphism. In Sections 4.2 and 4.3 we turn to curves of arbitrary genus. Section 4.2 studies the geometry of the moduli spaces $\mathcal{M}^{\mu,\delta,\mathrm{indec}}_{2,d}(\Sigma, L)$ when $\delta > 0$ using the $\mathbb{C}^*$-action: we show that limits at 0 always exist, and moreover how the resulting Bialynicki-Birula stratification can be used to compute the Poincaré series of the moduli space in the case where $L = K_\Sigma$. Finally in Section 4.3 we consider the moduli space $\mathcal{M}^{\mu\text{-s}}_{2,d}(\Sigma, L)$. As we will see, limits at 0 under the $\mathbb{C}^*$-action may no longer exist, so a different strategy to that used in the $\delta > 0$ case is required for computing the Poincaré series of the moduli space. Instead, we use the fact that the moduli space can be constructed using Non-Reductive GIT to obtain a partial compactification of the space. We then show how the cohomological formulae of Chapter 2 can be applied to compute the Poincaré series of this partial compactification.

4.1 Moduli spaces for unstable twisted Higgs bundles of rank 2 on $\mathbb{P}^1$

In Chapter 3 (see Theorem 3.2.9), we obtained a complete moduli-theoretic description of the refined Higgs Harder-Narasimhan of the stack of L-twisted Higgs bundles of rank 2 and odd degree on a smooth projective curve of arbitrary genus. The aim of this section is to describe the geometry and topology of the refined strata and of their corresponding coarse moduli spaces in the case where the smooth projective curve has genus zero[2]; we will consider the arbitrary genus case in the Sections 4.2 and 4.3.

For the remainder of this section we therefore fix $\Sigma = \mathbb{P}^1$ and $L = \mathcal{O}(t)$ where $t \geq 0$. To simplify notation, we omit this data when referring to the stack $\mathscr{H}_{2,d}(\mathbb{P}^1, \mathcal{O}(t))$, writing $\mathscr{H}_{2,d}$ instead, and similarly for the moduli spaces.

The advantage of the genus 0 case is that, by Grothendieck's splitting theorem[3], vector bundles over $\mathbb{P}^1$ have no continuous moduli and thus the moduli problem for Higgs bundles of rank 2 essentially reduces to the moduli problem for Higgs fields. That is, two Higgs fields for the same vector bundle E define isomorphic Higgs bundles if and only if there exists an automorphism ψ of E such that the equality $\psi \circ \phi \circ \psi^{-1} = \phi$ holds[4]. As a result, the moduli problem for Higgs bundles on $\mathbb{P}^1$ can be reduced to the problem of understanding the orbits for the conjugation action of automorphism groups of rank 2 bundles on $\mathbb{P}^1$ on Higgs fields for such bundles. Moreover, since vector bundles on $\mathbb{P}^1$ always split, the Higgs fields can be written in matrix form which simplifies their study. In the rank 2 case under consideration, this means that any $\mathcal{O}(t)$-twisted Higgs bundle on $\mathbb{P}^1$ of degree d can be written in the form

$$(E, \phi) = \left(\mathcal{O}(d_1) \oplus \mathcal{O}(d_2), \begin{pmatrix} a & b \\ c & d \end{pmatrix} \right)$$

where

(i) $d_1 + d_2 = d$;

(ii) $a, d \in \mathrm{Hom}(\mathcal{O}(d_1), \mathcal{O}(d_1) \otimes \mathcal{O}(t)) \cong \mathrm{Hom}(\mathcal{O}(d_2), \mathcal{O}(d_2) \otimes \mathcal{O}(t)) \cong H^0(\mathbb{P}^1, \mathcal{O}(t))$;

(iii) $c \in \mathrm{Hom}(\mathcal{O}(d_1), \mathcal{O}(d_2) \otimes \mathcal{O}(t)) \cong H^0(\mathbb{P}^1, \mathcal{O}(d_2 - d_1 + t))$; and

[2] I am grateful to Steven Rayan for suggesting that I consider unstable L-twisted Higgs bundles on $\mathbb{P}^1$, as a first step towards understanding the geometry of moduli spaces of unstable Higgs bundles.

[3] The splitting theorem (see [38]) states that any rank r vector bundle on $\mathbb{P}^1$ splits as a direct sum $\mathcal{O}(a_1) \oplus \cdots \oplus \mathcal{O}(a_r)$ and that the integers $a_1, \ldots, a_r$ are unique up to reordering.

[4] We are writing $\psi \circ \phi \circ \psi^{-1}$ for simplicity; what we really mean is $(\psi \otimes \mathrm{id}_L) \circ \phi \circ \psi^{-1}$ since ϕ maps E to $E \otimes L$, and therefore ψ must be interpreted as the automorphism $\psi \otimes \mathrm{id}_L$ of $E \otimes L$.

(iv) $b \in \mathrm{Hom}(\mathcal{O}(d_2), \mathcal{O}(d_1) \otimes \mathcal{O}(t)) \cong H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + t))$.

In Section 4.1.1 we explicitly describe the refined Higgs Harder-Narasimhan strata for the stack of rank 2 and odd degree $\mathcal{O}(t)$-twisted Higgs bundles on $\mathbb{P}^1$. This allows us to compute the dimensions of each of the refined strata. In Section 4.1.2 we study the geometry of the coarse moduli spaces of the unstable refined Higgs Harder-Narasimha strata. We provide an explicit description of the moduli spaces, which allows us to prove that they are smooth and to compute their dimensions. We also describe the image and fibre of the Hitchin morphism for each of these moduli spaces in the trace-free case. Finally, in Section 4.1.3 we compare the results obtained to known results from the semistable case relating to dimensions and the Hitchin fibration.

4.1.1 The stack and its refined Higgs Harder-Narasimhan strata

In this section we apply Theorem 3.2.9 to the case where $\Sigma = \mathbb{P}^1$ in order to obtain an explicit description of each of the Higgs Harder-Narasimhan strata appearing in the refined Higgs Harder-Narasimhan stratification. That is, we establish which of the strata are non-empty, and give an explicit description of the Higgs bundles which the non-empty strata parametrise (see Proposition 4.1.2). We then use this description to compute the dimension of the refined Higgs Harder-Narasimhan strata (see Proposition 4.1.5).

We start by describing the refined Higgs Harder-Narasimhan stratification for the stack of rank 2 and odd degree $\mathcal{O}(t)$-twisted Higgs bundles on $\mathbb{P}^1$. The statement of the result relies on the following

Notation 4.1.1. Given a Higgs bundle of the form $\left(\mathcal{O}(d_1) \oplus \mathcal{O}(d_2), \begin{pmatrix} a & b \\ c & d \end{pmatrix}\right)$, we let

$$d - a : H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2)) \to H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + t))$$

denote the map given by $\beta \mapsto \beta \otimes (d - a)$.

Proposition 4.1.2 (Refined Higgs Harder-Narasimhan stratification for the stack of rank 2 and odd degree $\mathcal{O}(t)$-twisted Higgs bundles on $\mathbb{P}^1$)**.** Suppose that d is odd. The refined Higgs

Harder-Narasimhan stratification of the stack[5] $\mathscr{H}_{2,d}$ is given by:

$$\mathscr{H}_{2,d} = \mathscr{H}_{2,d}^{ss} \sqcup \bigsqcup_{\mu=(d_1,d_2)\neq\mu_0} \left(\mathscr{H}_{2,d}^{\mu\text{-s}} \sqcup \mathscr{H}_{2,d}^{\mu,\delta_\mu,\text{indec}} \sqcup \mathscr{H}_{2,d}^{\mu,\text{dec}} \right),$$

where $\delta_\mu := d_1 - d_2 + 1$ given $\mu = (d_1, d_2)$, with each stratum non-empty and admitting a quasi-projective coarse moduli space. Moreover, given a Higgs Harder-Narasimhan type $\mu = (d_1, d_2)$, the strata refining $\mathscr{H}_{2,d}^{\mu}$ can be described as follows:

(i) the stratum $\mathscr{H}_{2,d}^{\mu\text{-}s}$ is open in $\mathscr{H}_{2,d}^{\mu}$ and parametrises isomorphism classes of Higgs bundles of the form

$$\left(\mathcal{O}(d_1) \oplus \mathcal{O}(d_2), \begin{pmatrix} a & b \\ 0 & d \end{pmatrix} \right)$$

where $a, d \in H^0(\mathbb{P}^1, \mathcal{O}(t))$ with $a \neq d$ and $b \in H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + t))$ is such that the image of b in $\operatorname{coker}(d - a)$ is non-zero;

(ii) the stratum $\mathscr{H}_{2,d}^{\mu,\delta_\mu,\text{indec}}$ is locally closed in $\mathscr{H}_{2,d}^{\mu}$ and parametrises isomorphism classes of Higgs bundles of the form

$$\left(\mathcal{O}(d_1) \oplus \mathcal{O}(d_2), \begin{pmatrix} a & b \\ 0 & a \end{pmatrix} \right)$$

where $a \in H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + t)) \setminus \{0\}$ and $b \neq 0$;

(iii) the stratum $\mathscr{H}_{2,d}^{\mu,\text{dec}}$ is closed in $\mathscr{H}_{2,d}^{\mu}$ and parametrises isomorphism classes of Higgs bundles of the form

$$\left(\mathcal{O}(d_1) \oplus \mathcal{O}(d_2), \begin{pmatrix} a & 0 \\ 0 & d \end{pmatrix} \right)$$

where $a, d \in H^0(\mathbb{P}^1, \mathcal{O}(t))$.

Remark 4.1.3 (Assumption that d is odd)**.** In the above Proposition 4.1.2, we have assumed that the degree d is odd. As per Remark 3.2.11, this assumption is only required to ensure that the moduli space of semistable Higgs bundles is a coarse moduli space and that it is smooth; the refinement of the unstable strata is independent of the parity of the degree, in the sense that even when d is even, the above unstable refined Higgs Harder-Narasimhan strata still admit quasi-projective coarse moduli spaces.

To prove Proposition 4.1.2, we will use the following

[5] We recall that we use the notation $\mathscr{H}_{2,d}$ in this section to denote the moduli stack $\mathscr{H}_{2,d}(\mathbb{P}^1, \mathcal{O}(t))$ of $\mathcal{O}(t)$-twisted Higgs bundles of rank 2 and degree d on $\mathbb{P}^1$.

Lemma 4.1.4. Let $\mu = (d_1, d_2) \neq \mu_0$ be a Higgs Harder-Narasimhan type of rank 2 and degree d and let $[(E, \phi)] \in \mathscr{H}^{\mu}_{2,d}$. Then $E \cong \mathcal{O}(d_1) \oplus \mathcal{O}(d_2)$ and, with respect to this identification, the Higgs field ϕ can be written as a matrix of the form $\begin{pmatrix} a & b \\ 0 & d \end{pmatrix}$ where $a, d \in H^0(\mathbb{P}^1, \mathcal{O}(t))$ and $b \in H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + t))$. Moreover, we have:

$$\mathrm{End}_{-1}(E, \phi) = \begin{cases} \{0\} & \text{if } a \neq d \\ \mathrm{End}_{-1}(E) & \text{if } a = d. \end{cases}$$

Proof. As we are considering the rank 2 case, by Proposition 3.1.8 we know that if (E, ϕ) has Higgs Harder-Narasimhan type μ then the underlying vector bundle E has the same Harder-Narasimhan type, and that its Higgs Harder-Narasimhan and Harder-Narasimhan filtrations coincide. Since E is a vector bundle on $\mathbb{P}^1$, it can be written as a direct sum of line bundles, say $\mathcal{O}(e_1) \oplus \mathcal{O}(e_2)$. Without loss of generality we can assume that $e_1 \geq e_2$, so it suffices to show that $e_1 = d_1$. Suppose that $e_1 \neq d_1$. Then since E has Harder-Narasimhan type (d_1, d_2), we must have that $e_1 < d_1$. Moreover, any line bundle on $\mathbb{P}^1$ of degree d_1 is isomorphic to $\mathcal{O}(d_1)$ so we can assume that there is an inclusion $\mathcal{O}(d_1) \hookrightarrow \mathcal{O}(e_1) \oplus \mathcal{O}(e_2)$. Post-composing this inclusion with the projection map to $\mathcal{O}(e_1)$ yields a morphism $\mathcal{O}(d_1) \to \mathcal{O}(e_1)$ which must be zero by degree considerations (since both bundles are stable). Thus $\mathcal{O}(d_1)$ is contained in the kernel of the projection map, namely in $\mathcal{O}(e_2)$, which for degree reasons is impossible since $e_2 \leq e_1 < d_1$ by assumption. Hence by contradiction $e_1 = d_1$. Moreover, since $[(E, \phi)]$ is unstable by assumption, writing ϕ in matrix form with respect to the splitting $\mathcal{O}(d_1) \oplus \mathcal{O}(d_2)$ of E, it must be that ϕ preserves $\mathcal{O}(d_1)$ so that ϕ can be written in the form $\begin{pmatrix} a & b \\ 0 & d \end{pmatrix}$ where $a, d \in H^0(\mathbb{P}^1, \mathcal{O}(t))$ and $b \in H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + t))$.

To prove the second statement, let $\psi \in \mathrm{End}_{-1}(E)$ with $\psi \neq 0$, and consider the diagram

$$\begin{array}{ccc} \mathcal{O}(d_1) \oplus \mathcal{O}(d_2) & \xrightarrow{\quad\psi\quad} & \mathcal{O}(d_1) \oplus \mathcal{O}(d_2) \\ \downarrow{\scriptstyle\phi} & & \downarrow{\scriptstyle\phi} \\ (\mathcal{O}(d_1) \oplus \mathcal{O}(d_2)) \otimes \mathcal{O}(t) & \xrightarrow{\psi\otimes\mathrm{id}} & (\mathcal{O}(d_1) \oplus \mathcal{O}(d_2)) \otimes \mathcal{O}(t). \end{array}$$

Writing $\psi = \begin{pmatrix} 0 & \beta \\ 0 & 0 \end{pmatrix}$ where $\beta \in \mathrm{Hom}(\mathcal{O}(d_2), \mathcal{O}(d_1)) \cong H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2))$, we see that the above diagram commutes if and only if the diagram

$$\begin{array}{ccc} \mathcal{O}(d_2) & \xrightarrow{\quad\beta\quad} & \mathcal{O}(d_1) \\ \downarrow{\scriptstyle a} & & \downarrow{\scriptstyle d} \\ \mathcal{O}(d_1) \otimes \mathcal{O}(t) & \xrightarrow{\beta\otimes\mathrm{id}} & \mathcal{O}(d_2) \otimes \mathcal{O}(t) \end{array}$$

commutes, or equivalently if and only if $a \otimes \beta = \beta \otimes d$ as elements of $H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + t))$. Since β is non-zero (as ψ is non-zero), it follows that for $\psi \in \operatorname{End}_{-1}(E)$ and $\psi \neq 0$, we have that $\psi \in \operatorname{End}_{-1}(E, \phi)$ if and only if $a = d$. We note that we have shown in Proposition 3.2.12 in the general case of L-twisted Higgs bundles of rank 2 (not necessarily on $\mathbb{P}^1$), that if $a \neq d$ then $\operatorname{End}_{-1}(E, \phi) = \{0\}$. □

We can now prove Proposition 4.1.2.

Proof of Proposition 4.1.2. Given $\mu = (d_1, d_2) \neq \mu_0$, it follows from the second part of Lemma 4.1.4 that if $[(E, \phi)]$ lies in $\mathscr{H}^{\mu}_{2,d}$ then $\dim \operatorname{End}_{-1}(E, \phi)$ is either 0 or $\delta_\mu = d_1 - d_2 + 1$. Thus for $\delta \neq 0$ and $\delta \neq d_1 - d_2 + 1$, the stratum $\mathscr{H}^{\mu,\delta,\text{indec}}_{2,d}$ is empty. Hence the refined Higgs Harder-Narasimhan stratification in rank 2 obtained in Theorem 3.2.9 simplifies to the desired stratification when $\Sigma = \mathbb{P}^1$. We now describe the unstable refined Higgs Harder-Narasimhan strata.

If $[(E, \phi)] \in \mathscr{H}^{\mu}_{2,d}$, then from the first part of Lemma 4.1.4 we can assume without loss of generality that $E = \mathcal{O}(d_1) \oplus \mathcal{O}(d_2)$ and that $\phi = \begin{pmatrix} a & b \\ 0 & d \end{pmatrix}$ where $a, d \in H^0(\mathbb{P}^1, \mathcal{O}(t))$ and $b \in H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + t))$. By the second part of Lemma 4.1.4, we have that $\operatorname{End}_1(E, \phi) = 0$ if and only if $a \neq d$. Thus if $[(E, \phi)] \in \mathscr{H}^{\mu\text{-s}}_{2,d}$ then ϕ must satisfy the condition that $a \neq d$. We now establish the condition on b required to ensure that $(E, \phi) \ncong \operatorname{gr}(E, \phi)$, so that $[(E, \phi)] \in \mathscr{H}^{\mu\text{-s}}_{2,d}$.

We have that $(E, \phi) \cong \operatorname{gr}(E, \phi)$ if and only if there exists an automorphism ψ of E such that the Higgs field $\psi \circ \phi \circ \psi^{-1}$ is diagonal. If $\psi \in \operatorname{Aut} E$, then by Proposition *A.0.1* (i) this automorphism can be represented in matrix form as $\begin{pmatrix} \alpha & \beta \\ 0 & \delta \end{pmatrix}$ where $\alpha, \delta \in \mathbb{G}_m$ and $\beta \in H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2))$. Conjugating ϕ by ψ yields the Higgs field

$$\psi \circ \phi \circ \psi^{-1} = \begin{pmatrix} a & \alpha\delta^{-1}b + \delta^{-1}\beta \otimes (d - a) \\ 0 & d \end{pmatrix}. \tag{4.1}$$

It follows then that (E, ϕ) is isomorphic to $\operatorname{gr}(E, \phi)$ if and only if there exists an automorphism $\psi = \begin{pmatrix} \alpha & \beta \\ 0 & \delta \end{pmatrix}$ of E such that $b = \alpha^{-1}\beta \otimes (d - a)$. Such an automorphism exists if and only if b lies in the image of the morphism $d - a : H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2)) \to H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + t))$ defined in Notation 4.1.1, in other words if and only if the image $\overline{b}$ of b in $\operatorname{coker}(d - a)$ is zero. It follows that $[(E, \phi)] \in \mathcal{M}^{\mu\text{-s}}_{2,d}$ if and only if $a \neq d$ and $\overline{b} \neq 0 \in \operatorname{coker}(d - a)$. Finally we note that since $t \geq 0$ and we are working over $\mathbb{P}^1$, the spaces $H^0(\mathbb{P}^1, \mathcal{O}(t))$ and $\operatorname{coker}(d - a)$ have positive dimension for any $a, d \in H^0(\mathbb{P}^1, \mathcal{O}(t))$ with $a \neq d$. This proves part (i).

Now suppose that $[(E,\phi)] \in \mathcal{M}_{2,d}^{\mu,\delta_\mu,\mathrm{indec}}$. From the second part of Lemma 4.1.4, we must have that $a = d$ to ensure that $\dim \mathrm{End}_{-1}(E,\phi) = \delta_\mu$. Moreover, we must have that $(E,\phi) \not\cong \mathrm{gr}(E,\phi)$. The condition that $b \neq 0$ is sufficient to ensure this, in contrast with the case above. Indeed, given $\psi = \begin{pmatrix} \alpha & \beta \\ 0 & \delta \end{pmatrix}$, we have:

$$\psi \circ \begin{pmatrix} a & b \\ 0 & a \end{pmatrix} \circ \psi^{-1} = \begin{pmatrix} a & \alpha\delta^{-1}b \\ 0 & a \end{pmatrix}.$$

If $b \neq 0$, then since $\alpha, \delta \in \mathbb{G}_m \cong \mathbb{C}^*$, the resulting Higgs field has non-zero top right entry. Thus $[(E,\phi)] \in \mathcal{M}_{2,d}^{\mu,\delta_\mu,\mathrm{indec}}$ if and only if $a = d$ and $b \neq 0$. Finally we note that since $t \geq 0$, the space $H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + t))$ has positive dimension. This proves part (ii).

Finally, suppose that $[(E,\phi)] \in \mathcal{M}_{2,d}^{\mu,\mathrm{dec}}$. Then $(E,\phi) \cong \mathrm{gr}(E,\phi)$ so there is a representative in the isomorphism class $[(E,\phi)]$ of the form $\left(\mathcal{O}(d_1) \oplus \mathcal{O}(d_2), \begin{pmatrix} a & 0 \\ 0 & d \end{pmatrix}\right)$. Moreover, any such Higgs bundle is decomposable. We also have that $h^0(\mathbb{P}^1, \mathcal{O}(t)) > 0$ since $t \geq 0$. This proves part (iii). □

We can use the analysis from the above proof to compute the dimensions of the refined Higgs Harder-Narasimhan strata.

Proposition 4.1.5 (Dimensions of the refined Higgs Harder-Narasimhan strata)**.** Let $\mu = (d_1, d_2)$ be an unstable Higgs Harder-Narasimhan type and let $\delta_\mu = d_1 - d_2 + 1$. Then the dimensions of the refined strata for the Higgs Harder-Narasimhan stratum $\mathscr{H}_{2,d}^{\mu}$ are as per the following table.

Refined stratum of $\mathscr{H}_{2,d}^{\mu}(\mathbb{P}^1, \mathcal{O}(t))$	Dimension
$\mathscr{H}_{2,d}^{\mu\text{-s}}(\mathbb{P}^1, \mathcal{O}(t))$	$3t$
$\mathscr{H}_{2,d}^{\mu,\delta_\mu,\mathrm{indec}}(\mathbb{P}^1, \mathcal{O}(t))$	$2t - 1$
$\mathscr{H}_{2,d}^{\mu,\mathrm{dec}}(\mathbb{P}^1, \mathcal{O}(t))$	$2t$

The proof of Proposition 4.1.5 relies on the following

Lemma 4.1.6 (Identification of a Higgs Harder-Narasimhan stratum as an irreducible quotient stack)**.** Let $\mu = (d_1, d_2)$ denote an unstable Higgs Harder-Narasimhan type of rank 2 and degree d. Then there is an isomorphism

$$\mathscr{H}_{2,d}^{\mu} \cong \left[\mathrm{Hom}^{\mu}(\mathcal{O}(d_2), \mathcal{O}(d_1 + t)) / \mathrm{Aut}(\mathcal{O}(d_1) \oplus \mathcal{O}(d_2))\right],$$

where $\mathrm{Hom}^{\mu}(\mathcal{O}(d_2), \mathcal{O}(d_1+t))$ denotes the subspace of $\mathrm{Hom}(\mathcal{O}(d_2), \mathcal{O}(d_1+t))$ consisting of Higgs fields which preserve $\mathcal{O}(d_1)$. Moreover, this quotient stack is irreducible and its dimension is given by

$$\dim \mathscr{H}^{\mu}_{2,d} = 3t.$$

Proof. To simplify notation, let E denote the vector bundle $\mathcal{O}(d_1) \oplus \mathcal{O}(d_2)$. By the first part of Lemma 4.1.4 we have that if $[(E', \phi)] \in \mathscr{H}^{\mu}_{2,d}$, then E' has Harder-Narasimhan type μ and therefore $E' \cong E$. Thus we have that $\mathscr{H}^{\mu}_{2,d} \subseteq F^{-1}([E])$ where F denotes the forgetful map $\mathscr{H}_{2,d} \to \mathscr{V}_{2,d}$. By Proposition 3.1.5, we have an isomorphism $F^{-1}(E) \cong [\mathrm{Hom}(E, E \otimes \mathcal{O}(t)) / \mathrm{Aut}(E)]$. Restricting this isomorphism to $\mathscr{H}^{\mu}_{2,d} = F^{-1}(E) \cap \mathscr{H}^{\mu}_{2,d}$, we obtain an isomorphism

$$\mathscr{H}^{\mu}_{2,d} \cong [\mathrm{Hom}^{\mu}(E, E \otimes \mathcal{O}(t)) / \mathrm{Aut}(E)] \tag{4.2}$$

where $\mathrm{Hom}^{\mu}(E, E \otimes \mathcal{O}(t))$ is the subspace of $\mathrm{Hom}(E, E \otimes \mathcal{O}(t))$ consisting of Higgs fields which preserve $\mathcal{O}(d_1)$. This subspace has dimension $2 \times h^0(\mathbb{P}^1, \mathcal{O}(t)) + h^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + t)) = 3t + d_1 - d_2 + 3$, and it is irreducible. Therefore $\mathscr{H}^{\mu}_{2,d}$ is also irreducible. Moreover, by Proposition A.0.1 (i), we have that

$$\dim \mathrm{Aut}(E) = \dim \left(H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2)) \rtimes (\mathbb{G}_m \times \mathbb{G}_m)\right) = d_1 - d_2 + 3.$$

Thus we have that

$$\begin{aligned} \dim \mathscr{H}^{\mu}_{2,d} &= \dim \mathrm{Hom}^{\mu}(E, E \otimes \mathcal{O}(t)) - \dim \mathrm{Aut}(E) \\ &= 3t + d_1 - d_2 + 3 - (d_1 - d_2 + 3) = 3t. \end{aligned}$$

□

We can now prove Proposition 4.1.5 which establishes the dimensions of the refined strata.

Proof of Proposition 4.1.5. We again let E denote the vector bundle $\mathcal{O}(d_1) \oplus \mathcal{O}(d_2)$. By construction, the stratum $\mathscr{H}^{\mu\text{-s}}_{2,d}$ is the open stratum in the refinement of the unstable stratum $\mathscr{H}^{\mu}_{2,d}$, which is irreducible by Lemma 4.1.6. Thus the open stratum $\mathscr{H}^{\mu\text{-s}}_{2,d}$ has the same dimension as $\mathscr{H}^{\mu}_{2,d}$, namely $3t$ by Lemma 4.1.6.

We now consider the locally closed stratum $\mathscr{H}^{\mu,\delta_\mu,\mathrm{indec}}_{2,d}$. We let $\mathrm{Hom}^{\mu,=,\mathrm{indec}}(E, E \otimes \mathcal{O}(t))$ denote the subspace of $\mathrm{Hom}(E, E \otimes \mathcal{O}(t))$ consisting of Higgs fields of the form $\begin{pmatrix} a & b \\ 0 & a \end{pmatrix}$ where

$b \neq 0$. This subspace is invariant under the action of $\operatorname{Aut} E$. Moreover, the isomorphism of (4.2) restricts to the isomorphism

$$\mathscr{H}_{2,d}^{\mu,\delta_\mu,\mathrm{indec}} \cong \left[\operatorname{Hom}^{\mu,=,\mathrm{indec}}(E, E \otimes \mathcal{O}(t)) / \operatorname{Aut} E\right].$$

Thus $\mathscr{H}_{2,d}^{\mu,\delta_\mu,\mathrm{indec}}$ has dimension equal to $d_1 - d_2 + t + 1 - (d_1 - d_2 + 3) = t - 2$.

For the closed stratum, we proceed by determining the dimension of an open dense substack of $\mathscr{H}_{2,d}^{\mu,\mathrm{dec}}$, consisting of those decomposable Higgs bundles with distinct entries on the diagonal. Under the identification of $\mathscr{H}_{2,d}^{\mu}$ as a quotient stack given in (4.2), we define the substack to be the quotient stack

$$\left[\operatorname{Hom}^{\mu,\neq,\mathrm{dec}}(E, E \otimes \mathcal{O}(t)) / \operatorname{Aut} E\right]$$

where $\operatorname{Hom}^{\mu,\neq,\mathrm{dec}}(E, E \otimes \mathcal{O}(t))$ denotes the subset of $\operatorname{Hom}^{\mu}(E, E \otimes \mathcal{O}(t))$ consisting of Higgs fields of the form $\begin{pmatrix} a & b \\ 0 & d \end{pmatrix}$ where $a \neq d$ and $b \in \operatorname{im}(d-a) \subseteq H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + t))$. This subset is indeed invariant under $\operatorname{Aut} E$, and dense in the subspace $\operatorname{Hom}^{\mu,\mathrm{dec}}(E, E \otimes \mathcal{O}(t))$ of $\operatorname{Hom}^{\mu}(E, E \otimes \mathcal{O}(t))$ consisting of Higgs fields of the form $\begin{pmatrix} a & b \\ 0 & d \end{pmatrix}$ where $b \in \operatorname{im}(d-a)$, namely those which give rise to decomposable Higgs bundles. Thus it suffices to compute the dimension of the above quotient stack, which is given by $(t+1+t+1+d_1-d_2+1)-(d_1+d_2+3) = 2t$. $\square$

4.1.2 Geometry of the coarse moduli spaces

In this section we study the coarse moduli spaces of the refined Higgs Harder-Narasimhan strata considered in the previous Section 4.1.1, with the aim of describing their geometry. Firstly, we explicitly describe the moduli spaces using the map taking a Higgs bundle to its associated graded (Theorem 4.1.7). These descriptions of the moduli spaces allow us to compute their dimensions and show that they are smooth (Corollary 4.1.10). Secondly, we study the Hitchin morphism for each of the moduli spaces and describe its image and fibres, showing moreover that the fibres are smooth (Proposition 4.1.12).

We start by stating Theorem 4.1.7.

Theorem 4.1.7 (Description of moduli spaces of unstable Higgs bundles of rank 2 on $\mathbb{P}^1$)**.** Let $\mu = (d_1, d_2)$ denote an unstable Higgs Harder-Narasimhan type of rank 2 and degree d, and let $\delta_\mu = d_1 - d_2 + 1$. Then the quasi-projective coarse moduli spaces for the refined Higgs Harder-Narasimhan strata are smooth and can be described as follows:

(i) there is an isomorphism

$$\mathcal{M}_{2,d}^{\mu,\mathrm{dec}} \cong \mathbb{C}^{t+1} \times \mathbb{C}^{t+1}$$
$$\left(\mathcal{O}(d_1) \oplus \mathcal{O}(d_2), \begin{pmatrix} a & 0 \\ 0 & d \end{pmatrix}\right) \mapsto (a, d),$$

using the identification of $H^0(\mathbb{P}^1, \mathcal{O}(t))$ with the space $\mathbb{C}^{t+1}$ parametrising polynomials of degree t.

(ii) the map sending a Higgs bundle to its associated graded induces a surjection

$$\mathrm{gr} : \mathcal{M}_{2,d}^{\mu\text{-s}} \twoheadrightarrow (\mathbb{C}^{t+1} \times \mathbb{C}^{t+1}) \setminus \Delta$$
$$\left(\mathcal{O}(d_1) \oplus \mathcal{O}(d_2), \begin{pmatrix} a & b \\ 0 & d \end{pmatrix}\right) \mapsto (a, d)$$

where Δ denotes the diagonal in the product. Moreover, for each $(a, d) \in (\mathbb{C}^{t+1} \times \mathbb{C}^{t+1}) \setminus \Delta$, there is an isomorphism $\mathrm{gr}^{-1}(a, d) \cong \mathbb{P}(\mathrm{coker}(d - a)) \cong \mathbb{P}^{t-1}$;

(iii) there is an isomorphism

$$\mathcal{M}_{2,d}^{\mu,\delta_\mu,\mathrm{indec}} \to \Delta \times \mathbb{P}^{d_1 - d_2 + t}$$
$$\left(\mathcal{O}(d_1) \oplus \mathcal{O}(d_2), \begin{pmatrix} a & b \\ 0 & a \end{pmatrix}\right) \mapsto ((a, a), [b])$$

where $[b]$ denotes the image of b in $\mathbb{P}(H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + 1)) \cong \mathbb{P}^{d_1 - d_2 + t}$.

Figure 4.1 illustrates the results of Theorem 4.1.7 above.

Remark 4.1.8 (The $\mathbb{C}^*$-action and semiprojectivity)**.** An unstable Higgs bundle $[(E, \phi)]$ of Higgs Harder-Narasimhan type $\mu = (d_1, d_2)$ is fixed by the Higgs field scaling $\mathbb{C}^*$-action if and only if the Higgs field is strictly upper triangular (we will study fixed points in more detail in Section 4.2.1, in the case of rank 2 Higgs bundles on a curve of arbitrary genus). By Theorem 4.1.7, the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}$ contains no Higgs bundles of this form. Thus we see that unlike the moduli space $\mathcal{M}_{2,d}^{s}$ of stable Higgs bundles, limits at 0 under the $\mathbb{C}^*$-action do not exist for the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}$. Note that we can see this directly by observing that the condition that the diagonal entries of the Higgs field are distinct (to ensure that the Higgs bundle lies in $\mathcal{M}_{2,d}^{\mu\text{-s}}$) is not preserved upon taking the limit as $t \to 0$. Therefore the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}$ cannot be semiprojective.

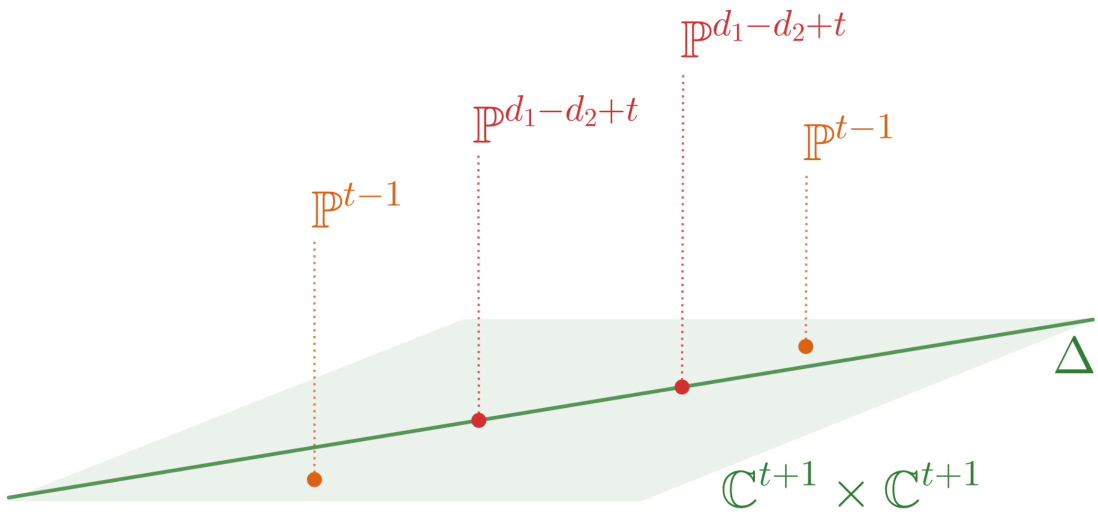

Figure 4.1: Moduli spaces for unstable $\mathcal{O}(t)$-twisted rank 2 Higgs bundles on $\mathbb{P}^1$ of Higgs Harder-Narasimhan type $\mu = (d_1, d_2) \neq \mu_0$. We let $\delta_\mu = d_1 - d_2 + 1$. The horizontal green plane represents the moduli space $\mathcal{M}_{2,d}^{\mu,\text{dec}} \cong \mathbb{C}^{t+1} \times \mathbb{C}^{t+1}$. The green line represents the diagonal Δ inside $\mathbb{C}^{t+1} \times \mathbb{C}^{t+1}$. The dotted lines represent the fibres of the map from the moduli spaces $\mathcal{M}_{2,d}^{\mu\text{-s}}$ and $\mathcal{M}_{2,d}^{\mu,\delta_\mu,\text{indec}}$ given by $[(E,\phi)] \mapsto [\text{gr}(E,\phi)]$. Under this map, we have: (i) the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}$ is mapped onto the complement of the diagonal, with fibres isomorphic to $\mathbb{P}^{t-1}$ (coloured in orange); (ii) the moduli space $\mathcal{M}_{2,d}^{\mu,\delta_\mu,\text{indec}}$ is mapped onto the diagonal Δ, with constant fibres isomorphic to $\mathbb{P}^{d_1-d_2+t}$ (coloured in red).

By contrast, the moduli space $\mathcal{M}_{2,d}^{\mu,\delta_\mu,\text{indec}}$ contains Higgs bundles with strictly upper triangular Higgs field, and two such Higgs bundles are isomorphic if and only if the top right entries of the Higgs fields are scalar multiples of each other (see the proof of Theorem 4.1.7). Thus the set of fixed points for the $\mathbb{C}^*$-action on $\mathcal{M}_{2,d}^{\mu,\delta_\mu,\text{indec}}$ can be identified with the projective space $\mathbb{P}(H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + t)))$. Moreover, it can be shown that any Higgs bundle $\left[\left(\mathcal{O}(d_1) \oplus \mathcal{O}(d_2), \begin{pmatrix} a & b \\ 0 & a \end{pmatrix}\right)\right] \in \mathcal{M}_{2,d}^{\mu,\delta_\mu,\text{indec}}$ has a well-defined limit at 0 under the $\mathbb{C}^*$-action, given by $\left[\left(\mathcal{O}(d_1) \oplus \mathcal{O}(d_2), \begin{pmatrix} 0 & b \\ 0 & 0 \end{pmatrix}\right)\right]$. It follows that the moduli space $\mathcal{M}_{2,d}^{\mu,\delta_\mu,\text{indec}}$ is semiprojective, just as the moduli space $\mathcal{M}_{2,d}^{s}$ is semiprojective.

We will generalise the above observations concerning semiprojectivity to moduli spaces of unstable Higgs bundles of rank 2 on curves of arbitrary genus in Section 4.2 (see Remarks 4.2.3, 4.2.7, and 4.2.10).

Remark 4.1.9 (Compactifying $\mathcal{M}_{2,d}^{\mu\text{-s}}$ and S-equivalence for Non-Reductive GIT)**.** One approach to compactifying the quasi-projective moduli space of semistable Higgs bundles is to ‘add in’ limits under the $\mathbb{C}^*$-action at the boundary (see [89, 46]); only limits as $t \to \infty$ need to be added since limits at 0 always exist. In the case of the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}$ of μ-stable Higgs bundles of Higgs Harder-Narasimhan type μ, limits can fail to exist even at 0, as seen in Remark 4.1.8 above. We can therefore hope to be able to compactify the quasi-projective moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}$ by first adding in limits at 0 under the $\mathbb{C}^*$-action (this step is not required for the moduli space of semistable Higgs bundles), and then adding in limits at ∞ (analogously to the semistable case). Figure 4.1 suggests how the first step might be achieved: it should consist in adding to the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}$ a boundary consisting of fibres lying over the diagonal in $\mathcal{M}_{2,d}^{\mu,\text{dec}} \cong \mathbb{C}^{t+1} \times \mathbb{C}^{t+1}$, of dimension $t-1$. These fibres would correspond to quotients of the fibres of the map $\text{gr} : \mathcal{M}_{2,d}^{\mu,\delta_\mu,\text{indec}} \to \mathcal{M}_{2,d}^{\mu,\text{dec}}$. In other words, a weaker equivalence relation than that of isomorphism would be defined for Higgs bundles parametrised by $\mathcal{M}_{2,d}^{\mu,\delta_\mu,\text{indec}}$, so that the new fibres could be interpreted as equivalence classes (rather than isomorphism classes) of Higgs bundles in $\mathcal{M}_{2,d}^{\mu,\delta_\mu,\text{indec}}$.

This strategy is reminiscent of the notion of S-equivalence in classical GIT (see [91, Prop 9]), which describes how, given the linear action of a reductive group G on a projective variety X, strictly semistable orbits must be identified in order to obtain a projective completion of the geometric quotient X^s/G. In particular, the example of $\mathcal{M}_{2,d}^{\mu\text{-s}}$ suggests that a similar notion of S-equivalence might hold for Non-Reductive GIT quotients, based on the fact that the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}$ can be constructed as an open subset of a Non-Reductive GIT quotient, as we have seen in Section 3.2 of Chapter 3. We note that the notion of S-equivalence is at present absent from Non-Reductive GIT.

The dimensions of the moduli spaces for the refined Higgs Harder-Narasimhan strata can be computed directly using the above Theorem 4.1.7, as per Corollary 4.1.10 below.

Corollary 4.1.10 (Dimensions of the moduli spaces for rank 2 unstable $\mathcal{O}(t)$-twisted Higgs bundles on $\mathbb{P}^1$)**.** Given an unstable Higgs Harder-Narasimhan type $\mu = (d_1, d_2)$ of rank 2 and degree d, the dimensions of the moduli spaces for the refined Higgs Harder-Narasimhan strata for $\mathscr{H}_{2,d}^{\mu}$ are as per the following table.

Moduli space	Dimension
$\mathcal{M}_{2,d}^{\mu\text{-s}}(\mathbb{P}^1, \mathcal{O}(t))$	$3t+1$
$\mathcal{M}_{2,d}^{\mu,\delta_\mu,\text{indec}}(\mathbb{P}^1, \mathcal{O}(t))$	$2t+d_1-d_2+1$
$\mathcal{M}_{2,d}^{\mu,\text{dec}}(\mathbb{P}^1, \mathcal{O}(t))$	$2t+2$

We now prove Theorem 4.1.7.

Proof of Theorem 4.1.7. Most of the analysis required to prove this result is contained in the proof of Proposition 4.1.2. To simplify notation, we let E denote the vector bundle $\mathcal{O}(d_1)\oplus\mathcal{O}(d_2)$ and ϕ a Higgs field $\begin{pmatrix} a & b \\ 0 & d \end{pmatrix}$ where $a, d \in H^0(\mathbb{P}^1, \mathcal{O}(t))$ and $b \in H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + t))$.

We start by defining two morphisms of stacks which we will use. Firstly we define a map $\text{gr} : \mathscr{H}_{2,d}^{\mu,\delta,\text{indec}} \to \mathscr{H}_{2,d}^{\mu,\text{dec}}$ given by $[(E, \phi)] \mapsto [\text{gr}(E, \phi)]$ at the level of geometric points. This map extends naturally to families. Moreover, an automorphism of a Higgs bundle induces an automorphism of the associated graded Higgs bundle (and this extends to automorphisms of families), thus we obtain a well-defined map of stacks which induces an algebraic map $\text{gr} : \mathcal{M}_{2,d}^{\mu,\delta,\text{indec}} \to \mathcal{M}_{2,d}^{\mu,\text{dec}}$ of the corresponding coarse moduli spaces. Secondly, we define a map $\Gamma : \mathscr{H}_{2,d}^{\mu,\text{dec}} \to \mathscr{H}_{1,d_1}^{ss} \times \mathscr{H}_{1,d_2}^{ss}$ given by $[(E, \phi)] \mapsto ([(\mathcal{O}(d_1), a)], [(\mathcal{O}(d_2), d)]$ at the level of geometric points. Just as with the map gr, the map Γ is a well-defined map of stacks. Since $H^0(\mathbb{P}^1, \mathcal{O}(t)) \times H^0(\mathbb{P}^1, \mathcal{O}(t))$ is the coarse moduli space for $\mathscr{H}_{1,d_1}^{ss} \times \mathscr{H}_{1,d_2}^{ss}$, while $\mathcal{M}_{2,d}^{\mu,\text{dec}}$ is the coarse moduli space for $\mathscr{H}_{2,d}^{\mu,\text{dec}}$, we obtain an induced algebraic map $\gamma : \mathcal{M}_{2,d}^{\mu,\text{dec}} \to H^0(\mathbb{P}^1, \mathcal{O}(t)) \times H^0(\mathbb{P}^1, \mathcal{O}(t))$.

We consider first the moduli space $\mathcal{M}_{2,d}^{\mu,\text{dec}}$, which we wish to show is isomorphic to the product $H^0(\mathbb{P}^1, \mathcal{O}(t)) \times H^0(\mathbb{P}^1, \mathcal{O}(t))$ via the map γ described above. It will then follow immediately that $\mathcal{M}_{2,d}^{\mu,\text{dec}}$ is smooth. We observe that the map $\gamma : \mathcal{M}_{2,d}^{\mu,\text{dec}} \to H^0(\mathbb{P}^1, \mathcal{O}(t)) \times H^0(\mathbb{P}^1, \mathcal{O}(t))$ is given by $[(E, \phi)] \mapsto (a, d)$; this is well-defined since by (4.1) we can see that conjugating ϕ by an automorphism of E preserves the diagonal entries. Moreover it is a bijection since given any $(a, d) \in H^0(\mathbb{P}^1, \mathcal{O}(t)) \times H^0(\mathbb{P}^1, \mathcal{O}(t))$, the Higgs bundle $[(E, \phi)]$, where ϕ is the diagonal matrix with (a, d) on the diagonal, lies in $\mathcal{M}_{2,d}^{\mu,\text{dec}}$. Thus fixing an isomorphism $H^0(\mathbb{P}^1, \mathcal{O}(t)) \times H^0(\mathbb{P}^1, \mathcal{O}(t)) \cong \mathbb{C}^{t+1} \times \mathbb{C}^{t+1}$, we have that γ is a bijective algebraic map $\mathcal{M}_{2,d}^{\mu,\text{dec}} \to \mathbb{C}^{t+1} \times \mathbb{C}^{t+1}$. Since $\mathcal{M}_{2,d}^{\mu,\text{dec}}$ is irreducible, since $\mathbb{C}^{t+1} \times \mathbb{C}^{t+1}$ is irreducible and smooth, and since we are working over a field of characteristic zero, it follows that γ is an isomorphism (using the result of [83, §2.4.4] for example).

We now turn to the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}$, which as seen at the beginning of the proof admits an algebraic map $\mathrm{gr} : \mathcal{M}_{2,d}^{\mu\text{-s}} \to \mathcal{M}_{2,d}^{\mu,\mathrm{dec}}$ given by $[(E,\phi)] \mapsto [\mathrm{gr}(E,\phi)]$. By Proposition 4.1.2 (i) we have that a Higgs bundle $[(E,\phi)]$ lies in $\mathcal{M}_{2,d}^{\mu\text{-s}}$ if and only if $a \neq d$ and the image of b in $\mathrm{coker}(d-a)$ is non-zero, where we recall that the morphism $d-a : H^0(\mathbb{P}^1, \mathcal{O}(d_1-d_2)) \to H^0(\mathbb{P}^1, \mathcal{O}(d_1-d_2+t))$ is given by $\beta \mapsto \beta \otimes (d-a)$. Thus under the identification $\mathcal{M}_{2,d}^{\mu,\mathrm{dec}} \cong \mathbb{C}^{t+1} \times \mathbb{C}^{t+1}$ proved above, the image of the morphism gr is contained in the complement of the diagonal $\Delta \subseteq \mathbb{C}^{t+1} \times \mathbb{C}^{t+1}$. To describe the fibre of this map, suppose that $(a,d) \in (\mathbb{C}^{t+1} \times \mathbb{C}^{t+1}) \setminus \Delta$. Then $\mathrm{gr}^{-1}(a,d)$ consists of isomorphism classes of Higgs bundles of the form $[(E,\phi)]$ where $b \in H^0(\mathbb{P}^1, \mathcal{O}(d_1-d_2+t)))$ satisfies the condition that its image in $\mathrm{coker}(d-a)$ is non-zero. We wish to determine when two such choices b and b' give rise to isomorphic Higgs bundles. By (4.1), this is true if and only if there exists an automorphism $\psi = \begin{pmatrix} \alpha & \beta \\ 0 & \delta \end{pmatrix}$ of E such that

$$b' = \alpha\delta^{-1} b + \delta^{-1}\beta \otimes (d-a). \tag{4.3}$$

We let $\bar{b}$ and $\overline{b'}$ denote the images of b and b' respectively in $\mathrm{coker}(d-a)$, which by assumption are non-zero. Thus we can consider their equivalence classes $[\bar{b}]$ and $[\overline{b'}]$ in $\mathbb{P}(\mathrm{coker}(d-a))$. Then (4.3) is satisfied for an automorphism ψ of E if and only if $[\bar{b}] = [\overline{b'}]$. It follows that there is an isomorphism $\mathbb{P}(\mathrm{coker}(d-a))) \to \mathrm{gr}^{-1}(a,d)$ induced by the morphism $H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2+t)) \setminus \{0\} \to \mathrm{gr}^{-1}(a,d)$ given by $b \mapsto \left[\left(E, \begin{pmatrix} a & b \\ 0 & d \end{pmatrix}\right)\right]$. Moreover, since $a \neq d$ (to ensure that the resulting Higgs bundle lies in the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}$), the morphism $d-a$ is injective. Therefore $\mathrm{coker}(d-a)$ has dimension equal to $d_1+d_2+t+1-(d_1+d_2+1) = t$, so that the fibre $\mathrm{gr}^{-1}(a,d)$ has dimension $t-1$ for any $(a,d) \in (\mathbb{C}^{t+1} \times \mathbb{C}^{t+1}) \setminus \Delta$. Thus $\mathcal{M}_{2,d}^{\mu\text{-s}}$ admits a surjective morphism onto the smooth variety $(\mathbb{C}^{t+1} \times \mathbb{C}^{t+1}) \setminus \Delta$ with equidimensional smooth fibres. It follows that $\mathcal{M}_{2,d}^{\mu\text{-s}}$ is smooth (by [90, Thm 3.3.27] for example).

Finally, we consider the moduli space $\mathcal{M}_{2,d}^{\mu,\delta_\mu,\mathrm{indec}}$. Again we can consider the map $\mathrm{gr} : \mathcal{M}_{2,d}^{\mu,\delta_\mu,\mathrm{indec}} \to \mathcal{M}_{2,d}^{\mu,\mathrm{dec}}$ taking a Higgs bundle to its associated graded. By Proposition 4.1.2 (ii), we have that $[(E,\phi)] \in \mathcal{M}_{2,d}^{\mu,\delta_\mu,\mathrm{indec}}$ if and only if $a = d$ and $b \neq 0$. Thus the map gr has image contained in the diagonal $\Delta \subseteq (\mathbb{C}^{t+1} \times \mathbb{C}^{t+1})$ under the identification $\mathcal{M}_{2,d}^{\mu,\mathrm{dec}} \cong \mathbb{C}^{t+1} \times \mathbb{C}^{t+1}$. Given $(a,a) \in \Delta$, the fibre $\mathrm{gr}^{-1}(a,a)$ consists of isomorphism classes of Higgs bundles of the form $[(E,\phi)]$ where $b \neq 0$. Moreover, using (4.1) we have that

$$\left[\left(E, \begin{pmatrix} a & b \\ 0 & a \end{pmatrix}\right)\right] = \left[\left(E, \begin{pmatrix} a & b' \\ 0 & a \end{pmatrix}\right)\right]$$

if and only if there exists an automorphism ψ of E such that $b' = \alpha\delta^{-1}b$. Since b and b' are non-zero by assumption, we can consider their equivalence classes $[b]$ and $[b']$ in $\mathbb{P}(H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + t)))$. Thus the map $H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + t)) \setminus \{0\} \to \mathrm{gr}^{-1}(a,a)$ given by $b \mapsto \left[\left(E, \begin{pmatrix} a & b \\ 0 & a \end{pmatrix}\right)\right]$ induces an isomorphism $\mathrm{gr}^{-1}(a,a) \cong \mathbb{P}(H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + t)))$ given by $[(E,\phi)] \mapsto [b]$. The fibres are therefore constant in dimension, and in fact we can identify $\mathcal{M}_{2,d}^{\mu,\delta_\mu,\mathrm{indec}}$ as a direct product via the map $\Delta \times (H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + 1)) \setminus \{0\}) \to \mathcal{M}_{2,d}^{\mu,\delta_\mu,\mathrm{indec}}$ given by $((a,a),b) \mapsto \left[\left(E, \begin{pmatrix} a & b \\ 0 & a \end{pmatrix}\right)\right]$ which descends to an isomorphism $\Delta \times \mathbb{P}(H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + 1))) \cong \mathcal{M}_{2,d}^{\mu,\delta_\mu,\mathrm{indec}}$. Fixing an isomorphism $H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + 1)) \cong \mathbb{C}^{d_1 - d_2 + t + 1}$, we obtain the desired isomorphism $\mathcal{M}_{2,d}^{\mu,\delta_\mu,\mathrm{indec}} \cong \Delta \times \mathbb{P}^{d_1 - d_2 + t}$, which in particular shows that the moduli space $\mathcal{M}_{2,d}^{\mu,\delta_\mu,\mathrm{indec}}$ is smooth. □

We now turn to describing the Hitchin fibration for the moduli spaces of unstable Higgs bundles of rank 2 on $\mathbb{P}^1$. As stated in the introduction to this Chapter 4, an important tool for studying the geometry of the moduli space of semistable Higgs bundles is the Hitchin morphism, which maps a Higgs bundle to the characteristic polynomial of its Higgs field; the Hitchin morphism is a surjective proper map onto an affine base with compact fibres which are generically smooth, and realises the moduli space as an integrable system (see [5] for the case of L-twisted Higgs bundles on $\mathbb{P}^1$). It is important to note that semistability is not required to define the Hitchin morphism, which can indeed be considered as a map from the stack of Higgs bundles. We can therefore consider the Hitchin morphism for each of the moduli spaces of unstable Higgs bundles which we have constructed in this thesis.

We now describe the fibres of the Hitchin morphism in the case of unstable Higgs bundles of rank 2 over $\mathbb{P}^1$. To do so, following [85] we restrict ourselves to the case of Higgs bundles with trace-free Higgs field[6]. That is, we consider the substack of $\mathscr{H}_{2,d}$ consisting of the fibre over 0 of the map $\mathscr{H}_{2,d} \to H^0(\mathbb{P}^1, \mathcal{O}(t))$ given by $[(E,\phi)] \mapsto \mathrm{tr}(\phi)$. Restricting our attention to trace-free Higgs fields means that the characteristic polynomial of a Higgs field is completely encoded by

[6] This simplification is made by Rayan in [85] as a means of obtaining a global description of the moduli space of $\mathcal{O}(2)$-twisted Higgs bundles of rank 2 and odd degree d on $\mathbb{P}^1$.

its determinant, which simplifies the analysis of the fibres of the Hitchin morphism. In the case of $\mathcal{O}(t)$-twisted Higgs bundles on $\mathbb{P}^1$ of rank 2, the determinant of the Higgs field is a section of $\mathcal{O}(2t)$, and so the base of the Hitchin fibration in this case is isomorphic to $H^0(\mathbb{P}^1, \mathcal{O}(2t))$.

Notation 4.1.11 (The trace-free case)**.** We let $\mathring{\mathscr{H}}_{2,d}$ denote the substack of $\mathscr{H}_{2,d}$ consisting of $\mathcal{O}(t)$-twisted Higgs bundles of rank 2 and degree d on $\mathbb{P}^1$ with trace-free Higgs field. Given an unstable Higgs Harder-Narasimhan type $\mu = (d_1, d_2)$, and setting $\delta_\mu = (d_1, d_2)$, we similarly define the subvarieties $\mathring{\mathcal{M}}^{\mu\text{-s}}_{2,d}$, $\mathring{\mathcal{M}}^{\mu,\delta_\mu,\text{indec}}_{2,d}$ and $\mathring{\mathcal{M}}^{\mu,\text{dec}}_{2,d}$ of $\mathcal{M}^{\mu\text{-s}}_{2,d}$, $\mathcal{M}^{\mu,\delta_\mu,\text{indec}}_{2,d}$ and $\mathcal{M}^{\mu,\text{dec}}_{2,d}$ respectively, consisting of those Higgs bundles with trace-free Higgs field. We let h denote the morphism from each of these moduli spaces to the Hitchin base $\mathcal{B} := H^0(\mathbb{P}^1, \mathcal{O}(2t))$ given by $[(E, \phi)] \mapsto \det(\phi)$. Moreover, we let $\mathcal{B}'$ denote the subset of $\mathcal{B}$ corresponding to the image of the map $H^0(\mathbb{P}^1, \mathcal{O}(t)) \to H^0(\mathbb{P}^1, \mathcal{O}(2t))$ given by $a \mapsto -a \otimes a$.

Proposition 4.1.12 (The Hitchin fibration for moduli spaces of unstable Higgs bundles in the trace-free case)**.** Let $\mu = (d_1, d_2)$ be an unstable Higgs Harder-Narasimhan type and let $\delta_\mu = d_1 - d_2 + 1$. The Hitchin morphism for the moduli spaces $\mathring{\mathcal{M}}^{\mu\text{-s}}_{2,d}$, $\mathring{\mathcal{M}}^{\mu,\delta_\mu,\text{indec}}_{2,d}$ and $\mathring{\mathcal{M}}^{\mu,\text{dec}}$ has fibres and images as follows:

(i) the image of $h : \mathring{\mathcal{M}}^{\mu\text{-s}}_{2,d} \to \mathcal{B}$ coincides with $\mathcal{B}' \setminus \{0\}$, and for each $-a \otimes a \in \mathcal{B}' \setminus \{0\}$, the fibre $h^{-1}(-a \otimes a)$ is isomorphic to the disjoint union $\mathbb{P}^{t-1} \sqcup \mathbb{P}^{t-1}$;

(ii) the image of $h : \mathring{\mathcal{M}}^{\mu,\delta_\mu,\text{indec}}_{2,d} \to \mathcal{B}$ is equal to $\{0\}$ and there is an isomorphism $h^{-1}(0) \cong \mathbb{P}^{d_1 - d_2 + t}$;

(iii) the Hitchin morphism h realises $\mathring{\mathcal{M}}^{\mu,\delta_\mu,\text{indec}}_{2,d}$ as a two-to-one branched cover of $\mathcal{B}'$, with a branched point over the origin.

Figure 4.2 illustrates the results of Proposition 4.1.12 above.

Remark 4.1.13 (Properness of the Hitchin fibration)**.** By Proposition 4.1.12 above we have that the fibres of the Hitchin morphism for $\mathring{\mathcal{M}}^{\mu\text{-s}}_{2,d}$ and $\mathring{\mathcal{M}}^{\mu,\delta_\mu,\text{indec}}_{2,d}$ are projective. We expect that the maps are in fact proper, and we hope to prove this in future work.

Proof of Proposition 4.1.12. To simplify notation, we let E denote the vector bundle $\mathcal{O}(d_1) \oplus \mathcal{O}(d_2)$. Suppose that $[(E, \phi)] \in \mathring{\mathcal{M}}^{\mu\text{-s}}_{2,d}$. Then by applying Proposition 4.1.2 (i) in the trace-free case, we know that ϕ must have diagonal given by $(a, -a)$ with $a \neq 0$, so that $h([(E, \phi)]) =$

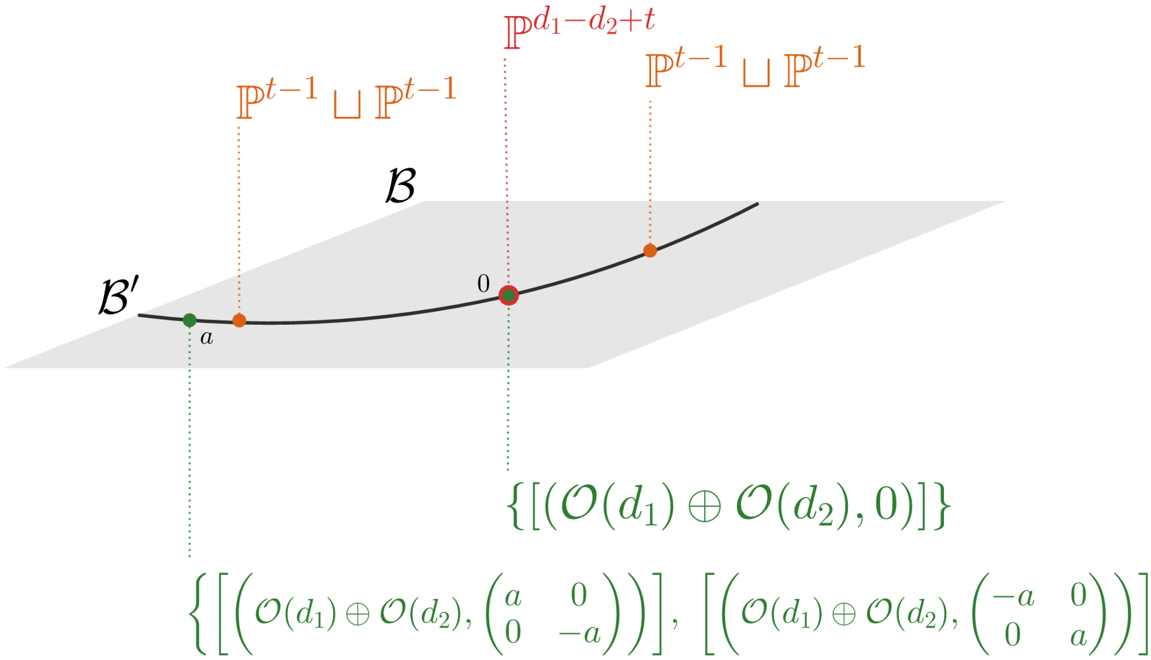

Figure 4.2: The Hitchin fibration for moduli spaces of unstable $\mathcal{O}(t)$-twisted rank 2 Higgs bundles on $\mathbb{P}^1$ with trace-free Higgs field. The horizontal grey plane represents the Hitchin base $\mathcal{B} = H^0(\mathbb{P}^1, \mathcal{O}(2t))$. The black line represents the subset $\mathcal{B}' \subseteq \mathcal{B}$ given by the image of the map $H^0(\mathbb{P}^1, \mathcal{O}(t)) \to H^0(\mathbb{P}^1, \mathcal{O}(2t))$ defined by $a \mapsto -a \otimes a$. The dotted lines represent the fibres of the Hitchin morphism for each of the three moduli spaces $\mathring{\mathcal{M}}_{2,d}^{\mu\text{-s}}$ (coloured in red), $\mathring{\mathcal{M}}_{2,d}^{\mu,\delta_\mu,\text{indec}}$ (coloured in orange) and $\mathring{\mathcal{M}}_{2,d}^{\mu,\text{dec}}$ (coloured in green). Under this map, we have: (i) the moduli space $\mathring{\mathcal{M}}_{2,d}^{\mu\text{-s}}$ is mapped onto $\mathcal{B}' \setminus \{0\}$ with fibre given by the disjoint union $\mathbb{P}^{t-1} \sqcup \mathbb{P}^{t-1}$; (ii) the moduli space $\mathring{\mathcal{M}}_{2,d}^{\mu,\delta_\mu,\text{indec}}$ is mapped onto $\{0\}$, and is isomorphic to $\mathbb{P}^{d_1-d_2+t}$; (iii) the moduli space $\mathring{\mathcal{M}}_{2,d}^{\mu,\text{dec}}$ is mapped onto the subset $\mathcal{B}' \subseteq \mathcal{B}$, with fibre over $\mathcal{B}' \setminus \{0\}$ consisting of two points, and fibre over $\{0\}$ consisting of a single point.

$\det(\phi) = -a \otimes a \in \mathcal{B}'$. Moreover, given $-a \otimes a \in \mathcal{B}' \setminus \{0\}$, the fibre of $h : \mathring{\mathcal{M}}_{2,d}^{\mu\text{-s}} \to \mathcal{B}'$ over $-a \otimes a$ consists of isomorphism classes of Higgs bundles of the form (E, ϕ) where $\phi = \begin{pmatrix} a & b \\ 0 & -a \end{pmatrix}$ or $\phi = \begin{pmatrix} -a & b \\ 0 & a \end{pmatrix}$ and the image of b in $\operatorname{coker}(2a)$ is non-zero. As seen in the proof of Proposition 4.1.2, the image of b in $\mathbb{P}(\operatorname{coker}(2a)) \cong \mathbb{P}^{t-1}$ (or equivalently in $\operatorname{coker}(-2a) \cong \mathbb{P}^{t-1}$) uniquely determines such an isomorphism class. Thus we can identify the fibre $h^{-1}(-a \otimes a)$ as a double cover of $\mathbb{P}^{t-1}$, with the preimage of any point in $\mathbb{P}^{t-1}$ consisting of the two choices for the top left entry of the Higgs field, a or $-a$. Since $a \neq 0$, the cover is unramified and therefore since it is a cover of $\mathbb{P}^{t-1}$ which is simply connected, the fibre is a disjoint union $\mathbb{P}^{t-1} \sqcup \mathbb{P}^{t-1}$.

Now suppose that $[(E, \phi)] \in \mathring{\mathcal{M}}_{2,d}^{\mu,\delta_\mu,\text{indec}}$. Then applying Proposition 4.1.2 (ii) in the trace-free case, we know that ϕ must have diagonal $(0,0)$, so that $h([(E,\phi)]) = \det\phi = 0 \in \mathcal{B}'$. The fibre of $h : \mathring{\mathcal{M}}_{2,d}^{\mu,\delta_\mu,\text{indec}} \to \mathcal{B}$ over $\{0\}$ is given by Higgs bundles of the form $[(E,\phi)]$ where $\phi = \begin{pmatrix} 0 & b \\ 0 & 0 \end{pmatrix}$ where $b \neq 0$. By the proof of Proposition 4.1.2 (ii), we know that the image of b in $\mathbb{P}(H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2 + t)))$ uniquely determines such an isomorphism class. Thus $h^{-1}(0) \cong \mathbb{P}^{d_1-d_2+t}$.

Finally, suppose that $[(E,\phi)] \in \mathring{\mathcal{M}}_{2,d}^{\mu,\text{dec}}$. Then the diagonal of ϕ is of the form $(a,-a)$ for any $a \in H^0(\mathbb{P}^1, \mathcal{O}(t))$. Thus $h : \mathring{\mathcal{M}}_{2,d}^{\mu,\text{dec}} \to \mathcal{B}$ has image equal to $\mathcal{B}'$, and for any $-a \otimes a \in \mathcal{B}'$, the fibre $h^{-1}(-a\otimes a)$ consists of the two isomorphism classes of Higgs bundles given by (E,ϕ) where $\phi = \begin{pmatrix} a & 0 \\ 0 & -a \end{pmatrix}$ or $\phi = \begin{pmatrix} -a & 0 \\ 0 & a \end{pmatrix}$, and these two isomorphism classes coincide if and only if $a = 0$. Thus $h^{-1}(-a\otimes a)$ can be identified as a branched double cover of $\mathcal{B}'$, with a branched point at 0. □

4.1.3 Comparison with the semistable case

We conclude this Section 4.1 by comparing the results obtained above regarding the geometry of the unstable refined Higgs Harder-Narasimhan strata and their coarse moduli spaces with the corresponding known results for the semistable stratum and its coarse moduli space.

Comparing dimensions. Combining Proposition 4.1.5 and Corollary 4.1.10 with known results from the semistable case (see [81] and [49] for example), we provide in Table 4.1 the dimensions of each of the refined Higgs Harder-Narasimhan strata appearing in the stack of $\mathcal{O}(t)$-twisted Higgs bundles on $\mathbb{P}^1$ of rank 2 and odd degree d, as well as the dimensions of their coarse moduli spaces. We can see from this table that just as for the semistable stratum, the dimensions of each of the refined strata for an unstable Higgs Harder-Narasimhan stratum depend only on the degree t of the twisting line bundle $\mathcal{O}(t)$. The largest of these refined strata parametrising unstable Higgs bundles is given by the substack of μ-stable Higgs bundles, and it has dimension t less than the semistable stratum.

Comparing the dimensions of the strata and of their coarse moduli spaces for the three unstable refined Higgs Harder-Narasimhan strata with the semistable case, we make the following three observations:

(i) The situation for the stratum $\mathscr{H}_{2,d}^{\mu\text{-s}}$ of μ-stable Higgs bundles is the same as that for the

Refined strata	Dimension	Coarse moduli space	Dimension
$\mathscr{H}_{2,d}^{ss}$	$4t$	$\mathcal{M}_{2,d}^{ss}$	$4t+1$
$\mathscr{H}_{2,d}^{\mu\text{-s}}$	$3t$	$\mathcal{M}_{2,d}^{\mu\text{-s}}$	$3t+1$
$\mathscr{H}_{2,d}^{\mu,\delta_\mu,\text{indec}}$	$2t-1$	$\mathcal{M}_{2,d}^{\mu,\delta_\mu,\text{indec}}$	$2t+d_1-d_2+1$
$\mathscr{H}_{2,d}^{\mu,\text{dec}}$	$2t$	$\mathscr{H}_{2,d}^{\mu,\text{dec}}$	$2t+2.$

Table 4.1: Dimensions of the refined Higgs Harder-Narasimhan strata and their coarse moduli spaces for the stack of rank 2 and odd degree d $\mathcal{O}(t)$-twisted Higgs bundles on $\mathbb{P}^1$. Here $\mu = (d_1, d_2)$ denotes a rank 2 and degree d unstable Higgs Harder-Narasimhan type and $\delta_\mu = d_1 - d_2 + 1$.

semistable stratum: the dimension of the coarse moduli space is one less than the dimension of the corresponding stratum. In the semistable case, this is because automorphisms of stable Higgs bundles are given by scalars (and since we have assumed that the degree d is odd, all semistable Higgs bundles are stable); as a result, the semistable stratum is a $\mathbb{G}_m$-gerbe over the moduli space of semistable Higgs bundles. The results of Sections 4.1.1 and 4.1.2 therefore show that this property extends to μ-stable Higgs bundles. That is, their automorphisms are also given by scalars, so that the stratum $\mathscr{H}_{2,d}^{\mu\text{-s}}$ is a $\mathbb{G}_m$-gerbe over the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}$. Thus μ-stability can be viewed as the analogue of stability when considering unstable Higgs bundles of Higgs Harder-Narasimhan type μ.

(ii) The above property regarding stabiliser groups fails for the stratum $\mathscr{H}_{2,d}^{\mu,\delta_\mu,\text{indec}}$, the coarse moduli space of which has dimension $2t + d_1 - d_2 + 1$ which increases as the slope of the destabilising subbundle increases. This is because indecomposable Higgs bundles (E, ϕ) of Higgs Harder-Narasimhan type μ and satisfying $\text{End}_{-1}(E, \phi) = \delta_\mu$ have a larger stabiliser group, given by a subgroup of

$$\text{Aut}(\mathcal{O}(d_1) \oplus \mathcal{O}(d_2)) \cong H^0(\mathbb{P}^1, \mathcal{O}(d_1 - d_2)) \rtimes (\mathbb{G}_m \times \mathbb{G}_m)$$

isomorphic to $H^0(\mathbb{P}^1, \mathcal{O}(d_1-d_2)) \rtimes \mathbb{G}_m$, with $\mathbb{G}_m$ acting non-trivially on $H^0(\mathbb{P}^1, \mathcal{O}(d_1-d_2))$. As a result, $\mathcal{M}_{2,d}^{\mu,\delta_\mu,\text{indec}}$ has larger dimension than expected.

(iii) The automorphism group of decomposable Higgs bundles consists of the product of the automorphism group of Higgs line bundles (given by scalars), and thus as expected the dimension of the moduli space of decomposable Higgs bundles of Higgs Harder-Narasimhan type μ is two less than the moduli space of the corresponding stratum.

Comparing the Hitchin fibrations. We now compare the behaviour of the Hitchin morphism in the semistable case to the behaviour in the unstable case established in Section 4.1.2 above. It follows from general results concerning the moduli space of semistable Higgs bundles (of any rank and over any smooth projective curve) that the Hitchin morphism[7]

$$h : \mathcal{M}^{ss}_{2,d} \to H^0(\mathbb{P}^1, \mathcal{O}(t)) \oplus H^0(\mathbb{P}^1, \mathcal{O}(2t))$$

given by $[(E, \phi)] \mapsto \operatorname{char}\phi$ where $\operatorname{char}\phi$ denotes the characteristic polynomial of ϕ, is proper (see [81] for example for the case of L-twisted Higgs bundles on a smooth projective curve). Thus in particular its fibres are compact. Moreover, the generic fibre of h is given by the Jacobian variety over the smooth spectral curve associated to the fixed characteristic polynomial[8] and contained in the total space of $\mathcal{O}(t)$. By [74, §2], the arithmetic genus of such a spectral curve is given by $t-1$, so that the fibre is generically a projective variety of dimension $t-1$. The fibres which are singular correspond to points in the Hitchin base which determine a singular spectral curve. The most special singular case is the fibre over 0, called the nilpotent cone. The nilpotent cone has been widely studied as it contains the fixed point locus for the Higgs field $\mathbb{C}^*$-action, and as a result encodes the topology of the moduli space (see for example [86] in the case of Higgs bundles on $\mathbb{P}^1$). The other singular fibres are given by non-zero elements of the Hitchin base which give rise to a singular spectral curve; in the trace-free rank 2 case these correspond to elements $\det\phi$ in the base with at least one repeated zero. The fibres in this case have been studied in detail in [36], which separates the study into two cases: when the spectral curve is irreducible, it is described as the compactification by rank 1 torsion-free sheaves of the Prym variety associated to the spectral curve (see [36, Thm 6.1]), and when the spectral curve is reducible, the fibre is described as a stratified space (see [36, Thm 7.7]).

As we have seen in Section 4.1.2, the moduli spaces $\mathring{\mathcal{M}}^{\mu\text{-s}}_{2,d}$, $\mathring{\mathcal{M}}^{\mu,\delta_\mu,\text{indec}}_{2,d}$ and $\mathring{\mathcal{M}}^{\mu,\text{dec}}_{2,d}$ also admit a Hitchin morphism to the Hitchin base $\mathcal{B} = H^0(\mathbb{P}^1, \mathcal{O}(2t))$, with compact fibres as in the semistable case. Nevertheless, unlike the semistable case, the morphism is not surjective. Rather, if we let $\mathcal{B}'$ denote the subset of $\mathcal{B}$ given by the image of $H^0(\mathbb{P}^1, \mathcal{O}(t))$ in $H^0(\mathbb{P}^1, \mathcal{O}(2t))$

[7]We can also consider this morphism in the trace-free case, in which case the image of the Hitchin morphism is simply given by $H^0(\mathbb{P}^1, \mathcal{O}(2t))$.

[8]Given $\pi : \operatorname{Tot}(\mathcal{O}(t)) \to \mathbb{P}^1$ the projection from the total space of $\mathcal{O}(t)$, and η the tautological section of the line bundle $\pi^*\mathcal{O}(2t)$ on $\operatorname{Tot}(\mathcal{O}(t))$, then the spectral curve associated to a characteristic polynomial with coefficients $(a_2, a_2) \in H^0(\mathbb{P}^1, \mathcal{O}(t)) \oplus H^0(\mathbb{P}^1, \mathcal{O}(2t))$ is given by the vanishing locus in $\operatorname{Tot}(\mathcal{O}(t))$ of the section $\eta^2 + \pi^* a_1 \eta + \pi^* a_2$ of $\pi^*(\mathcal{O}(2t))$.

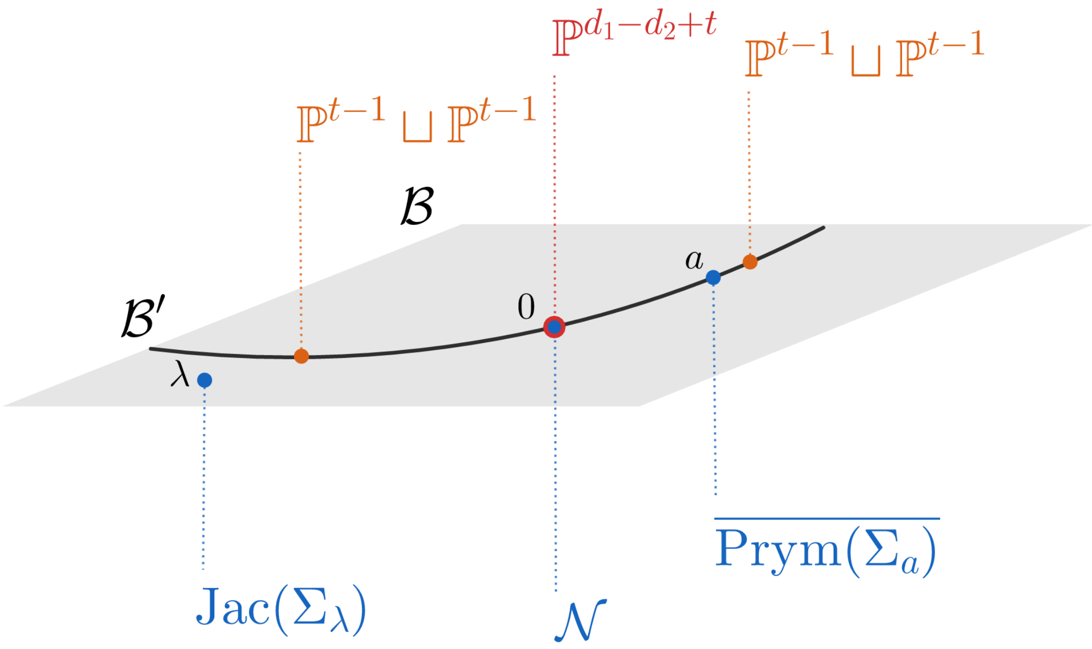

Figure 4.3: The Hitchin fibration for moduli spaces of $\mathcal{O}(t)$-twisted rank 2 Higgs bundles on $\mathbb{P}^1$ with trace-free Higgs field. The horizontal grey plane represents the Hitchin base $\mathcal{B} = H^0(\mathbb{P}^1, \mathcal{O}(2t))$. The black curve represents the subset $\mathcal{B}' \subseteq \mathcal{B}$ given by the image of the map $H^0(\mathbb{P}^1, \mathcal{O}(t)) \to H^0(\mathbb{P}^1, \mathcal{O}(2t))$ defined by $a \mapsto -a \otimes a$. The dotted lines represent the fibres of the Hitchin morphism for each of the three moduli spaces $\mathring{\mathcal{M}}^{ss}_{2,d}$ (coloured in blue), $\mathring{\mathcal{M}}^{\mu\text{-s}}_{2,d}$ (coloured in orange) and $\mathring{\mathcal{M}}^{\mu,\delta_\mu,\text{indec}}_{2,d}$ (coloured in red). Here $\mu = (d_1, d_2)$ denotes an unstable Higgs Harder-Narasimhan type of rank 2 and degree d, and we set $\delta_\mu = d_1 - d_2 + 1$. The Hitchin morphism is surjective onto $\mathcal{B}$ for the moduli space $\mathring{\mathcal{M}}^{ss}_{2,d}$: the fibre over a generic point λ (i.e. a characteristic polynomial λ corresponding to a smooth spectral curve Σ_λ) is given by the Jacobian $\mathrm{Jac}(\Sigma_\lambda)$ of the spectral curve Σ_λ; the fibre over a point a in $\mathcal{B}' \setminus \{0\}$ is isomorphic to a compactification $\overline{\mathrm{Prym}(\Sigma_a)}$ of the Prym variety over the spectral curve Σ_a; and the fibre over 0 is the nilpotent cone $\mathcal{N}$, studied in [86]). For the moduli spaces of unstable Higgs bundles, under the Hitchin morphism we have: (i) the moduli space $\mathring{\mathcal{M}}^{\mu\text{-s}}_{2,d}$ is mapped onto $\mathcal{B}' \setminus \{0\}$ with fibre given by the disjoint union $\mathbb{P}^{t-1} \sqcup \mathbb{P}^{t-1}$; (ii) the moduli space $\mathring{\mathcal{M}}^{\mu,\delta_\mu,\text{indec}}_{2,d}$ is mapped onto $\{0\}$, and is isomorphic to $\mathbb{P}^{d_1-d_2+t}$.

under the inclusion $a \mapsto -a \otimes a$, then the images of the Hitchin morphism for each of these three moduli spaces are given by $\mathcal{B}' \setminus \{0\}$, $\{0\}$ and $\mathcal{B}'$ respectively.

We note that each point in $b \in \mathcal{B}'$ corresponds to a singular spectral curve, which is irreducible if $b \neq 0$, and reducible if $b = 0$ (see [36, Rk 3.2]). Thus the fibres of the Hitchin morphism for the moduli space $\mathring{\mathcal{M}}^{ss}_{2,d}$ are singular for each $b \in \mathcal{B}'$. In contrast, by Proposition 4.1.12 the fibres of the Hitchin morphism for the moduli spaces parametrising unstable Higgs bundles are all smooth (if non-empty) over such points.

Figure 4.3 depicts the different fibres over the Hitchin base $\mathcal{B}$ for the three moduli spaces $\mathring{\mathcal{M}}^{ss}_{2,d}$, $\mathring{\mathcal{M}}^{\mu\text{-s}}_{2,d}$ and $\mathring{\mathcal{M}}^{\mu,\delta_\mu,\text{indec}}$. We omit the moduli space $\mathring{\mathcal{M}}^{\mu,\text{dec}}_{2,d}$, as its non-empty fibres consist of a finite number of points (see Proposition 4.1.12 and Figure 4.2).

4.2 Moduli spaces for 'μ-unstable' Higgs bundles on a curve of arbitrary genus

Section 4.1 above focused on describing the geometry of moduli spaces for rank 2 unstable Higgs bundles on $\mathbb{P}^1$. The appeal of working over $\mathbb{P}^1$ is that there are no moduli for the underlying bundle, and as a result the moduli spaces essentially reduce to moduli spaces of Higgs fields which can be described explicitly. In this section and the next we consider rank 2 unstable Higgs bundles on a curve of arbitrary genus. We have constructed moduli spaces for such Higgs bundles in Chapter 3 using Non-Reductive GIT. They are the moduli spaces $\mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}(\Sigma, L)$, for $\delta \geq 0$, which parametrise indecomposable Higgs bundles of Higgs Harder-Narasimhan type μ with a fixed dimension δ of 'backward endomorphisms' (see Theorem 3.2.2). The aim of this section is to study the geometry of these moduli spaces when $\delta > 0$.

In Section 4.2.1 we study the Higgs field scaling action on the moduli spaces $\mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}(\Sigma, L)$, bringing to light the differences between the $\delta > 0$ case (where limits at 0 always exist), and the $\delta = 0$ case (where limits at 0 may not always exist). In Section 4.2.2 we fix $L = K_\Sigma$ and study the geometry of the fixed point set (see Proposition 4.2.9). This enables us in Section 4.2.3 to describe the geometry of the moduli space $\mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}(\Sigma, L)$ when $\delta \neq 0$: we describe its Bialynicki-Birula stratification (see Theorem 4.2.12) and obtain a formula for the Poincaré series of the moduli space in terms of the Poincaré series of a Brill-Noether locus in the Picard variety $\mathrm{Pic}^{d_2-d_1}(\Sigma)$ and of the Picard variety $\mathrm{Pic}^{d_1}(\Sigma)$ where $\mu = (d_1, d_2)$ (see Corollary 4.2.13).

4.2.1 Limits under the $\mathbb{C}^*$-action

The aim of this section is first to describe the points in the moduli spaces $\mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}(\Sigma, L)$ fixed by the Higgs field scaling $\mathbb{C}^*$-action, and then to study the limits at 0 under the $\mathbb{C}^*$-action in these moduli spaces. More precisely, for $\delta \neq 0$ we show that the moduli space $\mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}(\Sigma, L)$ satisfies the first condition of semiprojectivity, namely that limits at 0 exist (see Proposition 4.2.2). This property will enable us to study the geometry of these moduli spaces in Section 4.2.3 below. For $\delta = 0$, we show that even the second condition, namely that the fixed point set is projective, can fail (see Proposition 4.2.6).

The first step for achieving these results is to describe Higgs bundles which are fixed points for the $\mathbb{C}^*$-action. Higgs bundles fixed by the $\mathbb{C}^*$-action have been extensively studied thanks in particular to the important role they play in describing the cohomology of the moduli space of semistable Higgs bundles. Higgs bundles which are fixed by the $\mathbb{C}^*$-action can be described as holomorphic chains (see [93, Lem 4.1] or [87, §3.1] for example). In the rank 2 case, semistable Higgs bundles which are fixed by the $\mathbb{C}^*$-action, and the correspondence with holomorphic chains, are particularly simple to describe (see [51, §7] or [35, §3] for example). That is, if (E, ϕ) is a rank 2 semistable Higgs bundle, then (E, ϕ) is fixed by $\mathbb{C}^*$ if and only if one of the following two conditions is satisfied:

(i) the Higgs field ϕ is zero, so that $(E, \phi) = (E, 0)$;

(ii) the underlying bundle E is unstable with $E \cong \operatorname{gr} E = E_1 \oplus E/E_1$, and with respect to this identification ϕ is of the form $\begin{pmatrix} 0 & 0 \\ \phi_{21} & 0 \end{pmatrix}$ where $\phi_{21} : E^1 \to E/E^1 \otimes L$ is the composition of the restriction of ϕ to E^1 with the quotient map $E \otimes L \to E/E^1 \otimes L$, so that

$$(E, \phi) \cong \left(E_1 \oplus E/E_1, \begin{pmatrix} 0 & 0 \\ \phi_{21} & 0 \end{pmatrix} \right).$$

This characterisation of fixed points and the correspondence with holomorphic chains in the case of semistable Higgs bundles of rank 2 can be extended to the case of unstable Higgs bundles of rank 2, as per Proposition 4.2.1 below.

Proposition 4.2.1 ($\mathbb{C}^*$-fixed points for unstable rank 2 Higgs bundles)**.** Let μ denote an unstable rank 2 Higgs Harder-Narasimhan type. Then a Higgs bundle (E, ϕ) is fixed by the Higgs field scaling $\mathbb{C}^*$-action if and only if one of the following two conditions is satisfied:

(i) the Higgs field ϕ is zero, in which case $(E, \phi) = (E, 0)$;

(ii) the underlying bundle E satisfies $E \cong \operatorname{gr} E = E_1 \oplus E/E^1$, and the image of the Higgs field is contained in $E_1 \otimes L \subseteq E \otimes L$ and restricts to zero on the subbundle E_1, giving an induced map $\phi_{12} : E/E_1 \to E_1 \otimes L$, in which case

$$(E, \phi) \cong \left(E_1 \oplus E/E_1, \begin{pmatrix} 0 & \phi_{12} \\ 0 & 0 \end{pmatrix} \right).$$

Proof. Let (E, ϕ) denote an unstable rank 2 Higgs bundle of Higgs Harder-Narasimhan type μ and suppose that it is fixed by the Higgs field scaling $\mathbb{C}^*$-action. By Proposition 3.1.8, since we are considering the rank 2 case and (E, ϕ) has Higgs Harder-Narasimhan type μ, then the Harder-Narasimhan type of E is also μ; we let $\operatorname{gr} E = E_1 \oplus E_2$. Since (E, ϕ) is fixed, the bundle E decomposes as a direct sum of eigenspaces for the limiting endomorphism of the family $\{\psi_t\}$ of automorphisms of E. Thus we can write $E = F_1 \oplus F_2$ where the F_i are eigenspaces (with F_2 trivial if E is itself an eigenspace). If $E = F_1$, then ϕ must be the zero map and thus $(E, \phi) = (E, 0)$.

Suppose then that $F_2 \neq \{0\}$. Since ψ_t must preserve the destabilising subbundle E_1, we must have that E_1 is contained in one of these eigenspaces. Without loss of generality we can therefore suppose that $E_1 = F_1$, so that $E_2 \cong F_2$; we fix this isomorphism and identify the two line bundles. Then ϕ either maps E_1 to $E_2 \otimes L$ and $E_2 \otimes L$ to 0, or E_2 to $E_1 \otimes L$ and $E_1 \otimes L$ to 0. Since by assumption (E, ϕ) is unstable, the destabilising subbundle E_1 must be preserved by ϕ. Thus we must be in the second case, and we let ϕ_{12} denote the map $E_2 \to E_1 \otimes L$ (noting that ϕ_{12} is induced from ϕ since the restriction of ϕ to E_1 is zero and the image of ϕ is contained in E_1). As a result, we have that $(E, \phi) \cong \left(E_1 \oplus E_2, \begin{pmatrix} 0 & \phi_{12} \\ 0 & 0 \end{pmatrix}\right)$. □

The two conditions required for a variety to be semiprojective with respect to a $\mathbb{C}^*$-action are firstly that limits at 0 exist, and secondly that the fixed point set is projective. For the remainder of this section we study the first condition of semiprojectivity for the moduli spaces $\mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}(\Sigma, L)$ (we will consider the second condition in Section 4.2.3 below, see Proposition 4.2.9). The behaviour is different according to whether δ is zero or strictly positive. We start with the case where δ is strictly positive.

Proposition 4.2.2 (Limits at 0 under the $\mathbb{C}^*$-action in $\mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}(\Sigma, L)$ when $\delta > 0$)**.** Let $[(E, \phi)] \in \mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}(\Sigma, L)$ where $\delta > 0$, and let $\left(E_1 \oplus E_2, \begin{pmatrix} \phi_{11} & 0 \\ 0 & \phi_{22} \end{pmatrix}\right)$ denote the associated graded of (E, ϕ). Then the limit $\lim_{t\to 0}[(E, t\phi)]$ exists. Moreover, we have:

(i) if $E \ncong \operatorname{gr} E$, then $\lim_{t\to 0}[(E, t\phi)] = [(E, 0)]$;

(ii) if $E \cong \operatorname{gr} E = E_1 \oplus E_2$, then $\lim_{t\to 0}[(E, t\phi)] = \left[\left(E_1 \oplus E_2, \begin{pmatrix} 0 & \phi_{12} \\ 0 & 0 \end{pmatrix}\right)\right]$ where $\phi_{12} : E_2 \to E_1 \otimes L$ is induced by ϕ.

Remark 4.2.3 (Semiprojectivity)**.** The above Proposition 4.2.2 implies that the first condition of semiprojectivity, namely that limits at 0 exist under the $\mathbb{C}^*$-action, is satisfied by the moduli space $\mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}(\Sigma, L)$ when $\delta > 0$ (provided the moduli space is non-empty). We will see in

Proposition 4.2.9 that the second condition of semiprojectivity, namely that the fixed points set is projective, fails in general for such moduli spaces. Nevertheless, the existence of limits at 0 suffices to provide a Bialynicki-Birula stratification for the moduli space, which can be used to study the geometry of the moduli space in terms of the geometry of the fixed point set, even if non-projective. This is the approach we take in Section 4.2.3 below.

The proof of Proposition 4.2.2 relies on the following

Lemma 4.2.4. Suppose that (E, ϕ) is a rank 2 unstable Higgs bundle of Higgs Harder-Narasimhan type μ, and let $\left(E_1 \oplus E_2, \begin{pmatrix} \phi_{11} & 0 \\ 0 & \phi_{22} \end{pmatrix}\right)$ denote its associated graded. Then we have:

(i) $\operatorname{End}_{-1}(E, \phi) \cong \operatorname{End}_{-1}(\operatorname{gr}(E, \phi))$, and similarly $\operatorname{End}_{-1}(E) \cong \operatorname{End}_{-1}(\operatorname{gr} E)$;

(ii) if $\dim \operatorname{End}_{-1}(E, \phi) = \delta > 0$, then $\phi_{11} = \phi_{22}$ and $\operatorname{End}_{-1}(E, \phi) \cong \operatorname{End}_{-1}(E)$, so that $\dim \operatorname{End}_{-1}(E) = \delta$.

Remark 4.2.5. Lemma 4.2.4 is necessary to ensure that the limit given in part (ii) of Proposition 4.2.2 is well-defined. Indeed, the map

$$\left[\left(E_1 \oplus E_2, \begin{pmatrix} \phi_{11} & \phi_{12} \\ 0 & \phi_{22} \end{pmatrix}\right)\right] \mapsto \left[\left(E_1 \oplus E_2, \begin{pmatrix} 0 & \phi_{12} \\ 0 & 0 \end{pmatrix}\right)\right]$$

is well-defined if and only if $\phi_{11} = \phi_{22}$ (see (4.1) in Section 4.1.1).

Proof of Lemma 4.2.4. We prove part (i) first. Let $0 = E^0 \subseteq E^1 \subseteq E^2 = E$ denote the Higgs Harder-Narasimhan filtration of (E, ϕ), and fix an isomorphism $\operatorname{gr}(E, \phi) \cong (E_1, \phi_1) \oplus (E_2, \phi_2)$ where $E_1 = E^1$ and $E_2 \cong E/E^1$. Let $\psi \in \operatorname{End}_{-1}(E, \phi)$; note that its image is contained in E^1. Since ψ is zero on $E^1 = E_1$, it factors through the quotient map $\pi : E \to E_2$ (let $\overline{\psi}$ denote the induced map $E_2 \to E$), and we have the following diagram:

$$\begin{array}{ccccc} E & \xrightarrow{\pi} & E_2 & \xrightarrow{\overline{\psi}} & E_1 \\ \downarrow{\scriptstyle \phi} & & \downarrow{\scriptstyle \phi_2} & & \downarrow{\scriptstyle \phi_1} \\ E \otimes L & \xrightarrow{\pi \otimes \mathrm{id}_L} & E_2 \otimes L & \xrightarrow{\overline{\psi}} & E_1 \otimes L. \end{array}$$

The outer square commutes since by assumption $\psi \in \operatorname{End}_{-1}(E, \phi)$. The left square commutes by definition of the map ϕ_2. Since the quotient map $\pi : E \to E_2$ is surjective, it follows that the right square also commutes. Thus we obtain a nilpotent endomorphism $\chi \in \operatorname{End}_{-1}(\operatorname{gr}(E, \phi))$ by setting $\chi = 0_{E_1} \oplus \overline{\psi}$. Conversely, given $\chi \in \operatorname{End}_{-1}(\operatorname{gr}(E, \phi))$, we can define $\psi \in \operatorname{End}_{-1}(E)$

by setting $\psi = \chi|_{E_2} \circ \pi$. Then $\overline{\psi} = \chi|_{E_2}$ and since both the left and right squares in the above diagram commute, it follows that the outer square commutes. Thus ψ commutes with the Higgs field so that $\psi \in \operatorname{End}_{-1}(E, \phi)$. This correspondence $\psi \leftrightarrow \chi$ defines the desired isomorphism, and proves part (i).

To prove part (ii), suppose that $\dim \operatorname{End}_{-1}(E, \phi) = \delta > 0$ and let $(E_1 \oplus E_2, \phi_{11} \oplus \phi_{22})$ denote its associated Higgs Harder-Narasimhan graded. By Lemma 4.2.4 (i), we have that $\dim \operatorname{End}_{-1}(E, \phi) = \dim \operatorname{End}_{-1} \operatorname{gr}(E, \phi)$. If $\phi_{11} \neq \phi_{22}$, then by the proof of Proposition 3.2.12, we have that $\operatorname{End}_{-1}(\operatorname{gr}(E, \phi)) = \{0\}$. Since $\delta \neq 0$ by assumption, we must therefore have that $\phi_{11} = \phi_{22}$.

If $\phi_{11} = \phi_{22}$, then any $\psi \in \operatorname{End}_{-1}(E_1 \oplus E_2)$ commutes with ϕ and thus gives an element of $\operatorname{End}_{-1}(\operatorname{gr}(E, \phi))$. Conversely, any element of $\operatorname{End}_{-1}(\operatorname{gr}(E, \phi))$ clearly gives an element of $\operatorname{End}_{-1}(E_1 \oplus E_2)$. Thus $\operatorname{End}_{-1}(E, \phi) = \operatorname{End}_{-1}(E_1 \oplus E_2)$, and since $\operatorname{End}_{-1}(E_1 \oplus E_2) \cong \operatorname{End}_{-1}(E)$, we obtain the desired isomorphism $\operatorname{End}_{-1}(E, \phi) \cong \operatorname{End}_{-1}(E)$. This proves part (ii). □

We can now prove Proposition 4.2.2.

Proof of Proposition 4.2.2. There are two cases to consider, when $E \ncong \operatorname{gr} E$ and when $E \cong \operatorname{gr} E$. We start with the first case, and hence suppose that $E \ncong \operatorname{gr} E$. To see that $[(E, 0)]$ is the limit of $[(E, t\phi)]$ under the $\mathbb{C}^*$-action, it suffices to show that $[(E, 0)] \in \mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}(\Sigma, L)$. Since (E, ϕ) has Higgs Harder-Narasimhan type μ, by Proposition 3.1.8 we have that E has Harder-Narasimhan type μ, and thus $[(E, 0)]$ has Higgs Harder-Narasimhan type μ. Moreover, since $E \ncong \operatorname{gr} E$, we have that $(E, 0) \ncong \operatorname{gr}(E, 0)$. Finally, we must show that $\dim \operatorname{End}_{-1}(E, 0) = \delta$. In order to do so, since $\operatorname{End}_{-1}(E, 0) = \operatorname{End}_{-1}(E)$ it suffices to show that the latter has dimension δ. By Lemma 4.2.4 (i) applied to Higgs bundles with zero Higgs field, we have that $\operatorname{End}_{-1}(E) \cong \operatorname{End}_{-1}(\operatorname{gr} E)$ and thus we need only show that $\dim \operatorname{End}_{-1}(E_1 \oplus E_2) = \delta$.

By Lemma 4.2.4 (ii) we have that $\operatorname{End}_{-1}(E, \phi) = \operatorname{End}_{-1}(E_1 \oplus E_2)$, and so $\dim \operatorname{End}_{-1}(E_1 \oplus E_2) = \delta$. Thus $\dim \operatorname{End}_{-1}(E, 0) = \delta$, and so $[(E, 0)] \in \mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}(\Sigma, L)$. Hence we have that if $E \ncong \operatorname{gr} E$, then $\lim_{t \to 0}[(E, t\phi)] = [(E, 0)]$ as required.

We now consider the second case, and hence suppose that $E \cong \operatorname{gr}(E)$. By Proposition 4.2.2 (ii) we can write $[(E, \phi)]$ as $\left[\left(E_1 \oplus E_2, \begin{pmatrix} \phi_{11} & \phi_{12} \\ 0 & \phi_{22} \end{pmatrix}\right)\right]$, where $\phi_{12} \neq 0$ since $[(E, \phi)] \in \mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}(\Sigma, L)$ by assumption and therefore $(E, \phi) \ncong \operatorname{gr}(E, \phi)$. Moreover, by Lemma 4.2.4 (ii) we have that $\phi_{11} = \phi_{22}$. For each $t \in \mathbb{G}_m$, we can consider the automorphism ψ_t of $E_1 \oplus E_2$ given by $\psi_t = \begin{pmatrix} 1 & 0 \\ 0 & t \end{pmatrix}$, and observe that $\psi_t \circ t\phi \circ \psi_t^{-1} = \begin{pmatrix} t\phi_{11} & \phi_{12} \\ 0 & t\phi_{22} \end{pmatrix}$. Thus we have:

$$\left(E_1 \oplus E_2, \begin{pmatrix} t\phi_{11} & t\phi_{12} \\ 0 & t\phi_{22} \end{pmatrix}\right) \cong \left(E_1 \oplus E_2, \begin{pmatrix} t\phi_{11} & \phi_{12} \\ 0 & t\phi_{22} \end{pmatrix}\right).$$

We note that since $\phi_{11} = \phi_{22}$, the isomorphism class $\left[\left(E_1 \oplus E_2, \begin{pmatrix} 0 & \phi_{12} \\ 0 & 0 \end{pmatrix}\right)\right]$ is uniquely determined by the isomorphism class $[(E, \phi)]$ (see Remark 4.2.5).

To show that the limit of $[(E, t\phi)]$ as t tends to 0 is indeed the Higgs bundle $[(E_0, \phi_0)] := \left[\left(E_1 \oplus E_2, \begin{pmatrix} 0 & \phi_{12} \\ 0 & 0 \end{pmatrix}\right)\right]$, it suffices to show that $[(E_0, \phi_0)] \in \mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}$. First, we note that since $E_0 = E_1 \oplus E_2$ has Harder-Narasimhan type μ, then $[(E_0, \phi_0)]$ has Higgs Harder-Narasimhan type also equal to μ because ϕ_0 preserves E_1. Moreover, since $\phi_{12} \neq 0$, we have that $[(E_0, \phi_0)] \ncong [\text{gr}(E_0, \phi_0)]$. Finally, we must show that $\dim \text{End}_{-1}(E_0, \phi_0) = \delta$. By Lemma 4.2.4 (i) we have that $\dim \text{End}_{-1}(E_0, \phi_0) = \dim \text{End}_{-1}(\text{gr}(E_0, \phi_0))$ and since $\text{gr}(E_0, \phi_0) = (E_1 \oplus E_2, 0)$, it suffices to show that $\dim \text{End}_{-1}(E_1 \oplus E_2) = \delta$. Since $[(E, \phi)] \in \mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}(\Sigma, L)$, by Lemma 4.2.4 we have that $\delta = \dim \text{End}_{-1}(E, \phi) = \dim \text{End}_{-1}(E)$. Thus $\dim \text{End}_{-1}(E_0, \phi_0) = \delta$ and so $[(E_0, \phi_0)] \in \mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}(\Sigma, L)$. As a result we obtain that $\lim_{t \to 0}[(E, \phi)] = [(E_0, \phi_0)]$ as required. □

We now turn to the $\delta = 0$ case, which corresponds to the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}(\Sigma, L)$ of μ-stable rank 2 Higgs bundles of Higgs Harder-Narasimhan type μ.

Proposition 4.2.6 (Limits at 0 for the $\mathbb{C}^*$-action on $\mathcal{M}_{2,d}^{\mu\text{-s}}(\Sigma, L)$)**.** Given $[(E, \phi)] \in \mathcal{M}_{2,d}^{\mu\text{-}s}(\Sigma, L)$, we have:

(i) if $\dim \text{End}_{-1}(E) \neq 0$ then the limit of $[(E, t\phi)]$ as t tends to zero does not exist;

(ii) if $E \ncong (\text{gr}\, E)$ and $\dim \text{End}_{-1}(\text{gr}\, E) = 0$ then the limit $\lim_{t \to 0}[(E, t\phi)]$ exists and equals $[(E, 0)]$;

(iii) if $E \cong \text{gr}\, E = E_1 \oplus E_2$ (so that we can write ϕ as a matrix (ϕ_{ij})) and if $\dim \text{End}_{-1}(\text{gr}\, E) = 0$, then the limit $\lim_{t \to 0}[(E, t\phi)]$ exists and equals $\left[\left(E_1 \oplus E_2, \begin{pmatrix} 0 & \phi_{12} \\ 0 & 0 \end{pmatrix}\right)\right]$.

Remark 4.2.7 (Semiprojectivity)**.** As we have seen in Proposition 4.2.2, the moduli spaces $\mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}(\Sigma, L)$ for $\delta > 0$ always satisfy the first of the two conditions required to be semiprojective with respect to the $\mathbb{C}^*$-action, namely that limits at 0 exist. By contrast, Proposition 4.2.6 shows that even the first condition can fail for the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}(\Sigma, L) = \mathcal{M}_{2,d}^{\mu,0,\text{indec}}(\Sigma, L)$.

Indeed, only in very special cases will the moduli space $\mathcal{M}^{\mu\text{-s}}_{2,d}(\Sigma, L)$ satisfy the property that the underlying bundle E of any Higgs bundle in the moduli space satisfies $\dim \mathrm{End}_{-1}(E) = 0$. For example, any vector bundle E given by an extension $0 \to E_1 \to E \to E_2 \to 0$ where $E_2^\vee \otimes E_1 \cong \mathcal{O}(D)$ for D an effective divisor on Σ will satisfy $\mathrm{End}_{-1}(E) \neq \{0\}$.

Proof. Parts (ii) and (iii) of Proposition 4.2.6 can be proved in exactly the same way as Proposition 4.2.2, since the condition that $\dim \mathrm{End}_{-1}(E) = 0$ ensures that $0 = \dim \mathrm{End}_{-1}(E, \phi) = \dim \mathrm{End}_{-1}(E)$ for any $[(E, \phi)] \in \mathcal{M}^{\mu\text{-s}}_{2,d}(\Sigma, L)$.

Thus we need only prove part (i). Let $[(E, \phi)] \in \mathcal{M}^{\mu\text{-s}}_{2,d}(\Sigma, L)$ such that $\dim \mathrm{End}_{-1}(\mathrm{gr}\, E) = \delta > 0$. Suppose, in order to reach a contradiction, that a limit $[(E_0, \phi_0)]$ exists in $\mathcal{M}^{\mu\text{-s}}_{2,d}(\Sigma, L)$. Then either $E_0 \ncong \mathrm{gr}\, E_0 := E_1^0 \oplus E_2^0$ and $\phi_0 = 0$, or $E_0 \cong \mathrm{gr}\, E_0 = E_1^0 \oplus E_2^0$ and ϕ_0 is of the form $\begin{pmatrix} 0 & \phi_{12} \\ 0 & 0 \end{pmatrix}$. If the first case holds, then we must have that $\dim \mathrm{End}_{-1}(\mathrm{gr}\, E_0) = 0$ since $[(E_0, 0)]$ lies in $\mathcal{M}^{\mu\text{-s}}_{2,d}(\Sigma, L)$. We now show that this leads to a contradiction.

Suppose that $\mu = (d_1, d_2)$. We can consider the map from the moduli space $\mathcal{M}^{\mu\text{-s}}_{2,d}(\Sigma, L)$ to the moduli space $\mathrm{Pic}^{d_1 - d_2}(\Sigma)$ of line bundles on Σ of degree $d_1 - d_2$ given by sending a Higgs bundle $[(E, \phi)]$ to the line bundle $E_2^\vee \otimes E_1$ where $\mathrm{gr}\, E = E_1 \oplus E_2$. The family $\{[(E, t\phi)]\}$ of Higgs bundles in $\mathcal{M}^{\mu\text{-s}}_{2,d}(\Sigma, L)$ parametrised by $t \in \mathbb{C}^*$ gives a corresponding family of line bundles in $\mathrm{Pic}^{d_1 - d_2}(\Sigma)$, each with dimension of global sections strictly greater than 0. Since by assumption the family $\{[(E, t\phi)]\}$ tends to $[(E_0, \phi_0)]$ as $t \to 0$, the corresponding family in $\mathrm{Pic}^{d_1 - d_2}(\Sigma)$ tends to $(E_2^0)^\vee \otimes E_1^0$ as $t \to 0$, since $\mathrm{gr}\, E_0 = E_1^0 \oplus E_2^0$. Using the fact that the dimension of $h^0(\Sigma, M)$ for $M \in \mathrm{Pic}^{d_1 - d_2}(\Sigma)$ is an upper semi-continuous function, we must have that $h^0(\Sigma, (E_2^0)^\vee \otimes E_1^0) > 0$, which contradicts the fact established above that $\dim \mathrm{End}_{-1}(\mathrm{gr}\, E_0) = 0$, since $\mathrm{End}_{-1}(\mathrm{gr}\, E_0) \cong H^0(\Sigma, (E_2^0)^\vee \otimes E_1^0)$.

The same contradiction can be reached if the second case holds. Indeed, in this case we also have that $\dim \mathrm{End}_{-1}(E_0, \phi_0) = \dim \mathrm{End}_{-1}(\mathrm{gr}\, E_0) = 0$ since $[(E_0, 0)]$ lies in $\mathcal{M}^{\mu\text{-s}}_{2,d}(\Sigma, L)$. And as before, this is impossible since by assumption $\dim \mathrm{End}_{-1}(\mathrm{gr}\, E) = \delta > 0$. □

4.2.2 Geometry of the fixed point set

In this section we consider Higgs bundles in the classical sense, namely when $L = K_\Sigma$ over a curve Σ of genus $g \geq 1$. To simplify notation, we omit the data corresponding to Σ and to L when referring to the stack and the moduli spaces. The aim of this section is to study the geometry of the moduli spaces $\mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}$ when $\delta > 0$, using the results of Section 4.2.1 above regarding the Higgs field scaling $\mathbb{C}^*$-action. We use the existence of limits under this action at 0 to construct the Bialynicki-Birula stratification of $\mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}$, and show that under suitable assumptions on g, μ and δ the strata are irreducible and smooth (Theorem 4.2.12). We then show how these results can be used to provide a formula for the Poincaré series of $\mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}$. Our reason for specialising to the case where $L = K_\Sigma$ is that Serre duality can then be used to simplify the calculation of the dimensions of the strata; nevertheless we expect similar results to hold for an arbitrary line bundle L satisfying $h^0(\Sigma, L) > 0$.

The assumptions on g, μ and δ which will be required to obtain our results are stated in Notation 4.2.8 below.

Notation 4.2.8 (The condition (BN) on g, μ and δ)**.** Given a smooth projective curve Σ of genus g, a rank 2 and degree d unstable Higgs Harder-Narasimhan type $\mu = (d_1, d_2)$, and $\delta \in \mathbb{N}$, we set $\rho(g, \mu, \delta) := g - \delta(\delta + d_2 - d_1 + g - 1)$.We say that g, μ and δ satisfy the condition (BN) if the following three conditions are satisfied[9]:

(i) $1 \leq d_1 - d_2 \leq 2g - 3$;

(ii) $1 \leq \delta \leq g$;

(iii) $0 < \rho \leq g - 1$.

To obtain the Bialynicki-Birula stratification for the $\mathbb{C}^*$-action on $\mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}$, we must first study the components of the fixed point set.

Proposition 4.2.9 (Fixed point set for the $\mathbb{C}^*$-action on $\mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}$ when $\delta \neq 0$)**.** Let $\mu = (d_1, d_2)$ be a rank 2 and degree d Higgs Harder-Narasimhan type and let $\delta \in \mathbb{N} \setminus \{0\}$. Suppose moreover that g, μ and δ satisfy the condition (BN). Then we have:

(i) the moduli space $\mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}$ is non-empty;

(ii) the fixed point set for the Higgs field scaling $\mathbb{C}^*$-action consists of two connected and irreducible components which are each smooth quasi-projective varieties:

1) a component $\mathcal{N}_{2,d}^{\mu,\delta,\mathrm{indec}}$ parametrising Higgs bundles of the form $[(E, 0)]$ with

$$\dim \mathcal{N}_{2,d}^{\mu,\delta,\mathrm{indec}} = 2g - (\delta - 1)(\delta + d_2 - d_1 + g - 1) - 1;$$

[9] The notation (BN) is in reference to 'Brill-Noether', since these conditions are obtained from rank 1 Brill-Noether theory.

2) a component $\mathcal{F}_{\text{dec}}^{\mu,\delta,\text{indec}}$ parametrising Higgs bundles of the form $\left[\left(E_1 \oplus E_2, \begin{pmatrix} 0 & \phi_{12} \\ 0 & 0 \end{pmatrix}\right)\right]$ where $\phi_{12} : E_2 \to E_1 \otimes K_\Sigma$, with

$$\dim \mathcal{F}_{\text{dec}}^{\mu,\delta,\text{indec}} = 3g + \delta(\delta + d_2 - d_1 + g - 1) + d_1 - d_2 - 2;$$

(iii) the fixed point set is projective if and only if the Brill-Noether locus $W_\delta \subseteq \text{Pic}^{d_1-d_2}(\Sigma)$, defined by

$$W_\delta^{d_1-d_2} := \{M \in \text{Pic}^{d_1-d_2}(\Sigma) \mid h^0(\Sigma, M) = \delta\},$$

is closed in $\text{Pic}^{d_1-d_2}(\Sigma)$.

Remark 4.2.10 (Quasi-projectivity of the fixed point set and failure of semiprojectivity of $\mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}$)**.** In general, given $\delta \in \mathbb{N}$ the Brill-Noether locus $W_\delta^{d_1-d_2} \subseteq \text{Pic}^{d_1-d_2}(\Sigma)$ is not closed but only locally closed: it is open in the closed subscheme

$$W_{\geq\delta}^{d_1-d_2} := \{[L] \in \text{Pic}^{d_1-d_2}(\Sigma) \mid h^0(\Sigma, K_\Sigma) \geq \delta\}.$$

Thus $W_\delta^{d_1-d_2}$ is closed in $\text{Pic}^{d_1-d_2}(\Sigma)$ if (and in general only if) δ is the maximal dimension amongst those δ such that $W_\delta^{d_1-d_2}$ is non-empty. We note that this is true in the case where line bundles in $\text{Pic}^{d_1-d_2}(\Sigma)$ all have the same fixed dimension of global sections (this happens for example if $d_1 - d_2 + 2g - 2$ is less than or equal to 0, by the Riemann-Roch formula).

Nevertheless, in general the fixed point set will not be projective, and thus we see that while the moduli spaces $\mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}$ satisfy the first condition of semiprojectivity (the existence of limits at 0 under the $\mathbb{C}^*$-action) when $\delta \neq 0$, they do not necessarily satisfy the second condition (having a projective fixed point set). Thus in contrast with the moduli space $\mathcal{M}_{2,d}^s$ of stable Higgs bundles, the moduli spaces $\mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}$ need not be semiprojective.

To prove the above Proposition 4.2.9, we will use the existence and properties of moduli spaces of rank 2 unstable vector bundles, first constructed in [15]. We thus first introduce the construction and results from [15] which we will use in our proof of Proposition 4.2.9.

Moduli spaces for rank 2 unstable vector bundles. A moduli space for rank 2 unstable vector bundles was first constructed by Brambila-Paz, Mata and Nitsure in [15]. More precisely, for any $\delta \in \mathbb{N}$ they construct a quasi-projective moduli space, which we denote by $\mathcal{N}_{2,d}^{\mu,\delta,\text{indec}}(\Sigma)$, parametrising isomorphism classes of rank 2 unstable vector bundles E of Harder-Narasimhan type μ such that $E \not\cong \mathrm{gr}(E)$ and such that $\dim \mathrm{End}_{-1}(E) = \delta$. Since Σ is a fixed curve, for simplicity we omit Σ from the notation.

The moduli space is constructed as a projective bundle over another moduli space: the quasi-projective moduli space $\mathcal{N}_{2,d}^{\mu,\delta,\mathrm{dec}}$ parametrising isomorphism classes of rank 2 vector bundles E of Harder-Narasimhan type μ such that $E \cong \mathrm{gr}\, E$ and $\dim \mathrm{End}_{-1}(E) = \delta$. We note that both moduli spaces $\mathcal{N}_{2,d}^{\mu,\delta,\mathrm{indec}}$ and $\mathcal{N}_{2,d}^{\mu,\delta,\mathrm{dec}}$ have been constructed by Jackson in [59], in the more general setting of sheaves with Harder-Narasimhan type of length 2 on an arbitrary smooth projective variety.

The moduli space $\mathcal{N}_{2,d}^{\mu,\delta,\mathrm{dec}}$ can be described explicitly as follows. As in Proposition 4.2.9, let $W_\delta^{d_1-d_2}$ denote the locally closed Brill-Noether subscheme of the Picard variety $\mathrm{Pic}^{d_1-d_2}$ consisting of line bundles on Σ with global sections of dimension δ. Then $\mathcal{N}_{2,d}^{\mu,\delta,\mathrm{dec}}$ corresponds to the preimage of $W_\delta^{d_1-d_2}$ in $\mathrm{Pic}^{d_1} \times \mathrm{Pic}^{d_2}$, under the map $\mathrm{Pic}^{d_1} \times \mathrm{Pic}^{d_2} \to \mathrm{Pic}^{d_1-d_2}$ given by $([E_1],[E_2]) \mapsto [E_2^\vee \otimes E_1]$. By [15, §7], which follows from standard results of rank 1 Brill-Noether theory, we have that if $\mathcal{N}_{2,d}^{\mu,\delta,\mathrm{dec}}$ is non-empty, then

$$\dim \mathcal{N}_{2,d}^{\mu,\delta,\mathrm{dec}} = \rho + g = 2g - \delta(\delta + d_2 - d_1 + g - 1) \tag{4.4}$$

where $\rho := g - \delta(\delta + d_2 - d_1 + g - 1)$.

The moduli space $\mathcal{N}_{2,d}^{\mu,\delta,\mathrm{indec}}$ is constructed as a projective bundle over $\mathcal{N}_{2,d}^{\mu,\delta,\mathrm{dec}}$ in such a way that the fibre over $[E_1 \oplus E_2] \in \mathcal{N}_{2,d}^{\mu,\delta,\mathrm{dec}}$ is isomorphic to the projective space $\mathbb{P}(H^1(\Sigma, E_2^\vee \otimes E_1))$, which parametrises isomorphism classes of vector bundles obtained as non-trivial extensions of E_2 by E_1. By the Riemann-Roch theorem for line bundles on curves, we have that $h^0(\Sigma, E_2^\vee \otimes E_1) - h^1(\Sigma, E_2^\vee \otimes E_1) = d_1 - d_2 + 1 - g$. Since $\mathrm{End}_{-1}(E_1 \oplus E_2) \cong H^0(\Sigma, E_2^\vee \otimes E_1)$, it follows that the fibres are equidimensional, of dimension equal to $d_2 - d_1 + \delta + g - 2$ (assuming this value is greater than or equal to zero).

By applying results from Brill-Noether theory, it is shown in [15, Thm 7.7] that if the conditions $1 \leq d_1 - d_2 \leq 2g - 3$ and $1 \leq \delta \leq g$ are satisfied, then the moduli space $\mathcal{N}_{2,d}^{\mu,\delta,\mathrm{indec}}$ is non-empty if and only if the inequality $0 \leq \rho \leq g - 1$ holds. We note that this condition also ensures the moduli space $\mathcal{N}_{2,d}^{\mu,\delta,\mathrm{dec}}$ is non-empty. Moreover, under these assumptions it is shown that the moduli space $\mathcal{N}_{2,d}^{\mu,\delta,\mathrm{indec}}$ is reduced and of pure dimension given by

$$\begin{aligned}\dim \mathcal{N}_{2,d}^{\mu,\delta,\mathrm{indec}} &= \dim \mathcal{N}_{2,d}^{\mu,\delta,\mathrm{indec}} + d_2 - d_1 + \delta + g - 2 \\ &= \rho + g + d_2 - d_1 + \delta + g - 2 \\ &= 2g - (\delta - 1)(\delta + d_2 - d_1 + g - 1) - 1. \end{aligned} \tag{4.5}$$

If in addition $0 < \rho$ (leading to the condition (BN) of Notation 4.2.8), then the moduli space is irreducible and smooth, also by [15].

We will use the above results in the proof of Proposition 4.2.9 below.

Proof of Proposition 4.2.9. Let $\delta > 0$ and suppose that $[(E, \phi)] \in \mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}$ is fixed by the $\mathbb{C}^*$-action. Then by Proposition 4.2.2, the Higgs bundle $[(E, \phi)]$ is in one of two forms. Either $E \ncong \operatorname{gr} E$, in which case $[(E, \phi)] = [(E, 0)]$, or $E \cong \operatorname{gr} E = E_1 \oplus E_2$, in which case $[(E, \phi)] \cong \left[\left(E_1 \oplus E_2, \begin{pmatrix} 0 & \phi_{12} \\ 0 & 0 \end{pmatrix}\right)\right]$ for some non-zero $\phi_{12} : E_2 \to E_1 \otimes K_\Sigma$. If $[(E, 0)] \in \mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}$, then by Proposition 3.1.8 the underlying bundle E must have Harder-Narasimhan type μ, and moreover we have that $E \ncong \operatorname{gr} E$ and $\dim \operatorname{End}_{-1}(E) = \delta$. Conversely, any rank 2 vector bundle E of Harder-Narasimhan type μ such that $\operatorname{gr} E \ncong E$ and $\dim \operatorname{End}_{-1}(E) = \delta$ gives rise to a Higgs bundle $[(E, 0)] \in \mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}$.

Thus the subset of the fixed point set of $\mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}$ consisting of Higgs bundles with zero Higgs field can be identified with the quasi-projective moduli space $\mathcal{N}_{2,d}^{\mu,\delta,\mathrm{indec}}$. Moreover, since g, μ and δ satisfy the condition (BN), then by the properties of $\mathcal{N}_{2,d}^{\mu,\delta,\mathrm{indec}}$ stated above we obtain that the subset is reduced, irreducible and smooth, and of pure dimension as given in (4.5) above.

This subset is open in the fixed point set by the openness of indecomposability (see [15, §4.1]). It is also closed: if $[(E, 0)]$ lies in $\mathcal{N}_{2,d}^{\mu,\delta,\mathrm{indec}}$, then any limit of this point inside the moduli space $\mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}$ must still have zero Higgs field. Moreover, the underlying bundle must remain indecomposable since there are no points in $\mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}$ of the form $[(E, 0)]$ where $E \cong \operatorname{gr}(E)$, and thus any point in the limit remains in $\mathcal{N}_{2,d}^{\mu,\delta,\mathrm{indec}}$. Therefore $\mathcal{N}_{2,d}^{\mu,\delta,\mathrm{indec}}$ represents a connected component of the fixed point set in $\mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}$ for the $\mathbb{C}^*$-action.

We now turn to the other type of fixed points in $\mathcal{M}_{2,d}^{\mu,\delta,\mathrm{indec}}$, those which are of the form $\left[\left(E_1 \oplus E_2, \begin{pmatrix} 0 & \phi_{12} \\ 0 & 0 \end{pmatrix}\right)\right]$ with $\phi_{12} \neq 0$. We let $\mathcal{F}_{2,d}^{\mu,\delta,\mathrm{indec}}$ denote the subset of such fixed points and let F denote the forgetful map $F : \mathcal{F}_{2,d}^{\mu,\delta,\text{indec}} \to \mathcal{N}_{2,d}^{\mu,\delta,\text{dec}}$ given by

$$\left[\left(E_1 \oplus E_2, \begin{pmatrix} 0 & \phi_{12} \\ 0 & 0 \end{pmatrix}\right)\right] \mapsto [E_1 \oplus E_2].$$

We note that the image does indeed lie in $\mathcal{N}_{2,d}^{\mu,\delta,\text{dec}}$, since for $\left[\left(E_1 \oplus E_2, \begin{pmatrix} 0 & \phi_{12} \\ 0 & 0 \end{pmatrix}\right)\right] \in \mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}$ we have that

$$\dim \operatorname{End}_{-1}(E_1 \oplus E_2) = \dim \operatorname{End}_{-1}\left(E_1 \oplus E_2, \begin{pmatrix} 0 & \phi_{12} \\ 0 & 0 \end{pmatrix}\right) = \delta.$$

To describe the fibres of the map F, we let $[(E_1 \oplus E_2)] \in \mathcal{N}_{2,d}^{\mu,\delta,\text{dec}}$. Then a Higgs bundle in the fibre is determined by choosing a Higgs field $\phi_{12} : E_2 \to E_1 \otimes L$, which must be non-zero to ensure that the resulting Higgs bundle is indecomposable. Such a non-zero Higgs field exists for the following reason. By assumption $h^0(\Sigma, E_2^\vee \otimes E_1) = \delta > 0$, since $[(E_1 \oplus E_2)] \in \mathcal{N}_{2,d}^{\mu,\delta,\text{dec}}$. Thus there exists a non-zero section $\sigma \in H^0(\Sigma, E_2^\vee \otimes E_1)$. Since $h^0(\Sigma, K_\Sigma) = g > 0$ (we recall that we are assuming $g \geq 1$), we can also pick a non-zero section $\tau \in H^0(\Sigma, K_\Sigma)$, so that $\sigma \otimes \tau$ gives a non-zero section of $E_2^\vee \otimes E_1 \otimes K_\Sigma$ under the tensor map $H^0(\Sigma, E_2^\vee \otimes E_1) \otimes H^0(\Sigma, K_\Sigma) \to H^0(\Sigma, E_2^\vee \otimes E_1 \otimes K_\Sigma)$. This section corresponds to a non-zero Higgs field $\phi_{12} : E_2 \to E_1 \otimes K_\Sigma$. Moreover, two Higgs fields $\phi_{12}, \phi_{12}' : E_2 \to E_1 \otimes K_\Sigma$ give rise to isomorphic Higgs bundles if and only if they are conjugate via an automorphism of $E_1 \oplus E_2$. Just as in the proof of Theorem 4.1.7, we obtain that this is true if and only if the Higgs field are scalar multiples of each other. Thus the fibre of F over $[E_1 \oplus E_2] \in \mathcal{N}_{2,d}^{\mu,\delta,\text{dec}}$ is given by the projective space $\mathbb{P}(H^0(\Sigma, E_2^\vee \otimes E_1 \otimes K_\Sigma))$.

By Serre duality, we have that $h^0(\Sigma, E_2^\vee \otimes E_1 \otimes K_\Sigma) = h^1(\Sigma, E_1^\vee \otimes E_2)$, and using the Riemann-Roch formula we obtain that $h^1(\Sigma, E_1^\vee \otimes E_2) = d_1 - d_2 + g - 1$ (using the fact that $h^0(\Sigma, E_1^\vee \otimes E_2) = 0$ since $d_1 > d_2$ and E_1, E_2 are line bundles, hence stable). Therefore $h^0(\Sigma, E_2^\vee \otimes E_1 \otimes K_\Sigma) = d_1 - d_2 + g - 1$. Thus the subset $\mathcal{F}_{2,d}^{\mu,\delta,\text{indec}}$ of fixed points in $\mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}$ surjects onto the moduli space $\mathcal{N}_{2,d}^{\mu,\delta,\text{indec}}$, with fibres that are projective spaces of dimension $d_1 - d_2 + g - 2$.

As seen above, since g, μ and δ satisfy the condition (BN), the moduli space $\mathcal{N}_{2,d}^{\mu,\delta,\text{dec}}$must be an irreducible and smooth quasi-projective variety in order for $\mathcal{N}_{2,d}^{\mu,\delta,\text{indec}}$ to be irreducible and smooth. Thus $\mathcal{F}_{2,d}^{\mu,\delta,\text{indec}}$ admits a surjective map onto an irreducible and smooth quasi-projective variety with irreducible and smooth equidimensional fibres. It follows that $\mathcal{F}_{2,d}^{\mu,\delta,\text{indec}}$ is also irreducible and smooth (see [98, 11.4.C] for irreducibility and [90, Thm 3.3.27] for smoothness). Thus we have:

$$\begin{aligned}\dim \mathcal{F}_{2,d}^{\mu,\delta,\text{indec}} &= \dim \mathcal{N}_{2,d}^{\mu,\delta,\text{dec}} + d_1 - d_2 + g - 2 \\ &= \rho + g + d_1 - d_2 + g - 2 \\ &= 3g + d_1 - d_2 - \delta(\delta + d_2 - d_1 + g - 1) - 2.\end{aligned}$$

Since the fixed point set is the union of $\mathcal{N}_{2,d}^{\mu,\delta,\text{indec}}$ and $\mathcal{F}_{2,d}^{\mu,\delta,\text{indec}}$, and we have shown above that $\mathcal{N}_{2,d}^{\mu,\delta,\text{indec}}$ is a connected component of the fixed points set, it follows that the fixed point set consists of two connected components. We have moreover shown that they are irreducible and smooth.

Thus the fixed point set is projective if and only if both $\mathcal{N}_{2,d}^{\mu,\delta,\text{indec}}$ and $\mathcal{F}_{2,d}^{\mu,\delta,\text{indec}}$ are projective. But both maps $\mathcal{N}_{2,d}^{\mu,\delta,\text{indec}} \to \mathcal{N}_{2,d}^{\mu,\delta,\text{dec}}$ and $\mathcal{F}_{2,d}^{\mu,\delta,\text{indec}} \to \mathcal{N}_{2,d}^{\mu,\delta,\text{dec}}$ can be constructed as projective bundles over $\mathcal{N}_{2,d}^{\mu,\delta,\text{dec}}$. Thus the fixed point set is projective if and only if the base $\mathcal{N}_{2,d}^{\mu,\delta,\text{dec}}$ is projective. As we have seen, the moduli space $\mathcal{N}_{2,d}^{\mu,\delta,\text{indec}}$ corresponds to the preimage in $\text{Pic}^{d_1} \times \text{Pic}^{d_2}$ of the Brill-Noether locus $W_\delta^{d_1-d_2}$ under the map given by $([E_1],[E_2]) \mapsto E_2^\vee \otimes E_1$. It follows that $\mathcal{N}_{2,d}^{\mu,\delta,\text{dec}}$ is projective if and only if $W_\delta^{d_1-d_2}$ is closed in $\text{Pic}^{d_1-d_2}$. □

4.2.3 Geometry of the moduli spaces $\mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}$ for $\delta > 0$

Proposition 4.2.9 above shows that the fixed point set for the $\mathbb{C}^*$-action on the moduli space $\mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}$ when $\delta > 0$ consists of two smooth irreducible components. This allows us in Theorem 4.2.12 below to describe the Bialynicki-Birula stratification for the $\mathbb{C}^*$-action on $\mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}$ when $\delta > 0$ and to show that the strata are irreducible and smooth. It can also be used to compute the Poincaré series of $\mathcal{M}_{2,d}^{\mu\text{-s}}$ in terms of the Poincaré series of a Brill-Noether locus and of a Picard variety (see Corollary 4.2.13).

The statement of Proposition 4.2.9 relies on the following notation.

Notation 4.2.11. Let $\mathcal{M}_{\text{indec}}^{\mu,\delta,\text{indec}}$ and $\mathcal{M}_{\text{dec}}^{\mu,\delta,\text{indec}}$ denote the subvarieties of $\mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}$ consisting of Higgs bundles $[(E,\phi)]$ satisfying $E \not\cong \text{gr}(E)$ and $E \cong \text{gr}(E)$ respectively.

Theorem 4.2.12 (Bialynicki-Birula stratification for $\mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}(\Sigma, K_\Sigma)$ when $\delta > 0$)**.** Let Σ be a smooth projective curve of genus $g \geq 1$. Let $\mu = (d_1, d_2) \neq \mu_0$ be a rank 2 and degree d Higgs

Harder-Narasimhan type and let $\delta \in \mathbb{N} \setminus \{0\}$. Suppose that g, μ and δ satisfy the condition (BN). Then the Bialynicki-Birula stratification for the $\mathbb{C}^*$-action on $\mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}$ for $\delta > 0$ is given by

$$\mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}(\Sigma, L) = \mathcal{M}_{\text{indec}}^{\mu,\delta,\text{indec}}(\Sigma, L) \sqcup \mathcal{M}_{\text{dec}}^{\mu,\delta,\text{indec}}(\Sigma, L)$$

and satisfies the following properties:

(i) the open stratum is $\mathcal{M}_{\text{indec}}^{\mu,\delta,\text{indec}}$ which fibres over the component $\mathcal{N}_{2,d}^{\mu,\delta,\text{indec}}$ of the fixed point set, with constant fibre isomorphic to $H^0(\Sigma, K_\Sigma)$, and satisfies

$$\dim \mathcal{M}_{\text{indec}}^{\mu,\delta,\text{indec}} = 3g - (\delta - 1)(\delta + d_2 - d_1 + g - 1) - 1;$$

(ii) the closed stratum is $\mathcal{M}_{\text{dec}}^{\mu,\delta,\text{indec}}$, which fibres over the component $\mathcal{F}_{2,d}^{\mu,\delta,\text{indec}}$ of the fixed point set, with constant fibre isomorphic to $H^0(\Sigma, K_\Sigma)$, and satisfies

$$\dim \mathcal{M}_{\text{dec}}^{\mu,\delta,\text{indec}} = 4g - \delta(\delta + d_2 - d_1 + g - 1) + d_1 - d_2 - 2;$$

(iii) the two strata are irreducible and smooth.

Figure 4.4 illustrates Proposition 4.2.9 and Theorem 4.2.12 by depicting the Bialynicki-Birula stratification for the $\mathbb{C}^*$-action on $\mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}$ for $\delta > 0$, as well as the structure of the fixed point set. Before proving Theorem 4.2.12, we first state and prove the conclusion that we can draw from Theorem 4.2.12 regarding the geometry of the moduli space $\mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}$ for $\delta > 0$.

Corollary 4.2.13 (Geometry and Poincaré series of $\mathscr{M}_{2,d}^{\mu,\delta,\text{indec}}(\Sigma, K_\Sigma)$ when $\delta > 0$)**.** Under the assumptions of Theorem 4.2.12, the moduli space $\mathscr{M}_{2,d}^{\mu,\delta,\text{indec}}$ where $\delta > 0$ consists of two irreducible connected components which are both smooth, and the Poincaré polynomial of $\mathscr{M}_{2,d}^{\mu,\delta,\text{indec}}$ is given by

$$P_t\left(\mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}\right) = P_t(W_\delta^{d_2-d_1}) P_t(\text{Pic}^{d_1}) \left(P_t(\mathbb{P}^{d_2-d_1+g-2+\delta}) + P_t(\mathbb{P}^{d_1-d_2+g-2})\right),$$

where $W_\delta^{d_2-d_1}$ denotes the locally closed subvariety of $\text{Pic}^{d_2-d_1}$ given by

$$W_\delta^{d_2-d_1} := \{[M] \in \text{Pic}^{d_2-d_1} \mid h^0(\Sigma, M) = \delta\}.$$

Proof of Corollary 4.2.13. From Theorem 4.2.12 we have a decomposition

$$\mathcal{M}_{2,d}^{\mu,\delta,\text{indec}} = \mathcal{M}_{\text{indec}}^{\mu,\delta,\text{indec}} \sqcup \mathcal{M}_{\text{dec}}^{\mu,\delta,\text{indec}}$$

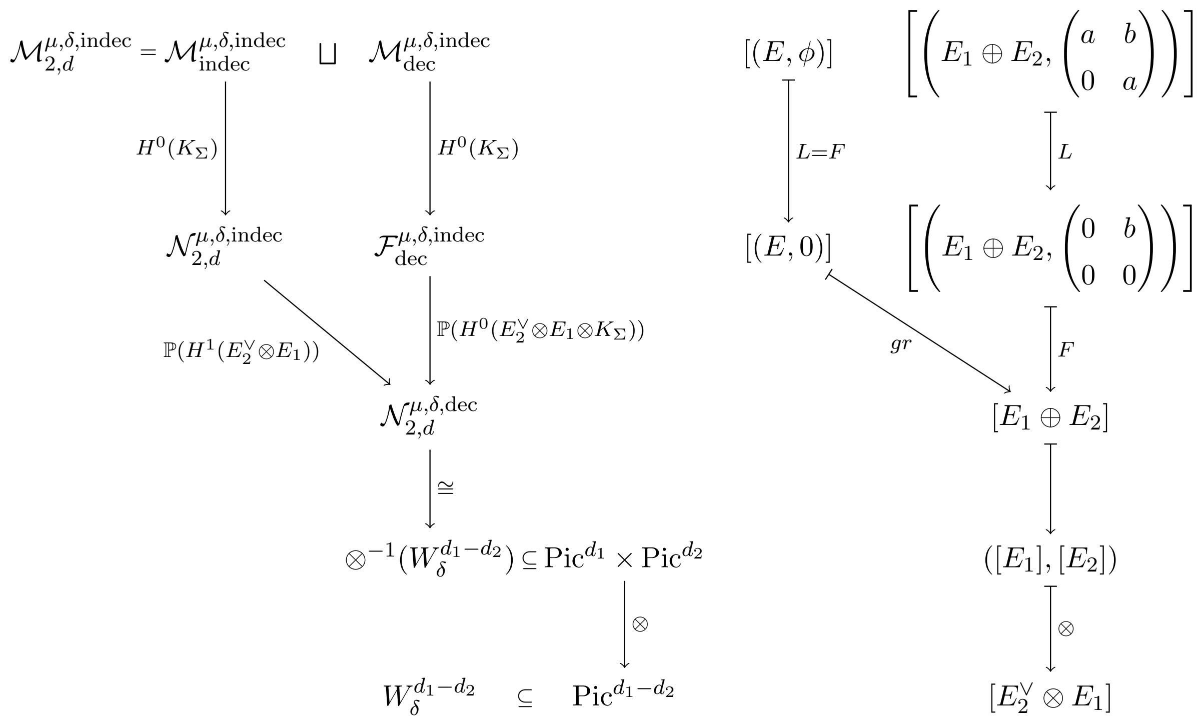

Figure 4.4: The Bialynicki-Birula stratification for the $\mathbb{C}^*$-action on $\mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}(\Sigma, K_\Sigma)$ when $\delta > 0$, and the structure of the fixed point set. Since Σ is fixed throughout, for simplicity we omit Σ from the notation. The left-hand side of the diagram describes the stratification and the structure of the fixed point set, while the right-hand side gives the definition of the maps. The stratification consists of two strata: the open stratum $\mathcal{M}_{\text{indec}}^{\mu,\delta,\text{indec}}$ parametrising Higgs bundles with indecomposable underlying bundle, and the closed stratum $\mathcal{M}_{\text{dec}}^{\mu,\delta,\text{indec}}$ parametrising Higgs bundles with decomposable underlying bundle. The labels on the arrows of the left-hand side of the diagram correspond to the fibres over a point in the image of the map, of the form provided in the right-hand side of the diagram. The map L denotes the map taking a Higgs bundle to its limit at 0 under the $\mathbb{C}^*$-action while the map F denotes the map forgetting the Higgs field. The map gr is defined by taking a vector bundle to its Harder-Narasimhan graded. The fixed point set consists of two smooth irreducible components $\mathcal{N}_{2,d}^{\mu,\delta,\text{indec}}$ and $\mathcal{F}_{\text{dec}}^{\mu,\delta,\text{indec}}$ (provided g, μ and δ satisfy the condition (BN)). We recall that $W_\delta^{d_1-d_2} \subseteq \text{Pic}^{d_1-d_2}$ is the Brill-Noether locus consisting of line bundles $[M] \in \text{Pic}^{d_1-d_2}$ with $h^0(\Sigma, M) = \delta$.

where $\mathcal{M}^{\mu,\delta,\text{indec}}_{\text{indec}}$ is open in $\mathcal{M}^{\mu,\delta,\text{indec}}_{2,d}$ and $\mathcal{M}^{\mu,\delta,\text{indec}}_{\text{dec}}$ is closed in $\mathcal{M}^{\mu,\delta,\text{indec}}_{2,d}$. To show that $\mathcal{M}^{\mu,\delta,\text{indec}}_{2,d}(\Sigma, L)$ is not irreducible, it suffices to show that the dimension of the open stratum is smaller than or equal to that of the closed stratum. By Theorem 4.2.12, we have that $\dim \mathcal{M}^{\mu,\delta,\text{indec}}_{\text{dec}} - \dim \mathcal{M}^{\mu,\delta,\text{indec}}_{\text{indec}} = d_1 - d_2 - \delta$. Thus it suffices to show that $d_1 - d_2 \geq \delta$, and this follows from the assumption that the condition (BN) holds for g, μ and δ. Indeed, by the condition (BN) we have that $\delta + d_2 - d_1 + g - 1 < g/\delta$, and since $\delta \geq 1$ by assumption we obtain that $\delta + d_2 - d_1 + g - 1 < g$, or equivalently that $d_1 - d_2 \geq \delta$ as required.

By Proposition 4.2.9, the fixed point set consists of two irreducible connected components. Thus the moduli space $\mathcal{M}^{\mu,\delta,\text{indec}}_{2,d}$ also consists of two connected components, each irreducible and smooth by Theorem 4.2.12.

It follows that the Poincaré series of $\mathcal{M}^{\mu,\delta,\text{indec}}_{2,d}$ is given by

$$P_t(\mathcal{M}^{\mu,\delta,\text{indec}}_{2,d}) = P_t(\mathcal{M}^{\mu,\delta,\text{indec}}_{\text{indec}}) + P_t(\mathcal{M}^{\mu,\delta,\text{indec}}_{\text{dec}}).$$

Moreover, using the retraction of each component onto the corresponding component of the fixed point set, we have that $P_t(\mathcal{M}^{\mu,\delta,\text{indec}}_{\text{indec}}) = P_t(\mathcal{N}^{\mu,\delta,\text{indec}}_{2,d})$ and $P_t(\mathcal{M}^{\mu,\delta,\text{indec}}_{\text{dec}}) = t^{2g} P_t(\mathcal{F}^{\mu,\delta,\text{indec}}_{2,d})$.

Both components $\mathcal{N}^{\mu,\delta,\text{indec}}_{2,d}$ and $\mathcal{F}^{\mu,\delta,\text{indec}}_{2,d}$ of the fixed point set fibre over $\mathcal{N}^{\mu,\delta,\text{dec}}_{2,d}$, as established in Proposition 4.2.9. By using the description of the fibres obtained there as well as the formula for the Poincaré series of a fibration (see [16]), we obtain:

$$P_t(\mathcal{F}^{\mu,\delta,\text{indec}}_{2,d}) = P_t(\mathcal{N}^{\mu,\delta,\text{dec}}_{2,d}) P_t(\mathbb{P}^{d_1-d_2+g-2}) \text{ and}$$
$$P_t(\mathcal{N}^{\mu,\delta,\text{indec}}_{2,d}) = P_t(\mathcal{N}^{\mu,\delta,\text{dec}}_{2,d}) P_t(\mathbb{P}^{d_2-d_1+\delta+g-2}).$$

It remains only to consider $P_t(\mathcal{N}^{\mu,\delta,\text{dec}}_{2,d})$, for which we use the existence of an isomorphism (see [15]) $\mathcal{N}^{\mu,\delta,\text{dec}}_{2,d} \cong W^{d_2-d_1}_{\delta} \times \text{Pic}^{d_1}$ given by $[E] \mapsto ([(E/E_1)^\vee \otimes E_1], [E_1])$ where E_1 denotes the canonically destabilising subbundle of E, with inverse given by $([L], [E_1]) \mapsto [E_1 \oplus (E_1 \otimes L^\vee)]$. As a result, we have that $P_t(\mathcal{N}^{\mu,\delta,\text{dec}}_{2,d}) = P_t(W^{d_2-d_1}_{\delta}) P_t(\text{Pic}^{d_1})$. Combining the above formulae yields the desired formula

$$P_t\left(\mathcal{M}^{\mu,\delta,\text{indec}}_{2,d}\right) = P_t(W^{d_2-d_1}_{\delta}) P_t(\text{Pic}^{d_1}) \left(P_t(\mathbb{P}^{d_2-d_1+g-2+\delta}) + P_t(\mathbb{P}^{d_1-d_2+g-2})\right)$$

for the Poincaré series of $\mathcal{M}^{\mu,\delta,\text{indec}}_{2,d}$ when $\delta > 0$. □

We conclude this section with the proof of Theorem 4.2.12.

Proof of Theorem 4.2.12. By Proposition 4.2.9, we have that the fixed point set for the $\mathbb{C}^*$-action consists of two smooth irreducible components: $\mathcal{N}_{2,d}^{\mu,\delta,\text{indec}}$ and $\mathcal{F}_{\text{dec}}^{\mu,\delta,\text{indec}}$. Thus the Bialynicki-Birula stratification is given by two strata: those Higgs bundles with limit in $\mathcal{N}_{2,d}^{\mu,\delta,\text{indec}}$ and those with limit in $\mathcal{F}_{2,d}^{\mu,\delta,\text{indec}}$. From Proposition 4.2.2 we know that the former is characterised by having an underlying bundle not isomorphic to its Harder-Narasimhan graded, while the second stratum is characterised by having its underlying bundle isomorphic to its graded. By the openness of indecomposability, proved in [15, §4.1], we obtain that the former is the open stratum, and the latter is the closed stratum.

We now describe these strata by studying the fibres of the retraction map onto the fixed point set. We consider the closed stratum $\mathcal{M}_{\text{dec}}^{\mu,\delta,\text{indec}}$ first. Given $\left[\left((E_1 \oplus E_2), \begin{pmatrix} 0 & \phi_{12} \\ 0 & 0 \end{pmatrix}\right)\right] \in \mathcal{M}_{\text{dec}}^{\mu,\delta,\text{indec}}$, a point in the fibre of the retraction map $\mathcal{M}_{\text{dec}}^{\mu,\delta,\text{indec}}$ is given by a Higgs bundle of the form $\left[\left((E_1 \oplus E_2), \begin{pmatrix} \phi_{11} & \phi_{12} \\ 0 & \phi_{22} \end{pmatrix}\right)\right]$ where $\phi_{11} = \phi_{22} \in H^0(\Sigma, K_\Sigma)$, under the identification of $H^0(\Sigma, E_1^\vee \otimes E_1 \otimes K_\Sigma)$ and $H^0(\Sigma, E_2^\vee \otimes E_2 \otimes K_\Sigma)$ with $H^0(\Sigma, K_\Sigma)$. Thus any choice of section of K_Σ gives rise to a Higgs bundle in the fibre, and moreover distinct sections give rise to non-isomorphic Higgs bundles. Indeed, the diagonal entries of the Higgs field remain unchanged after conjugating by an automorphism of $E_1 \oplus E_2$. Thus the fibre is isomorphic to the vector space $H^0(\Sigma, K_\Sigma)$ (which has dimension g), and therefore constant for each point in the fixed point set $\mathcal{F}_{\text{dec}}^{\mu,\delta,\text{indec}}$. By Proposition 4.2.9, we have that $\mathcal{F}_{\text{dec}}^{\mu,\delta,\text{indec}}$ is an irreducible and smooth quasi-projective variety. Since the fibres of the retraction map $\mathcal{M}_{\text{dec}}^{\mu,\delta,\text{indec}} \to \mathcal{F}_{\text{dec}}^{\mu,\delta,\text{indec}}$ are irreducible, smooth and of constant dimension, it follows that $\mathcal{M}_{\text{dec}}^{\mu,\delta,\text{indec}}$ is also irreducible and smooth (see [98, 11.4.C] for irreducibility and [90, Thm 3.3.27] for smoothness). Its dimension is given by:

$$\begin{aligned}\dim \mathcal{M}_{\text{dec}}^{\mu,\delta,\text{indec}} &= \dim \mathcal{F}_{\text{dec}}^{\mu,\delta,\text{indec}} + h^0(\Sigma, K_\Sigma) \\ &= (3g + \delta(\delta + d_2 - d_1 + g - 1) + d_1 - d_2 - 2) + g \text{ by Proposition 4.2.9} \\ &= 4g + \delta(\delta + d_2 - d_1 + g - 1) + d_1 - d_2 - 2.\end{aligned}$$

We now consider the open stratum $\mathcal{M}_{\text{indec}}^{\mu,\delta,\text{indec}}$. The fibre over a vector bundle $[E] \in \mathcal{N}_{2,d}^{\mu,\delta,\text{indec}}$ is determined by a choice of Higgs field $\phi \in \text{Hom}^\mu(E, E \otimes K_\Sigma)$ (we recall that this is the subspace of Higgs fields preserving the destabilising subbundle E_1). Two such Higgs fields ϕ and ϕ' give isomorphic Higgs bundles if and only if there exists an automorphism ψ of E such that $\phi' = \psi \circ \phi \circ \psi^{-1}$.

To determine when two Higgs fields are conjugate via an automorphism of E, it is helpful to work in the complex analytic setting, by letting $\mathcal{E}$, $\mathcal{E}_1$ and $\mathcal{E}_2$ denote the C^∞-bundles underlying E, E_1 and E_2 respectively. Indeed, there is a C^∞-splitting of $\mathcal{E}$ so that $\mathcal{E} = \mathcal{E}_1 \oplus \mathcal{E}_2$, and by Proposition A.0.1 we can write an automorphism of E in matrix form with respect to this splitting, namely in the form $\begin{pmatrix} \alpha & \beta \\ 0 & \alpha \end{pmatrix}$ where $\alpha \in \mathbb{G}_m$ and $\beta \in H^0(\Sigma, E_2^\vee \otimes E_1)$. We can also write a Higgs field ϕ giving rise to a Higgs bundle $[(E,\phi)]$ in the fibre of the retraction map over $[E]$ in matrix form with respect to this splitting, namely in the form $\begin{pmatrix} \phi_{11} & \phi_{12} \\ 0 & \phi_{22} \end{pmatrix}$. Since $[(E,\phi)] \in \mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}$ and $\delta > 0$, by Lemma 4.2.4 (ii) we must have that $\operatorname{End}_{-1}(E,\phi) = \operatorname{End}_{-1}(E)$. Using the same argument as that used in the proof of Proposition 3.2.12, we can conclude that the equality $\phi_{11} = \phi_{22}$ must hold.

By conjugating the matrix for a Higgs field ϕ as above by the matrix corresponding to an automorphism of E, we obtain that the diagonal entry $\phi_{11} = \phi_{22}$ is invariant upon conjugation, and that the entry ϕ_{12} is preserved only up to scaling by α^{-1}. Thus the quotient of $\operatorname{Hom}^\mu(E, E \otimes K_\Sigma)$ by the action of $\operatorname{Aut} E$ can be identified with the vector space $H^0(\Sigma, K_\Sigma)$. By Proposition 4.2.9, the component $\mathcal{N}_{2,d}^{\mu,\delta,\text{indec}}$ of the fixed point is irreducible and smooth, and therefore since $\mathcal{M}_{\text{indec}}^{\mu,\delta,\text{indec}}$ fibres over this component with equidimensional irreducible and smooth fibres, we obtain that the open stratum $\mathcal{M}_{\text{indec}}^{\mu,\delta,\text{indec}}$ is irreducible and smooth as required (see [98, 11.4.C] for irreducibility and [90, Thm 3.3.27] for smoothness). Finally, we have that

$$\begin{aligned}
\dim \mathcal{M}_{\text{indec}}^{\mu,\delta,\text{indec}} &= \dim \mathcal{N}_{2,d}^{\mu,\delta,\text{indec}} + h^0(\Sigma, K_\Sigma) \\
&= (2g - (\delta - 1)(\delta + d_2 - d_1 + g - 1) - 1) + g \\
&= 3g - (\delta - 1)(\delta + d_2 - d_1 + g - 1) - 1.
\end{aligned}$$

□

4.3 Moduli space for 'μ-stable' Higgs bundles on a curve of arbitrary genus

The previous Section 4.2 studied the geometry of the moduli spaces $\mathcal{M}_{2,d}^{\mu,\delta,\text{indec}}(\Sigma, L)$ when $\delta \neq 0$ and $L = K_\Sigma$. The condition that $\delta \neq 0$ is crucial for obtaining our results in Section 4.2: it ensures that the moduli spaces contain limits at 0 under the Higgs field scaling $\mathbb{C}^*$-action, endowing them with a Bialynicki-Birula stratification which can then be used to describe their geometry. Nevertheless, it can be argued that the case where $\delta = 0$ is the most interesting and the most important to understand. Indeed, the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}(\Sigma, L) = \mathcal{M}_{2,d}^{\mu,0,\text{indec}}(\Sigma, L)$ is the coarse moduli space for the open stratum in the refined Higgs Harder-Narasimhan stratification of a given unstable Higgs Harder-Narasimhan stratum $\mathscr{H}_{2,d}^{\mu}(\Sigma, L)$, parametrising μ-stable Higgs bundles of Higgs Harder-Narasimhan type μ. From this perspective, it can be said that the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}(\Sigma, L)$ is to unstable Higgs bundles of Higgs Harder-Narasimhan type μ what the moduli space $\mathcal{M}_{2,d}^{s}(\Sigma, L)$ is to all Higgs bundles. In other words, just as the concept of stability picks out an open subset of the parameter space for all Higgs bundles for which a quasi-projective coarse moduli space can be constructed, given a Higgs Harder-Narasimhan type μ, the concept of μ-stability picks out an open subset of the parameter space for Higgs bundles of Higgs Harder-Narasimhan type μ (itself locally closed in the parameter space for all Higgs bundles) for which a quasi-projective coarse moduli space can be constructed.

The aim of this section is to describe a strategy for studying the geometry of the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}(\Sigma, L)$ which does not rely on the Higgs field scaling $\mathbb{C}^*$-action. Instead, we use the fact that the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}(\Sigma, L)$ can be constructed as a non-reductive GIT quotient, as shown in Chapter 3. This strategy reflects the point emphasised in the introduction to Chapter 2, namely that GIT and Non-Reductive GIT are useful not just for constructing moduli spaces, but also for studying the geometry of these spaces. More precisely, knowing that a given moduli space can be constructed as a classical or non-reductive GIT quotient enables the application of existing formulae for computing the cohomology of such quotients, as seen in Chapter 2. We show in this section how these formulae could be applied to compute the Poincaré series of a partial compactification of the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}(\Sigma, L)$, a partial compactification which we construct using Non-Reductive GIT.

To this end, in Section 4.3.1 we review the Non-Reductive GIT construction of the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}(\Sigma, L)$. In Section 4.3.2 we construct a partial compactification of $\mathcal{M}_{2,d}^{\mu\text{-s}}(\Sigma, L)$ using Non-Reductive GIT. Finally, in Section 4.3.3 we explain how the cohomological formula from Section 2.3.3, given in Corollary 2.3.5, can be used to compute the Poincaré series of the partial compactification of the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}(\Sigma, L)$.

While in the previous Section 4.2.3 we considered twisted Higgs bundles in the classical case, namely where the twisting line bundle $L \to \Sigma$ is the canonical line bundle K_Σ of Σ, in this section we return to the general case where L is any line bundle on Σ. We keep Σ and L fixed throughout and thus omit this data from the notation from here on.

4.3.1 Review of the construction of $\mathcal{M}_{2,d}^{\mu\text{-s}}$ using Non-Reductive GIT

In this section we review the construction of the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}$, achieved in Section 3.2 of Chapter 3 using the spectral correspondence and the application of Non-Reductive GIT to the construction of moduli spaces for sheaves. For the remainder of the section we fix a Higgs Harder-Narasimhan type $\mu \neq \mu_0$ of rank 2 and degree d.

The set-up leading to the construction of $\mathcal{M}_{2,d}^{\mu\text{-s}}$ can be summarised as follows (we use the notation introduced in Sections 3.2.1 and 3.2.2):

(i) there is an isomorphism between $\mathscr{H}_{2,d}^{\mu}(\Sigma, L)$ and a quotient stack $\left[\mathring{Q}_n^{\mu}/G_n\right]$ for n sufficiently large, where the centre of $G_n = \mathrm{GL}(V_n)$ acts trivially on $\mathring{Q}_n^{\mu}$;

(ii) the scheme $\mathring{Q}_n^{\mu}$ is a locally closed subscheme of $\mathring{Q}_n \subseteq \mathrm{Quot}(V_n \otimes \mathcal{O}_Y(-n), \mathcal{P})$, with $\mathring{Q}_n$ parametrising sheaves with support not intersecting $D = Y \setminus \mathrm{Tot}\, L$ where Y is a projective completion of $\mathrm{Tot}\, L$;

(iii) there exists an $\mathrm{SL}(V_n)$-equivariant embedding of $\mathring{Q}_n$ inside a big projective space $X_{n,m}$ (provided $m \gg n \gg 1$) admitting a linear action of $\mathrm{SL}(V_n)$, so that there is a linear action of $\mathrm{SL}(V_n)$ on the closure $\overline{\mathring{Q}}_n$ of $\mathring{Q}_n$ inside $X_{n,m}$;

(iv) given the GIT-instability stratification $\overline{\mathring{Q}_n} = \bigsqcup_{\beta \in \mathcal{B}} S_\beta$ associated to the linear action of $\mathrm{SL}(V_n)$ (and to a choice of invariant inner product on the Lie algebra of $\mathrm{SL}(V_n)$) on $\overline{\mathring{Q}_n}$, there is a correspondence $\mu \mapsto \beta_{n,m}(\mu)$ between Harder-Narasimhan types μ and GIT-instability types β, such that any point in $\mathring{Q}_n$ corresponding to a Higgs bundle of Higgs Harder-Narasimhan type μ lies in the stratum indexed by $\beta_{n,m}(\mu)$ for $m \gg n \gg 1$.

As a result of the above set-up, the problem of constructing a moduli space for Higgs bundles of Higgs Harder-Narasimhan type μ can be reduced to that of constructing a quotient for the action of $\mathrm{SL}(V_n)$ on the unstable stratum $S_{\beta_{n,m}(\mu)} \subseteq \overline{\mathring{Q}_n^{\mu}}$ (to simplify notation we let $\beta = \beta_{n,m}(\mu)$ from here onwards in this section).

This is a typical setting in which Non-Reductive GIT can be applied. Indeed, as we have seen in Section 1.1.1, the problem can be further reduced to that of constructing a quotient for the action of P_β on Y_β^{ss}, where $S_\beta \cong \mathrm{SL}(V_n) \times_{P_\beta} Y_\beta^{ss}$ and $P_\beta \subseteq \mathrm{SL}(V_n)$ is the parabolic (and therefore non-reductive) subgroup determined by the instability index β. Non-Reductive GIT can then be applied to the linear action of P_β on the closure $\overline{Y_\beta^{ss}}$ of Y_β^{ss} in $\overline{\mathring{Q}_n}$, and this application produces an open P_β-invariant subset of Y_β^{ss} which admits a quasi-projective geometric quotient, as well as a projective completion of this geometric quotient.

The moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}$ is an open subset of this quotient, obtained by restricting the open P_β-invariant subset of $\overline{Y_\beta^{ss}}$ to its intersection with $\mathring{Q}_n^\mu$. More precisely, if we set

$$X' := \overline{Y_\beta^{ss}} \cap \mathring{Q}_n^\mu \text{ and } X := \overline{Y_\beta^{ss}},$$

then the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}$ is given by the geometric $P_{\beta(\mu)}$-quotient of the open subset of X' corresponding to the intersection of X' with the stable locus (in the non-reductive sense) for the action of P_β on X (to which Non-Reductive GIT applies since X is projective).

4.3.2 Partial compactification of $\mathcal{M}_{2,d}^{\mu\text{-s}}$ using Non-Reductive GIT

The above Non-Reductive GIT set-up enables the construction not just of the moduli space $\mathcal{M}_{2,d}^{\mu\text{-s}}$, but also of a partial compactification of this space. We describe this partial compactification using the notation introduced in Section 4.3.1 above.

As we have seen in Section 1.2, the construction of projective Non-Reductive GIT quotients depends on the conditions satisfied by the group action, in particular on whether or not semistability coincides with stability, in both the non-reductive (for the action of the externally graded unipotent group contained in the non-reductive group) and classical sense (for the residual reductive group action on the intermediate $\widehat{U}$-quotient). Indeed, if semistability does not coincide with stability in the non-reductive sense, then blow-ups are necessary to obtain a projective quotient. Moreover, as we have seen in Section 1.2.2, the blow-up construction can vary depending on the conditions satisfied by the unipotent radical of the non-reductive group and those satisfied by the action of its semi-direct product with the grading $\mathbb{G}_m$ on the variety.

Since the partial compactification of $\mathcal{M}_{2,d}^{\mu\text{-s}}$ which we will construct below is determined by a projective Non-Reductive GIT quotient associated to the action of P_β on X, we now study in more detail the conditions satisfied by the action of P_β on X.

The parabolic subgroup P_β can be written as a semidirect product $U_\beta \rtimes L_\beta$ where U_β denotes its unipotent radical and L_β a Levi subgroup (see Section 1.1.1). We let $\lambda_\beta : \mathbb{G}_m \to$

$Z(L_\beta)$ denote the grading one-parameter subgroup. By [59, Prop 5.3.1.4] (this result applies to sheaves, and can in turn be applied to Higgs bundles via the spectral correspondence), we have an isomorphism $\operatorname{Lie}\operatorname{Stab}_{U_\beta}(E,\phi) \cong \operatorname{End}_{-1}(E,\phi)$, and the latter group is not necessarily trivial. As a result, the condition ($ss = s \neq \emptyset[\widehat{U}]$) is not in general satisfied for the action of $\widehat{U}_\beta := U_\beta \rtimes \lambda_\beta(\mathbb{G}_m) \subseteq P_\beta$ on X, which means that blow-ups are required to construct a quotient, as per Theorem 1.2.3.

Nevertheless, the action of P_β on X satisfies the following three conditions, which place us in the simplest situation for performing the blow-ups from Non-Reductive GIT (see Theorem 1.2.10):

(1) by Proposition 3.2.12 there exists a point in $Z_{\min}$ (in fact in $Z_{\min} \cap X'$) with trivial unipotent stabiliser. This places us in the simpler case of the $\widehat{U}$-theorem with blow-ups (Theorem 1.2.3), for which Blow-up Construction 1 applies;

(2) the unipotent radical U_β of P_β is abelian and $\lambda_\beta(\mathbb{G}_m)$ acts with a single positive weight on $\operatorname{Lie} U_{\beta(\mu)}$ via the adjoint action. This follows from the correspondence $\mu \mapsto \beta_{n,m}(\mu) =: \beta$ of [55, §6.1] and the fact that $\mu = (d_1, d_2)$ is a Higgs Harder-Narasimhan type of rank 2 (so that $d_1 > d_2$), which together ensure that $U_\beta \subseteq \operatorname{SL}(V_n)$ consists of upper triangular unipotent matrices of the form

$$\left(\begin{array}{ccc|cc} 1 & & 0 & & \\ & \ddots & & \multicolumn{2}{c}{U_{\beta(\mu)}} \\ 0 & & 1 & & \\ \hline & 0 & & 1 & 0 \\ & & & 0 & 1 \end{array}\right),$$

and that the one-parameter subgroup $\lambda_\beta(\mathbb{G}_m)$ acts with a single positive weight on U_β. This condition on U_β and the adjoint action of λ_β on $\operatorname{Lie} U_\beta$ corresponds to the simplest case for obtaining cohomological formulae of Non-Reductive GIT quotients (see Theorem 2.3.2);

(3) semistability coincides with stability for the induced action of $L_\beta/\lambda_\beta(\mathbb{G}_m)$ on the intersection of $Z_{\min}$ with X'. This is because points in the U_β-sweep of $Z_{\min} \cap X'$ can be interpreted as Higgs bundles which are isomorphic to their Higgs Harder-Narasimhan

graded (this follows from [55, Lem 6.7], by identifying Higgs bundles as sheaves under the spectral correspondence). The stabiliser of such Higgs bundles in L_β consists of a product of two copies of $\mathbb{G}_m$, representing the product of the automorphism groups of each rank 1 Higgs bundle making up the Higgs Harder-Narasimhan graded. The stabiliser of such a Higgs bundle in the Levi subgroup L_β of $P_{\beta(\mu)} \subseteq \mathrm{SL}(V_n)$ is isomorphic to a single copy of $\mathbb{G}_m$, since we are working with $\mathrm{SL}(V_n)$ rather than $\mathrm{GL}(V_n)$, and thus further quotienting by $\lambda_\beta(\mathbb{G}_m)$ produces a trivial stabiliser group. This condition means that only the $\widehat{U}$-blow-ups need to be performed for the action of P_β on X in order to obtain a projective geometric quotient; the partial desingularisation construction is not needed since semistability already coincides with stability for the induced linear action of L_β/λ_β on $\widehat{X}//\widehat{U}_\beta$.

By (1) we can apply the alternative Blow-up Construction 1 to the action of $\widehat{U}_\beta \subseteq P_\beta$ on X, resulting in a variety $\widehat{X}$ with a linear action of P_β satisfying the condition ($ss = s \neq \emptyset[\widehat{U}]$). Thus there is an associated projective Non-Reductive GIT quotient $\widehat{X}//P_\beta$. We let $\widehat{X}'$ denote the proper transform in $\widehat{X}$ of $X' \subseteq X$, noting that $\widehat{X}'$ is non-empty since X is never contained in the centre of the blow-ups, thanks to (1). The subvariety $\widehat{X}'$ is a P_β-invariant subset of $\widehat{X}$. We note that the restriction of the quotient $\widehat{X}//P_\beta$ to $\widehat{X}'$ is a geometric quotient, by (3). The partial compactification of $\mathcal{M}_{2,d}^{\mu\text{-s}}$ is defined to be the image of $\widehat{X}'$ in the quotient $\widehat{X}//P_\beta$ and we denote it by $\mathcal{M}_{2,d}^{\mu\text{-ss}}$. The reason that $\mathcal{M}_{2,d}^{\mu\text{-ss}}$ is only a partial compactification, rather than a projective completion, is that $\widehat{X}'$ may not be closed in $\widehat{X}$ (since X' may not be closed in X), and therefore its image inside the projective quotient $\widehat{X}//P_\beta$ may not be closed. Remark 4.3.1 below justifies the choice of notation for the partial compactification.

Remark 4.3.1 (Analogy with the partial compactification $\mathcal{M}_{2,d}^{ss}$ of $\mathcal{M}_{2,d}^{s}$)**.** Suppose that the degree d is even, so that the inclusion of the moduli space $\mathcal{M}_{2,d}^{s}$ of stable Higgs bundles inside the moduli space $\mathcal{M}_{2,d}^{ss}$ of semistable Higgs bundles is strict. Since $\mathcal{M}_{2,d}^{ss}$ is quasi-projective, it is only a partial compactification of $\mathcal{M}_{2,d}^{s}$ (we note that compactifications of $\mathcal{M}_{2,d}^{ss}$ have been studied in [46] and [89] for example). The partial compactification $\mathcal{M}_{2,d}^{\mu\text{-ss}}$ of $\mathcal{M}_{2,d}^{\mu\text{-s}}$ can be thought of as the analogue for the unstable case of the partial compactification of $\mathcal{M}_{2,d}^{s}$ given by $\mathcal{M}_{2,d}^{ss}$. It is for this reason that we use the superscript 'μ-ss' to denote it. Nevertheless, there is a key difference between the partial compactifications of $\mathcal{M}^{s}_{2,d}$ and of $\mathcal{M}^{\mu\text{-s}}_{2,d}$: the former has a clear moduli-theoretic interpretation, which is not the case a priori for the latter.

4.3.3 Geometry of the partial compactification of $\mathcal{M}^{\mu\text{-s}}_{2,d}$

The three conditions, described in Section 4.3.2 above, that are satisfied by the action of P_β on X' (the subvariety of X parametrising Higgs bundles of Higgs Harder-Narasimhan type μ) are important because they are conditions which must be satisfied if we wish to apply the cohomological formula of Corollary 2.3.5. The aim of this section is to show how this formula can be applied to compute the Poincaré series of the partial compactification $\mathcal{M}^{\mu\text{-ss}}_{2,d}$ of $\mathcal{M}^{\mu\text{-s}}_{2,d}$ obtained in Section 4.3.2 above.

There are two obstacles to overcome if we wish to apply the formula of Corollary 2.3.5 to the action of P_β on X'; they concern projectivity and smoothness. Indeed, Corollary 2.3.5 is formulated for group actions on smooth projective varieties, whereas X' is not projective, and a priori we do not know that it is smooth. However, we do have that $Z_{\min} \cap X' \subseteq Z_{\min}$ is smooth. This is because it parametrises rank 2 Higgs bundles of Higgs Harder-Narasimhan type μ which are isomorphic to their graded, and such Higgs bundles admit a smooth moduli space consisting of a product of moduli spaces for Higgs bundles of rank 1, which are smooth.

We expect that Corollary 2.3.5 can be generalised so that it can be applied in this situation. The generalisation we have in mind, and which would enable its application to the situation we are considering, is the following: given $H = U \rtimes R$ (satisfying the conditions of Corollary 2.3.5) acting linearly on an irreducible projective variety X (not necessarily smooth) and $X' \subseteq X$ a quasi-projective subvariety such that $X' \cap Z_{\min}$ is non-empty and smooth, then the formula of Corollary 2.3.5 is still valid if we replace all relevant subvarieties of X by their intersection with X'.

Assuming this conjecture, then as illustrated in Figure 2.4, the Poincaré series of the projective completion $\mathcal{M}^{\mu\text{-ss}}_{2,d}$ of $\mathcal{M}^{\mu\text{-s}}_{2,d}$ can then be computed from the Poincaré series of $Z_{\min} \cap X')/\!/(L_\beta/\lambda_\beta(\mathbb{G}_m))$ and from those of the images inside $(Z_{\min} \cap X')/\!/(L_\beta/\lambda_\beta(\mathbb{G}_m))$ of the centres of the blow-ups of the alternative Blow-up Construction 1 which produces the projective completion. We conclude this section by describing these auxiliary spaces in a moduli-theoretic way, in order to show that they are spaces for which we can feasibly expect to be able to compute the Poincaré series.

We start with the space $(Z_{\min} \cap X')//(L_\beta/\lambda_\beta(\mathbb{G}_m))$. As noted in (3) above, the U_β-sweep of $Z_{\min} \cap X'$ is the locus in X' parametrising Higgs bundles which are isomorphic to their Higgs Harder-Narasimhan graded. Since two Higgs bundles in X' are isomorphic if and only if they lie in the same P_β-orbit, the GIT quotient $(Z_{\min} \cap X')//(L_\beta/\lambda_\beta(\mathbb{G}_m))$ can be interpreted as a moduli space for Higgs bundles of Higgs Harder-Narasimhan type μ which are isomorphic to their Higgs Harder-Narasimhan graded. This moduli space is isomorphic to a product of moduli spaces of rank 1 Higgs bundles of degree d_1 and d_2 respectively, where $\mu = (d_1, d_2)$. In other words, we have an isomorphism

$$(Z_{\min} \cap X')//(L_\beta/\lambda_\beta(\mathbb{G}_m)) \cong \mathcal{M}^{ss}_{1,d_1} \times \mathcal{M}^{ss}_{1,d_2} \tag{4.6}$$

where $\mathcal{M}^{ss}_{1,d_i}$ denotes the moduli space of semistable Higgs bundles of rank 1 and degree d_i (the semistability condition is not in fact necessary since all Higgs bundles of rank 1 are stable). These moduli spaces are easy to describe: they are given by the product of a Picard variety with the space of global sections of the twisting line bundle L (see [34] for example). Thus computing the Poincaré series of $(Z_{\min} \cap X')//(L_\beta/\lambda_\beta(\mathbb{G}_m))$ reduces to the problem of computing the Poincaré series of Picard varieties. We recall that the Poincaré series of Picard varieties also appear as a key ingredient in the formula for the Poincaré series of $\mathcal{M}^{\mu,\delta,\text{indec}}_{2,d}(\Sigma, K_\Sigma)$ obtained in Corollary 4.2.13.

It remains only to describe the centres of the blow-ups from the alternative Blow-up Construction 1, more precisely their images inside the GIT quotient $(Z_{\min} \cap X')//(L_\beta/\lambda_\beta(\mathbb{G}_m))$. By definition, the centre of the first blow-up, which we denote by $C_{\max}$, corresponds to the closed subvariety of X' consisting of the locus of points in the U_β-sweep of $Z_{\min} \cap X'$ with maximal dimension stabiliser group in $\widehat{U}_\beta := U_\beta \rtimes \lambda_\beta(\mathbb{G}_m)$. As we have already seen, the Lie algebra of the unipotent stabiliser for a Higgs bundle $(E, \phi) \in X'$ coincides with the group $\text{End}_{-1}(E, \phi)$ of nilpotent endomorphisms of (E, ϕ). And as we have seen above, the U_β-sweep of $Z_{\min} \cap X'$ parametrises Higgs bundles of Higgs Harder-Narasimhan type μ which are isomorphic to their Higgs Harder-Narasimhan graded. Thus the image of the centre of the first blow-up in the quotient $(Z_{\min} \cap X')//(L_\beta/\lambda_\beta(\mathbb{G}_m))$, which we identify with the product $\mathcal{M}^{ss}_{1,d_1} \times \mathcal{M}^{ss}_{1,d_2}$ using (4.6), consists of pairs $([(E_1, \phi_1)], [(E_2, \phi_2)]) \in \mathcal{M}^{ss}_{1,d_1} \times \mathcal{M}^{ss}_{1,d_2}$ such that $\dim \text{End}_{-1}(E_1 \oplus E_2, \phi_1 \oplus \phi_2)$

is maximal.

By Lemma 4.1.4, we have that $\mathrm{End}_{-1}(E_1 \oplus E_2, \phi_1 \oplus \phi_2) = \{0\}$ if $\phi_1 \neq \phi_2$. Thus to determine the locus yielding the maximal dimension, we can assume that $\phi_1 = \phi_2$. In this case, again by Lemma 4.1.4, we have that $\mathrm{End}_{-1}(E_1 \oplus E_2, \phi_1 \oplus \phi_2) = \mathrm{End}_{-1}(E_1 \oplus E_2)$, and the latter group can be identified with the group $H^0(\Sigma, E_2^\vee \otimes E_1)$. As in Section 4.2.3, we let $\otimes$ denote the map $\mathrm{Pic}^{d_1} \times \mathrm{Pic}^{d_2} \to \mathrm{Pic}^{d_1-d_2}$ given by $([E_1],[E_2]) \mapsto [E_2^\vee \otimes E_1]$.

Let $\delta_{\max} = d_{(0)} > d_{(1)} > \cdots > d_{(r+1)}$ denote in decreasing order the integers $\delta \in \mathbb{N}$ such that there exists a line bundle M of degree $d_1 - d_2$ in the image of $\otimes$ satisfying $h^0(\Sigma, M) = \delta$. Then we have that $\delta_{\max} = \max\left\{\delta \in \mathbb{N} \mid \text{there exists } M \in \mathrm{Im}(\otimes) \text{ such that } h^0(\Sigma, M) = \delta\right\}$. For each $\delta \in \mathbb{N}$, we let $W^{d_1-d_2}_{\geq\delta} \subseteq \mathrm{Pic}^{d_1-d_2}$ and $W^{d_1-d_2}_{\delta} \subseteq \mathrm{Pic}^{d_1-d_2}$ denote the Brill-Noether loci consisting of isomorphism classes of line bundles M such that $h^0(\Sigma, M) \geq \delta$ and $h^0(\Sigma, M) = \delta$ respectively.

Thus, given $([E_1,\phi_1],[E_2,\phi_2]) \in \mathcal{M}^{ss}_{1,d_1} \times \mathcal{M}^{ss}_{1,d_2}$, we have that a Higgs bundle $([E_1,\phi_1],[E_2,\phi_2])$ lies in $C_{\max}/\!/(L_{\beta(\mu)}/\lambda_{\beta(\mu)}(\mathbb{G}_m))$ if and only if $\phi_1 = \phi_2$ and ${E_2}^\vee \otimes E_1 \in W_{\delta_{\max}} \subseteq \mathrm{Pic}^{d_1-d_2}$. As a result, we obtain a map

$$\otimes^{-1}(W_{\delta_{\max}}) \times H^0(\Sigma, L) \to C_{\max}/\!/(L_{\beta(\mu)}/\lambda_{\beta(\mu)}(\mathbb{G}_m))$$
$$(([E_1],[E_2]),\phi) \mapsto ([(E_1,\phi)],[(E_2,\phi)])$$

which is an isomorphism. Thanks to this identification, we see that the Poincaré series of the centre $C_{\max}$ can be computed from the Poincaré series of $\otimes^{-1}(W_{\delta_{\max}})$, which we expect to be able to compute from the Poincaré series of a suitable Brill-Noether locus in a Picard variety. This conclusion mirrors that obtained in the case of the moduli spaces $\mathcal{M}^{\mu,\delta,\mathrm{indec}}_{2,d}(\Sigma, K_\Sigma)$ for $\delta > 0$, for which the Poincaré series also involves the Poincaré series of Brill-Noether loci in Picard varieties (see Corollary 4.2.13).

To compute the Poincaré series of $\mathcal{M}^{\mu\text{-ss}}_{2,d}$, we must describe the centre not only of the first blow-up, as we have done above, but also of the blow-ups at each stage of the alternative Blow-up Construction 1. But these can similarly be described in a moduli-theoretic way thanks to Theorem 2.3.2 which establishes that at each stage of the alternative Blow-up Construction 1, the centre of the blow-up corresponds to a resolution of singularities of a closed subvariety of the original variety. Applying this result in the current setting, we obtain that at the i-th stage of the alternative Blow-up Construction 1, the quotient $C^{(i)}_{\max}//(L_{\beta(\mu)}/\lambda_{\beta(\mu)}(\mathbb{G}_m))$ of the centre of the blow-up corresponds to a resolution of singularities of the subvariety of $\mathcal{M}^{ss}_{1,d_1} \times \mathcal{M}^{ss}_{1,d_2}$ given by $\otimes^{-1}(W_{\geq d_{i-1}}) \times H^0(\Sigma, L)$, where an element $(([E_1],[E_2]),\phi)$ corresponds to the pair $([(E_1,\phi)],[(E_2,\phi)])$. This implies that the Poincaré series of the centres of the subsequent blow-ups could be computed from the Poincaré series of suitable desingularisations of Brill-Noether loci.

We aim in future work to turn the above argument into an explicit formula for the Poincaré series of the partial compactification $\mathcal{M}^{\mu\text{-ss}}_{2,d}$ in terms of Poinaré series of Picard varieties and Brill-Noether loci, analogous to the formula obtained for the Poincaré series of the moduli spaces $\mathcal{M}^{\mu,\delta,\mathrm{indec}}_{2,d}(\Sigma, K_\Sigma)$ for $\delta > 0$ in Corollary 4.2.13.

Remark 4.3.2 (Partial desingularisations of Brill-Noether loci)**.** Following from the above analysis, we note that using the alternative Blow-up Construction 1 as part of the Non-Reductive GIT construction of the moduli space of Higgs bundles (or even just vector bundles, as in [59]), can provide a way of desingularising Brill-Noether loci for line bundles on smooth projective curves (the reason line bundles are involved, rather than higher rank bundles, is that we are considering the rank 2 case). Extending the above results to the higher rank case could therefore provide a way of desingularising Brill-Noether loci for higher rank bundles on smooth projective curves. We hope to pursue this idea in future work, since Brill-Noether loci and their singularities represent an active area of research, dating back to Riemann (see [17] and the references therein).

Appendix A

Automorphism groups of unstable vector bundles of rank 2

The purpose of this appendix is to prove Proposition A.0.1 which describes the automorphism groups of unstable vector bundles of rank 2. The description differs based on whether or not the vector bundle is isomorphic to its associated Harder-Narasimhan graded. This result is well-known if we assume that the vector bundle is isomorphic to its Harder-Narasimhan graded (see for example [24]), but we have not found a reference in the literature for the case where the vector bundle is not isomorphic to its Harder-Narasimhan graded.

We note that Proposition A.0.1 is needed to obtain the results of Sections 4.1 and 4.2 in Chapter 4.

Proposition A.0.1 (Automorphisms of unstable vector bundles of rank 2)**.** Let E be an unstable vector bundle of rank 2, let $0 = E^0 \subsetneq E^1 \subsetneq E^2 = E$ denote its Harder-Narasimhan filtration and $\mathrm{gr}(E) = E_1 \oplus E_2$ its Harder-Narasimhan graded where $E_1 \cong E^1/E^0$ and $E_2 \cong E^2/E^1$. We have:

(i) if $E \cong \mathrm{gr}\, E$, then $\mathrm{Aut}(E)$ is isomorphic to the group of matrices $\begin{pmatrix} \alpha & \beta \\ 0 & \delta \end{pmatrix}$ where $\alpha \in \mathrm{Aut}(E_1) \cong \mathbb{G}_m$, $\beta \in \mathrm{Aut}(E_2) \cong \mathbb{G}_m$ and $\beta \in H^0(\Sigma, E_2^\vee \otimes E_1)$, so that

$$\mathrm{Aut}(E) \cong H^0(\Sigma, E_2^\vee \otimes E_1) \rtimes (\mathbb{G}_m \times \mathbb{G}_m);$$

(ii) if $E \ncong \mathrm{gr}\, E$, then the restriction map $r\colon \mathrm{Aut}(E) \to \mathrm{Aut}(E^1)$ induces a short exact sequence $1 \to H^0(\Sigma, (E/E^1)^\vee \otimes E^1) \to \mathrm{Aut}(E) \to \mathrm{Aut}(E^1) \cong \mathbb{G}_m \to 1$ which splits as a direct product, so that

$$\mathrm{Aut}(E) \cong H^0(\Sigma, E_2^\vee \otimes E_1) \times \mathbb{G}_m.$$

Proof. Suppose first that $E \cong \operatorname{gr} E = E_1 \oplus E_2$ and let $\psi \in \operatorname{Aut}(E_1 \oplus E_2)$. Restricting ψ to E_1 gives an automorphism of E_1, since E_1 is the canonical destabilising subbundle of E and therefore must be preserved by ψ. Moreover, the composition of the restriction of ψ to E_2 with the projection map $E_1 \oplus E_2 \to E_2$ gives an automorphism of E_2. Thus there is a map $r : \operatorname{Aut}(E_1 \oplus E_2) \to \operatorname{Aut}(E_1) \times \operatorname{Aut}(E_2) \cong \mathbb{G}_m \times \mathbb{G}_m$, and this map is surjective. Moreover, if ψ lies in the kernel of r, then $\psi - \operatorname{id}$ gives a map $E_2 \to E_1$ and thus a section in $H^0(\Sigma, E_2^\vee \otimes E_1)$. And any such section gives via addition to the identity map an automorphism of ψ lying in the kernel of r. Thus $\ker r \cong H^0(\Sigma, E_2^\vee \otimes E_1)$, and by taking the natural inclusion $\operatorname{Aut}(E_1) \times \operatorname{Aut}(E_2) \to \operatorname{Aut}(E_1 \oplus E_2)$ given by $(\phi_{11}, \phi_{22}) \mapsto \phi_{11} \oplus \phi_{22}$ we have a splitting of the short exact sequence $1 \to H^0(\Sigma, E_2^\vee \otimes E_1) \to \operatorname{Aut}(E_1 \oplus E_2) \to \mathbb{G}_m \times \mathbb{G}_m \to 1$ so that

$$\operatorname{Aut} E \cong H^0(\Sigma, E_2^\vee \otimes E_1) \rtimes (\mathbb{G}_m \times \mathbb{G}_m).$$

The matrix-type description of this group comes from identifying E with its graded $E_1 \oplus E_2$.

Next suppose that $E \ncong \operatorname{gr} E$. As in the case above, restricting an automorphism of E to E^1 gives an automorphism of E^1, so that there is a map $r : \operatorname{Aut}(E) \to \operatorname{Aut}(E^1)$. Since E^1 is stable, its automorphism group is isomorphic to $\mathbb{G}_m$. The map r is therefore surjective, since the automorphism $\lambda \operatorname{id}_E$ of E restricts to the automorphism $\lambda \operatorname{id}_{E^1}$ of E^1 for any $\lambda \in \mathbb{G}_m$. Thus there is a short exact sequence $1 \to \ker r \to \operatorname{Aut} E \to \operatorname{Aut}(E^1) \to 1$.

We can show that $\ker r \cong H^0(\Sigma, (E/E^1)^\vee \otimes E^1)$ by constructing an explicit isomorphism. Let $f \in \ker r$. Then the restriction of $f - \operatorname{id}_E$ to E^1 is the zero map, so that it induces a map $\overline{f - \operatorname{id}} : E/E^1 \to E$. Considering the post-composition of this map with the quotient map $\pi : E \to E/E^1$, we obtain a map $E/E^1 \to E/E^1$ which must be given by scalar multiplication since E/E^1 is stable. Since E is indecomposable, the short exact sequence $E^1 \to E \to E/E^1$ of vector bundles does not split, and therefore any morphism $E/E^1 \to E$ must produce the zero map upon post-composition with the quotient map $E \to E/E^1$. Thus $\overline{f - \operatorname{id}} : E/E^1 \to E$ must have image contained in E^1 and so we have a map $E/E^1 \to E^1$, or equivalently an element of $H^0(\Sigma, (E/E^1)^\vee \otimes E^1)$.

We now construct an inverse to the map $\ker r \to H^0(\Sigma, (E/E^1)^\vee \otimes E^1)$ defined above. Given $g : E/E^1 \to E^1$, we define a map $f : E \to E$ by setting $f = \operatorname{id}_E + g \circ \pi$. Its restriction to E^1 is the identity map, and moreover it has an inverse given by $\mathrm{id}_E - g \circ \pi$.

Finally, we note that if $f \in \ker r$, then $\mathrm{id}_E + \overline{f - \mathrm{id}_E} \circ \pi = f$ so that the maps $f \mapsto \overline{f - \mathrm{id}_E}$ and $g \mapsto \mathrm{id}_E + g \circ \pi$ are inverse to each other.

It remains only to show that the exact sequence $1 \to \ker r \to \mathrm{Aut}(E) \to \mathrm{Aut}(E^1) \to 1$ splits as a direct product, and to do so it suffices to construct a map $\mathrm{Aut}\, E \to \ker r$ such that post-composing it with the natural inclusion $\ker r \to \mathrm{Aut}(E)$ yields the identity map. Given an automorphism ψ of E, its restriction to E^1 is an automorphism of E^1 and thus is given by non-zero scalar multiplication. If λ is the scalar, then we can define the automorphism $\psi\lambda^{-1}$ of E, which lies in the kernel of r. This gives a splitting of the short exact sequence, so that we obtain the desired isomorphism $\mathrm{Aut}\, E \cong H^0(\Sigma, (E/E^1)^\vee \otimes E^1) \times \mathbb{G}_m)$. □

Bibliography

[1] J. Alper. On the local quotient structure of Artin stacks. *Journal of Pure and Applied Algebra*, 214(9):1576 – 1591, 2010.

[2] J. Alper. Good moduli spaces for Artin stacks. *Annales de l'Institut Fourier*, 63(6):2349–2402, 2013.

[3] J. Alper, D. Halpern-Leistner, and J. Heinloth. Existence of moduli spaces for algebraic stacks. arXiv:1812.01128 [math.AG], 2018.

[4] M. Atiyah and R. Bott. The Yang-Mills equations over Riemann surfaces. *Philosophical Transactions of the Royal Society*, 308:523–615, 1983.

[5] A. Beauville. Jacobiennes des courbes spectrales et systèmes hamiltoniens complètement intégrables. *Acta Math.*, 164:211–235, 1990.

[6] G. Bérczi, T. Hawes, F. Kirwan, and B. Doran. Geometric Invariant Theory for graded unipotent groups and applications. *Journal of Topology*, 11:826–855, 2018.

[7] G. Bérczi, T. Hawes, F. Kirwan, and B. Doran. Graded linearisations. *Proceedings of Symposia in Pure Mathematics*, 99, 2018.

[8] G. Bérczi, T. Hawes, F. Kirwan, and B. Doran. Projective completions of graded unipotent quotients. arXiv:1607.04181 [math.AG], 2020.

[9] G. Bérczi, V. Hoskins, and F. Kirwan. Stratifying quotient stacks and moduli stacks. In J. A. Christophersen and K. Ranestad, editors, *Geometry of Moduli*, pages 1–33, Cham, 2018. Springer International Publishing.

[10] G. Bérczi, J. Jackson, and F. Kirwan. Variation of Non-Reductive Geometric Invariant Theory. *Surveys in Differential Geometry*, XXII:49–69, 2018.

[11] G. Bérczi, J. Jackson, F. Kirwan, and V. Hoskins. Quotients of unstable varieties and moduli spaces of sheaves of fixed Harder-Narasimhan type, 2. In preparation.

[12] G. Bérczi and F. Kirwan. Moment maps and cohomology for graded unipotent quotients. arXiv:1909.11495 [math.AG], September 2019.

[13] A. Bialynicki-Birula. Some theorems on actions of algebraic groups. *Annals of Mathematics*, 98(3):480–497, 1973.

[14] H. Boden and K. Yokogawa. Moduli spaces of parabolic Higgs bundles and parabolic K(D) pairs over smooth curves: I. *International Journal of Mathematics*, 07:573–598, 11 1996.

[15] L. Brambila-Paz, O. Gutiérrez, and N. Nitsure. Moduli stacks and moduli schemes for rank 2 unstable bundles. arXiv:0911.2301 [math.AG], 11 2009.

[16] M. Brion and E. Peyre. The Virtual Poincaré Polynomials of Homogeneous Spaces. *Compositio Mathematica*, 134(3):319–335, 2002.

[17] S. Casalaina-Martin and M. Teixidor i Bigas. Singularities of Brill-Noether loci for vector bundles on a curve. *Mathematische Nachrichten*, 284(14-15):1846–1871, October 2011.

[18] S. Casalaina-Martin and J. Wise. *The Geometry, Topology and Physics of Moduli Spaces of Higgs Bundles*, volume 36 of *Lecture Note Series, Institute for Mathematical Sciences, National University of Singapore*, chapter An Introduction to Moduli Stacks, with a View towards Higgs Bundles on Algebraic Curves, pages 199–399. World Scientific Publishing Company, 09 2018.

[19] H. Derksen and G. Kemper. *Computational Invariant Theory*, volume 130 of *Encyclopaedia of Mathematical Sciences*. Springer-Verlag, 2007.

[20] I. Dolgachev and Y. Hu. Variation of Geometric Invariant Theory quotients. *Publications mathématiques de l'I.H.E.S*, 87:5–51, 1998.

[21] R. Donagi and T. Pantev. Langlands duality for Hitchin systems. *Inventiones mathematicae*, 189:653–735, April 2012.

[22] R. Donagi and E. Witten. Supersymmetric Yang-Mills theory and integrable systems. *Nucl. Phys.*, B460:299–334, 1996.

[23] B. Doran and F. Kirwan. Towards Non-Reductive Geometric Invariant Theory. *Pure and Applied Mathematics Quarterly*, 3(1):61–105, 2007. Special Issue: In honour of Robert MacPherson, Part 3 of 3.

[24] J.-M. Drézet and G. Trautmann. Moduli spaces of decomposable morphisms of sheaves and quotients by non-reductive groups. *Annales de l'Institut Fourier*, 53(1):107–192, 2003.

[25] R. Earl and F. Kirwan. Complete Sets of Relations in the Cohomology Rings of Moduli Spaces of Holomorphic Bundles and Parabolic Bundles Over a Riemann Surface. *Proceedings of the London Mathematical Society*, 89(3):570–622, 11 2004.

[26] D. Edidin and D. Rydh. Canonical reduction of stabilisers for Artin stacks with good moduli spaces. arXiv:1710.03220 [math.AG], October 2017.

[27] Y. Fan. Construction of the Moduli Space of Higgs Bundles using Analytic Methods. arXiv:2004.07182 [math.DG], 2020.

[28] J. Fogarty. Fixed point schemes. *Bull. Amer. Math. Soc.*, 77(2):203–204, 03 1971.

[29] E. Franco, O. Garcia-Prada, and P. E. Newstead. Higgs bundles over elliptic curves. *Illinois J. Math.*, 58(1):43–96, 2014.

[30] D. Gaiotto, G. W. Moore, and A. Neitzke. Wall-crossing, Hitchin Systems, and the WKB Approximation. *Advances in Mathematics*, 234, 07 2009.

[31] O. García–Prada and S. Ramanan. Twisted Higgs bundles and the fundamental group of compact Kähler manifolds. *Mathematical Research Letters*, 7:517–535, 2000.

[32] O. García-Prada. Higgs bundles and higher Teichmüller spaces. arXiv:1901.09086 [math.AG], 01 2019.

[33] W. Goldman. Higgs bundles and geometric structures on surfaces. *The Many Facets of Geometry: A Tribute to Nigel Hitchin*, 05 2008.

[34] W. M. Goldman and E. Z. Xia. *Rank one Higgs bundles and representations of fundamental groups of Riemann surfaces.* Number 904 in Memoirs of the American Mathematical Society. American Mathematical Society, 2008.

[35] P. Gothen and R. Zuniga-Rojas. Stratifications on the Moduli Space of Higgs Bundles. *Portugaliae Mathematica*, 74, 11 2015.

[36] P. B. Gothen and A. G. Oliveira. The Singular Fiber of the Hitchin Map. *International Mathematics Research Notices*, 2013(5):1079–1121, 02 2012.

[37] H. Gray and B. Mansoulié. The Higgs boson: the hunt, the discovery, the study and some future perspectives. ATLAS Feature, CERN, July 2018. https://atlas.cern/updates/atlas-feature/higgs-boson.

[38] A. Grothendieck. Sur La Classification Des Fibrés Holomorphes Sur La Sphère de Riemann. *American Journal of Mathematics*, 79(1):121–138, 1957.

[39] A. Grothendieck. Techniques de construction et théorèmes d'existence en géometrie algébrique iv: les schémas de Hilbert. *Séminaire Bourbaki*, 6:249–276, 1960-1961.

[40] S. Gurjar and N. Nitsure. Schematic Harder-Narasimhan stratification for families of principal bundles and Λ -modules. *Proc. Indian Acad. Sci. (Math. Sci.)*, 124(3):315–332, August 2014.

[41] P. Hacking, R. Laza, and D. Oprea. *Compactifying Moduli Spaces*, volume VII of *2297-0304*. Birkhäuser Basel, 2016.

[42] F. Haiden, L. Katzarkov, M. L. Kontsevich, and P. Pandit. Semistability, modular lattices, and iterated logarithms. arXiv:1706.01073 [math.RT], June 2017.

[43] D. Halpern-Leistner. The equivariant Verlinde formula on the moduli of Higgs bundles. arXiv: 1608.01754 [math.AG], 8 2016.

[44] D. Halpern-Leistner. On the structure of instability in moduli theory. arXiv:1411.0627v4 [math.AG], July 2018.

[45] E. Hamilton. Stratifications and quasi-projective coarse moduli spaces for the stack of Higgs bundles. arXiv:1911.13194 [math.AG], December 2019.

[46] T. Hausel. Compactification of moduli of Higgs bundles. *J. Reine Angew. Math*, 503:169–192, 1998.

[47] T. Hausel and F. Rodriguez Villegas. Cohomology Of Large Semiprojective Hyperkähler Varieties. *Asterisque*, 2015(370), 09 2013.

[48] T. Hausel and M. Thaddeus. Mirror symmetry, Langlands duality, and the Hitchin system. *Inventiones mathematicae*, 153(1):197–229, July 2003.

[49] J. Heinloth. A conjecture of hausel on the moduli space of higgs bundles on a curve. *Asterisque*, 2015:157–175, 01 2015.

[50] J. Heinloth. Hilbert-Mumford stability on algebraic stacks and applications to $\mathcal{G}$-bundles on curves. *Épijournal de Géométrie Algébrique*, 1(11), 09 2017. arXiv:1609.06058 [math.AG].

[51] N. Hitchin. The self-duality equations on a Riemann surface. *Proceedings of the London Mathematical Society*, 3(55):59–126, 1987.

[52] N. Hitchin. Stable bundles and integrable systems. *Duke Mathematical Journal*, 54(1):91–114, 1987.

[53] V. Hoskins. *Moduli spaces of complexes of sheaves.* PhD thesis, University of Oxford, 2011.

[54] V. Hoskins. Stratifications for moduli of sheaves and moduli of quiver representations. arXiv:1407.4057 [math.AG], July 2014.

[55] 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*, 3(105):852–890, 2012.

[56] D. Huybrechts and M. Lehn. *The Geometry of Moduli Spaces of Sheaves.* Cambridge Mathematical Library. Cambridge University Press, 2 edition, 2010.

[57] B. Iversen. A fixed point formula for action of tori on algebraic varieties. *Inventiones mathematicae*, 16:229–236, 1972.

[58] I. Iwanari. Stable points on algebraic stacks. *Advances in Mathematics*, 223(1):257 – 299, 2010.

[59] J. Jackson. *Moduli Spaces of Unstable Curves and Sheaves via Non-Reductive Geometric Invariant Theory*. PhD thesis, University of Oxford, 2018.

[60] A. Jaffe and E. Witten. Quantum Yang-Mills Theory, 2000.

[61] L. Katzarkov, D. Orlov, and T. Pantev. Notes on Higgs Bundles and D-Branes. Available at http://archive.schools.cimpa.info/archivesecoles/20131126152208/tony_pantev.pdf, 2013.

[62] S. Keel and S. Mori. Quotients by groupoids. *Annals of Mathematics*, 145(1):193–213, 1997.

[63] A. D. King. Moduli of representations of finite-dimensional algebras. *The Quarterly Journal of Mathematics*, 45(4):515–530, 12 1994.

[64] F. Kirwan. *Cohomology of Quotients in Symplectic and Algebraic Geometry*. Mathematical Notes, Princeton University Press, 1984.

[65] F. Kirwan. On spaces of maps from Riemann surfaces to Grassmannians and applications to the cohomology of moduli of vector bundles. *Arkiv för Matematik*, 24(1-2):221–275, 1985.

[66] F. Kirwan. Partial desingularisations of quotients of nonsingular varieties and their Betti numbers. *Annals of Mathematics*, 122:41–85, 1985.

[67] F. Kirwan. On the Homology of Compactifications of Moduli Spaces of Vector Bundles Over a Riemann Surface. *Proceedings of the London Mathematical Society*, s3-53(2):237–266, September 1986.

[68] F. Kirwan. Rational intersection cohomology of quotient varieties I. *Inventiones mathematicae*, 86:471–506, 1986.

[69] F. Kirwan. Rational intersection cohomology of quotient varieties II. *Inventiones mathematicae*, 90:153–167, 1987.

[70] F. Kirwan. Moduli spaces of bundles over Riemann surfaces and the Yang-Mills stratification revisited. *Clay Math. Proc.*, 3:239–283, 2004.

[71] F. C. Kirwan. Cohomology of Moduli Spaces. arXiv:0305344 [math.AG], 2002.

[72] F. C. Kirwan. *The Breadth of Symplectic and Poisson Geometry. Progress in Mathematics.*, volume 232, chapter Refinements of the Morse stratification of the normsquare of the moment map, pages 327–362. Birkhäuser Boston, 2005.

[73] M. Kontsevich and Y. Soibelman. *Homological Mirror Symmetry and Tropical Geometry. Lecture Notes of the Unione Matematica Italiana,*, chapter Wall-Crossing Structures in Donaldson-Thomas Invariants, Integrable Systems and Mirror Symmetry. Springer, Cham, 03 2014.

[74] E. Markman. Spectral curves and integrable systems. *Compositio Mathematica*, 93(3):255–290, 1994.

[75] G. D. Mostow. On a Conjecture of Montgomery. *Annals of Mathematics*, 65(3):513–516, 1957.

[76] D. Mumford, J. Fogarty, and F. Kirwan. *Geometric Invariant Theory.* Springer, 1994.

[77] M. Nagata. On the 14-th problem of hilbert. *American Journal of Mathematics*, 81(3):766–772, 1959.

[78] M. Nagata. Invariants of group in an affine ring. *J. Math. Kyoto Univ.*, 3(3):369–378, 1963.

[79] N. A. Nekrasov and S. L. Shatashvili. Quantum integrability and supersymmetric vacua. *Prog. Theor. Phys. Suppl.*, 177:105–119, 2009.

[80] P. Newstead. *Introduction to Moduli Problems and Orbit Spaces.* Narosa Publishing House, 1978 (reprint 2012).

[81] N. Nitsure. Moduli space of semistable pairs on a curve. *Proceedings of the London Mathematical Society*, s3(62):275–300, March 1991.

[82] N. Nitsure. Schematic Harder-Narasimhan Stratification. *International Journal of Mathematics*, 22(10):1365–1373, 2011.

[83] V. Platonov, A. Rapinchuk, and R. Rowen. *Algebraic Groups and Number Theory.* Academic Press, 1993.

[84] S. Ramanan, M. Narasimhan, and A. Beauville. Spectral curves and the generalised Theta divisor. *Journal für die reine und angewandte Mathematik*, 398:169–179, 1989.

[85] S. Rayan. Co-Higgs bundles on $\mathbb{P}^1$. *New York J. Math.*, 19:925–945, 2013.

[86] S. Rayan. The quiver at the bottom of the twisted nilpotent cone on $\mathbb{P}^1$. *European Journal of Mathematics*, 3, 09 2016.

[87] S. Rayan. Aspects of the Topology and Combinatorics of Higgs Bundle Moduli Spaces. *Symmetry Integrability and Geometry-methods and Applications*, 14:129, 2018.

[88] Z. Reichstein. Stability and equivariant maps. *Inventiones mathematicae*, 96(2):349–383, Jun 1989.

[89] A. Schmitt. Projective Moduli for Hitchin Pairs. *International Journal of Mathematics*, 9(1):107–118, 1998.

[90] H. Schoutens. *The Use of Ultraproducts in Commutative Algebra.* Lecture Notes in Mathematics 1999. Springer-Verlag Berlin Heidelberg, 1 edition, 2010.

[91] C. S. Seshadri. Geometric reductivity over arbitrary base. *Advances in Mathematics*, 26:225–274, 1977.

[92] S. Shatz. The decomposition and specialization of algebraic families of vector bundles. *Compositio Mathematica*, 35(2):163–187, 1977.

[93] C. Simpson. Higgs bundles and local systems. *Publications mathématiques de l'I.H.E.S*, 75:5–95, 1992.

[94] C. Simpson. Moduli of representations of the fundamental group of a smooth projective variety II. *Publications Mathématiques de l'IHÉS*, 80:5–79, 1994.

[95] C. Simpson. The Hodge filtration on nonabelian cohomology. *Algebraic Geometry*, 2:217–284, 1997.

[96] C. T. Simpson. Moduli of representations of the fundamental group of a smooth projective variety I. *Publications mathématiques de l'IHÉS*, 79, 1994.

[97] M. Thaddeus. Geometric Invariant Theory and flips. *Journal of the American Mathematical Society*, 9(3):691–723, July 1996.

[98] R. Vakil. The Rising Sea. Foundations of Algebraic Geometry. math216.wordpress.com, November 2017.

[99] R. Wentworth. *Geometry and Quantization of Moduli Spaces*, chapter Higgs Bundles and Local Systems on Riemann Surfaces. Birkhäuser, Cham, 2016.

[100] E. Witten. Mirror Symmetry, Hitchin's Equations, And Langlands Duality. *The Many Facets of Geometry: A Tribute to Nigel Hitchin*, 03 2008.