Small projective spaces and Stillman uniformity for sheaves

Daniel Erman, Steven V Sam and Andrew Snowden

ABSTRACT

We prove an analog of Ananyan–Hochster’s small subalgebra theorem in the context of sheaves on projective space, and deduce from this a version of Stillman’s conjecture for cohomology tables of sheaves. The main tools in the proof are Draisma’s **GL**-noetherianity theorem and the BGG correspondence.

1. Introduction

Ananyan and Hochster [AH20] recently established strong finiteness results on the Betti tables of graded modules, including Stillman’s conjecture. In Boij–Söderberg theory, Betti tables of graded modules play a dual role to cohomology tables of coherent sheaves [ES09, SE10]. It is therefore natural to look for analogs of the results of Ananyan–Hochster for cohomology tables. The purpose of this paper is to establish such results.

1.1 The work of Ananyan–Hochster

To put our results in context, we first review the work of Ananyan–Hochster. Fix a field $\mathbf{k}$, and let R be a finitely generated $\mathbf{k}$ -algebra with $R_0 = \mathbf{k}$ and $R_n = 0$ for $n < 0$. Let M be a finitely generated graded R -module, generated in degrees greater than or equal to d . Define $\beta_{i,j}(M) = \dim_{\mathbf{k}} \operatorname{Tor}_i^R(M, \mathbf{k})_j$, where the subscript here indicates the j th graded piece. These numbers are arranged into a table—the *Betti table* of M —as follows:

$$\beta(M) = \begin{array}{cccccc} \beta_{0,d} & \beta_{1,d+1} & \beta_{2,d+2} & \beta_{3,d+3} & \beta_{4,d+4} & \cdots \\ \beta_{0,d+1} & \beta_{1,d+2} & \beta_{2,d+3} & \beta_{3,d+4} & \beta_{4,d+5} & \cdots \\ \vdots & & \vdots & & \vdots & \\ \beta_{0,n} & \beta_{1,n+1} & \beta_{2,n+2} & \beta_{3,n+3} & \beta_{4,n+4} & \cdots \\ \vdots & & \vdots & & \vdots & \end{array}$$

Some important properties of M can be read off from the Betti table: the first two columns record the degrees of generators and relations; the Castelnuovo–Mumford regularity is encoded by the lowest row with a nonzero entry, that is, by $\max\{j - i \mid \beta_{i,j} \neq 0\}$; and the projective

Received 10 February 2020, accepted in final form 17 October 2020.

2020 *Mathematics Subject Classification* 14F05, 13D02.

Keywords: Stillman’s conjecture, sheaf cohomology.

This journal is © **Foundation Compositio Mathematica** 2021. This article is distributed with Open Access under the terms of the **Creative Commons Attribution Non-Commercial License**, which permits non-commercial reuse, distribution, and reproduction in any medium, provided that the original work is properly cited. For commercial re-use, please contact the **Foundation Compositio Mathematica**.

D. Erman was partially supported by NSF DMS-1601619. S. V Sam was partially supported by NSF DMS-1849173 and a Sloan Fellowship. A. Snowden was supported by NSF DMS-1453893.

dimension is the index of the last nonzero column. Moreover, the graded Betti numbers provide a refinement of the Hilbert polynomial of M .

Betti tables are invariant under flat extension of scalars: that is, if $R \rightarrow S$ is a flat morphism of $\mathbf{k}$ -algebras (of the sort under consideration) and M is an R -module, then $\beta(M) = \beta(S \otimes_R M)$. Therefore, to control $\beta(M)$ one could try to descend M to a more manageable ring. This kind of descent is exactly what Ananyan–Hochster achieve.

THEOREM 1.1 (Ananyan–Hochster). *Let $\mathbf{b}$ and $\mathbf{b}'$ be column vectors. Then there exist integers n and d with the following property: Suppose that S is a standard-graded polynomial algebra over $\mathbf{k}$ (in however many variables) and M is a graded S -module such that the first two columns of $\beta(M)$ are $\mathbf{b}$ and $\mathbf{b}'$. Then there exist a graded polynomial algebra S_0 , generated by at most n variables of degrees at most d , a flat $\mathbf{k}$ -algebra homomorphism $S_0 \rightarrow S$, and a graded S_0 -module M_0 such that $M \cong S \otimes_{S_0} M_0$.*

Over an algebraically closed field, this result follows by applying [AH20, Theorem B] to the vector space spanned by the polynomials in a presentation matrix of M ; see [ESS18, §6.6] for the case of an arbitrary field. As suggested, the theorem allows one to control Betti tables, via standard noetherianity arguments over S_0 . Alternately, one could adapt [ESS19a, Theorem 5.14]. The precise statement is as follows.

COROLLARY 1.2. *Let $\mathbf{b}$ and $\mathbf{b}'$ be given. Then there exists a finite set $\mathcal{B} = \mathcal{B}(\mathbf{b}, \mathbf{b}')$ of tables such that if S and M are as in Theorem 1.1, then $\beta(M) \in \mathcal{B}$. In other words, given the first two columns of $\beta(M)$, there are only finitely many possibilities for the whole table.*

Stillman’s conjecture asserts that the projective dimension of a module over a standard-graded polynomial ring can be bounded from the first two columns of its Betti table. It therefore follows immediately from the corollary. We emphasize that the main point of the corollary and Stillman’s conjecture is that the bounds are independent of the number of variables in the polynomial ring, and depend only on the number and degrees of generators and relations of the module.

1.2 Main results

We establish analogs of the above results in the setting of coherent sheaves. As noted above, this is a natural setting to seek such analogs, due to the duality in Boij–Söderberg theory [ES09, SE10, EE17]. (Koszul duality is another natural setting to ask about such analogs, but McCullough provided a negative result in that context [McC19].)

We begin with a definition. Let X be a projective variety over $\mathbf{k}$ equipped with an ample line bundle $\mathcal{O}_X(1)$, and let $\mathcal{E}$ be a coherent sheaf on X . Define $\gamma_{i,j}(\mathcal{E}) = \dim_{\mathbf{k}} H^i(X, \mathcal{E}(j))$. These numbers are arranged into a table—the *cohomology table* of $\mathcal{E}$ —as follows:

$$\gamma(\mathcal{E}) = \begin{array}{ccccccc} \cdots & \gamma_{n,-n-2} & \gamma_{n,-n-1} & \gamma_{n,-n} & \gamma_{n,1-n} & \gamma_{n,2-n} & \cdots \\ \cdots & \gamma_{n-1,-n-1} & \gamma_{n-1,-n} & \gamma_{n-1,1-n} & \gamma_{n-1,2-n} & \gamma_{n-1,3-n} & \cdots \\ & \vdots & & \vdots & & \vdots & \\ \cdots & \gamma_{1,-3} & \gamma_{1,-2} & \gamma_{1,-1} & \gamma_{1,0} & \gamma_{1,1} & \cdots \\ \cdots & \gamma_{0,-2} & \gamma_{0,-1} & \gamma_{0,0} & \gamma_{0,1} & \gamma_{0,2} & \cdots \end{array}$$

Unlike Betti tables, cohomology tables extend infinitely to the left and right. The number of nonzero rows is bounded by the dimension of the support of $\mathcal{E}$. We let $\gamma^k(\mathcal{E})$ denote the k th column of $\gamma(\mathcal{E})$, that is, the sequence $(\gamma_{0,k}(\mathcal{E}), \gamma_{1,k-1}(\mathcal{E}), \dots, \gamma_{i,k-i}(\mathcal{E}), \dots)$.

Cohomology tables are invariant under finite pushforwards: that is, if $(Y, \mathcal{O}_Y(1))$ is a second projective variety and $i: X \rightarrow Y$ is a finite morphism such that $i^*(\mathcal{O}_Y(1)) = \mathcal{O}_X(1)$, then $\gamma(i_*(\mathcal{E})) = \gamma(\mathcal{E})$. Thus, to control the cohomology table of $\mathcal{E}$, one could try to realize $\mathcal{E}$ as a pushforward from a simpler variety. This is exactly our main result.

THEOREM 1.3. *Let $\mathbf{b}$ and $\mathbf{b}'$ be column vectors. Then there exists an integer $N = N(\mathbf{b}, \mathbf{b}')$ with the following property: Suppose that $\mathcal{E}$ is a coherent sheaf on $\mathbf{P}^n$, for some n , such that the zeroth and first columns of $\gamma(\mathcal{E})$ are $\mathbf{b}$ and $\mathbf{b}'$. Then there exist a linear embedding $i: \mathbf{P}^{n_0} \rightarrow \mathbf{P}^n$ with $n_0 \leq N$ and a coherent sheaf $\mathcal{E}_0$ on $\mathbf{P}^{n_0}$ such that $\mathcal{E} \cong i_*(\mathcal{E}_0)$.*

Note that the conclusion of this theorem is stronger than that of Theorem 1.1 as it realizes the sheaf as the pushforward under a *linear* embedding. As before, this theorem allows us to control cohomology tables.

COROLLARY 1.4. *Let $\mathbf{b}$ and $\mathbf{b}'$ be given. Then there exists a finite set $\Gamma = \Gamma(\mathbf{b}, \mathbf{b}')$ of tables such that if $\mathcal{E}$ is as in Theorem 1.3, then $\gamma(\mathcal{E}) \in \Gamma$. In other words, given the zeroth and first columns of $\gamma(\mathcal{E})$, there are only finitely many possibilities for the whole table.*

As a special case of this corollary, we obtain a Stillman-like conjecture for sheaves.

COROLLARY 1.5. *Let $\mathbf{b}$ and $\mathbf{b}'$ be given. Then there exists an integer $r = r(\mathbf{b}, \mathbf{b}')$ with the following property: if $\mathcal{E}$ is as in Theorem 1.3, then the Castelnuovo–Mumford regularity of $\mathcal{E}$ is at most r .*

In fact, we prove Corollary 1.5 first and use it in our proof of Theorem 1.3.

Remark 1.6. In Theorem 1.3, we fixed the zeroth and first columns of $\beta(\mathcal{E})$. In fact, one could fix any two consecutive columns. This follows from Theorem 1.3 and the fact that the cohomology table of a twist $\mathcal{E}(n)$ is obtained from that of $\mathcal{E}$ by shifting all columns to the left by n . For Betti tables, however, one must use the first two columns (see Remark 4.7).

1.3 Overview of the proof

The basic idea is to define an infinite-dimensional parameter space for the sheaves $\mathcal{E}$ under consideration and use Draisma’s **GL**-noetherianity result [Dra19] to control this space. We provide some more details.

First, a coherent sheaf $\mathcal{E}$ on projective space can be encoded by its associated Tate resolution $\mathbf{T}(\mathcal{E})$, which is a doubly infinite, everywhere exact complex over an exterior algebra. This Tate resolution can, in turn, be encoded by any single differential. This is reviewed in Section 2.

The initial data $\mathbf{b}, \mathbf{b}'$ determines the degrees of the entries of the zeroth differential of the Tate resolution. We use this in Section 3.1 to define a parameter space for coherent sheaves with these columns in their cohomology table. This parameter space, which is infinite-dimensional since we do not a priori fix the number of variables, is contained in a finite direct sum of exterior powers.

Inside of our parameter space, we define Z_i to be the locus where the regularity of the corresponding sheaf is strictly bigger than i . In Corollary 3.10, we show that Z_i is Zariski closed. This is one of the main technical points in the paper and where we diverge from McCullough’s counterexamples [McC19]. If we were to study the analogous loci in the context of [McC19], then the result would fail. As McCullough explicitly shows, for a 1×1 matrix consisting of a single quadratic form in an exterior algebra, the locus in $\bigwedge^2 \mathbf{k}^\infty$ where the cokernel has regularity at most i is actually Zariski open, not closed. The key fact that enables us to prove that Z_i is Zariski

closed in our case is that the matrices that arise in a Tate resolution are not generic, and so our ambient parameter space is a very specific sublocus inside of a direct sum of exterior powers; see Lemma 2.3.

We then have a chain of closed subsets $Z_1 \supseteq Z_2 \supseteq \cdots$. Draisma’s noetherianity result [Dra19] implies that $Z_r = Z_{r+1}$ for $r \gg 0$. Since any coherent sheaf on a projective space has finite regularity, we see that $Z_r = \emptyset$ for $r \gg 0$. This gives a bound on regularity, yielding Corollary 1.5.

We then deduce Theorem 1.3 from Corollary 1.5. For this, we use the fact that differentials beyond the regularity of a coherent sheaf in the Tate resolution are linear maps. (This is shown in [EFS03] and reviewed in Section 2.) Using the regularity bound above, we can examine a fixed differential for all of our coherent sheaves. Repeated applications of Corollary 3.9 (which also relies on Draisma’s noetherianity result) show that the possible sizes of this differential come from a finite list, and hence there is a maximal number of linear forms used in this fixed differential. This yields Theorem 1.3.

1.4 Open questions

Our work naturally suggested several further questions. First, can one prove Theorem 1.3 via more elementary methods, for example, without appealing to infinite-dimensional noetherianity results? We have been able to handle some special cases directly, such as when the sheaf $\mathcal{E}$ is the structure sheaf of a smooth subvariety, which suggests that this indeed may be possible. Second, could one prove Theorem 1.3 via an ultraproduct argument, similar to what is done in [ESS19a]? Finally, it would be interesting to better understand how our result and that of Ananyan–Hochster interact with the duality between Betti tables and cohomology tables appearing in Boij–Söderberg theory.

1.5 Outline

In Section 2, we review the necessary material, with a specific emphasis on the Bernstein–Gel’fand–Gel’fand (BGG) correspondence and on properties of Tate resolutions. In Section 3, we construct the parameter spaces $X_{\mathbf{d}}^0$ which parametrize the coherent sheaves $\mathcal{E}$ with $\gamma^k(\mathcal{E}) = \mathbf{b}$ and $\gamma^{k+1}(\mathcal{E}) = \mathbf{b}'$. We also prove a number of key semicontinuity results about these spaces, including the fact that the space Z_i mentioned above is Zariski closed. Finally, Section 4 contains the proofs of Theorem 1.3 and Corollary 1.4.

2. Background

Throughout, we work over an arbitrary field $\mathbf{k}$.

2.1 The BGG correspondence

Here we give an overview of the key results of [EFS03] that will be relevant to us. Let $V_n = \langle x_0, \dots, x_n \rangle$ be a graded vector space of dimension $n + 1$ over $\mathbf{k}$, where $\deg(x_i) = 1$, and let $\mathrm{Sym}(V_n)$ be the symmetric algebra on V_n . Let $W_n = \langle e_0, \dots, e_n \rangle$ be the dual graded vector space with corresponding dual basis, where $\deg(e_i) = -1$. Let $\bigwedge W_n$ be the corresponding exterior algebra. We often write $S = \mathrm{Sym}(V_n)$ or $E = \bigwedge W_n$ when the number of variables is fixed.

The BGG correspondence provides a correspondence between complexes of S -modules and complexes of E -modules. It was introduced in [Bei78, BGG78], though the treatment developed in [EFS03] will be most relevant for our purposes. There is a functor $\mathbf{R}$ which transforms

a graded S -module $M = \bigoplus_d M_d$ into a graded free linear complex $\mathbf{R}(M)$ of E -modules, where $\mathbf{R}(M)^i = M_i \otimes_{\mathbf{k}} E(-i)$. The differential $\partial^i: M_i \otimes_{\mathbf{k}} E(-i) \rightarrow M_{i+1} \otimes_{\mathbf{k}} E(-i-1)$ is determined by the image of the generators $M_i \rightarrow M_{i+1} \otimes W_n$ via the formula $\partial(m) = \sum_{i=0}^n (-1)^i x_i m \otimes e_i$. Alternately, one may view the map $M_i \rightarrow M_{i+1} \otimes W_n$ as the adjoint of the multiplication map $M_i \otimes V_n \rightarrow M_{i+1}$ given by the S -module structure on M . There is a similar functor $\mathbf{L}$ which transforms a graded E -module into a graded free linear complex of S -modules. The functors $\mathbf{R}$ and $\mathbf{L}$ induce equivalences between the derived categories of S -modules and of E -modules.

2.2 Tate resolutions

Let $\mathcal{E}$ be a coherent sheaf on $\mathbf{P}^n = \text{Proj}(\text{Sym}(V_n))$. Work of Eisenbud, Fløystad, and Schreyer shows that $\gamma(\mathcal{E})$ can be viewed as the Betti table of a minimal, free, doubly infinite, everywhere exact complex on $E = \bigwedge W_n$, known as a *Tate resolution of $\mathcal{E}$* and denoted by $\mathbf{T}(\mathcal{E})$; see [EFS03, Theorem 4.1]. To be more specific, if $\mathbf{T}(\mathcal{E}) = [\cdots \rightarrow \mathbf{T}(\mathcal{E})^k \rightarrow \mathbf{T}(\mathcal{E})^{k+1} \rightarrow \cdots]$ is a Tate resolution for $\mathcal{E}$, then in cohomological degree k , we have the free module

$$\mathbf{T}(\mathcal{E})^k = \bigoplus_{i+j=k} E(-j)^{\gamma_{i,j}(\mathcal{E})}. \quad (2.1)$$

The graded Betti numbers of the free module $\mathbf{T}(\mathcal{E})^k$ are thus given by the entries of the cohomology table in a single column.

Example 2.2. Let $C \subseteq \mathbf{P}^2$ be a cubic curve. The cohomology table of its structure sheaf is

		$\gamma^{-2}(\mathcal{O}_C)$	$\gamma^{-1}(\mathcal{O}_C)$	$\gamma^0(\mathcal{O}_C)$	$\gamma^1(\mathcal{O}_C)$	$\gamma^2(\mathcal{O}_C)$	
$\gamma(\mathcal{O}_C) =$	2	$\cdots$	0	0	0	0	$\cdots$
	1	$\cdots$	9	6	3	1	$\cdots$
	0	$\cdots$	0	0	1	3	$\cdots$

A Tate resolution $\mathbf{T}(\mathcal{O}_C)$ for $\mathcal{O}_C$ has the form

$$\cdots \xrightarrow{\partial^{-3}} E(3)^9 \xrightarrow{\partial^{-2}} E(2)^6 \xrightarrow{\partial^{-1}} E(1)^3 \oplus E^1 \xrightarrow{\partial^0} E^1 \oplus E(-1)^3 \xrightarrow{\partial^1} E(-2)^6 \rightarrow \cdots.$$

By fixing the entries in columns 0 and 1 of $\mathbf{T}(\mathcal{E})$ as in Theorem 1.3, we are fixing the degrees of the entries of a matrix representing the differential ∂^0 from $\mathbf{T}(\mathcal{E})$. McCullough's examples in [McC19, Theorem 4.1] show that this alone is not sufficient to obtain finiteness results about the cokernel of a matrix over E . In particular, he shows that for each $\ell \geq 1$, the ideal generated by the single quadric $e_1 \wedge e_2 + e_3 \wedge e_4 + \cdots + e_{2\ell-1} \wedge e_{2\ell}$ has regularity $\ell - 2$. Thus, even for a 1×1 matrix ϕ containing a single quadric in the exterior algebra, there is no way to bound the regularity of $\text{coker } \phi$ solely in terms of the degrees of the entries of ϕ (that is, without reference to the number of variables).

But the matrices that arise in Tate resolutions turn out to be far from generic. To describe this, we let $\partial^k: \mathbf{T}(\mathcal{E})^k \rightarrow \mathbf{T}(\mathcal{E})^{k+1}$ be one of the differentials in a Tate resolution and let $\phi^\vee$ be a matrix representing the dual of ∂^k , where the duality functor is $\text{Hom}_E(-, E)$. As the following lemma shows, it is always the case that $\text{coker } \phi^\vee$ has regularity $-k$, where we mean regularity as a graded E -module. (This should not be confused with the Castelnuovo–Mumford regularity of the sheaf $\mathcal{E}$.) In other words, the Betti table of $\text{coker } \phi^\vee$ has a built-in floor, and so we avoid McCullough's counterexamples. The lemma follows easily from [EFS03, Theorem 4.1]. But it will play an important role, and so we give an independent statement and proof.

LEMMA 2.3. Fix a coherent sheaf $\mathcal{E}$ on $\mathbf{P}^n$, and fix some integer k . Let ∂^k denote the differential $\mathbf{T}(\mathcal{E})^k \rightarrow \mathbf{T}(\mathcal{E})^{k+1}$, and let $\phi^\vee: (\mathbf{T}(\mathcal{E})^{k+1})^* \rightarrow (\mathbf{T}(\mathcal{E})^k)^*$ be a matrix representing the dual of ∂^k . The regularity of $\text{coker}(\phi^\vee)$ as a graded E -module is $-k$.

In particular, if $\mathbf{G}$ is a minimal free resolution of $\text{coker}(\phi^\vee)$, then $\mathbf{G}_j$ is generated entirely in degree $-j - k$ for $j \gg 0$.

Proof. It suffices to prove that $\mathbf{G}_j$ is generated entirely in degree $-j - k$ for $j \gg 0$, as this will imply that the regularity of $\text{coker}(\phi^\vee)$ is $-k$. Now, since $\mathbf{T}(\mathcal{E})$ is doubly exact and since E is self-injective, we have that $\mathbf{G}_j = (\mathbf{T}(\mathcal{E})^{k+j})^*$. It thus suffices to show that $\mathbf{T}(\mathcal{E})^{k+j}$ is generated entirely in degree $k+j$ for fixed k and $j \gg 0$. This follows by combining (2.1) with Serre vanishing for $\mathcal{E}$. $\square$

2.3 Cone-stable invariants

We will consider numerical functions ν whose domain is the set of pairs (E, M) , where E is an exterior algebra over $\mathbf{k}$ generated by finitely many elements all of the same degree and M is a finitely generated graded E -module, and whose target is $\mathbf{Z} \cup \{\infty\}$. Then ν is a *module invariant* if $\nu(M)$ depends only on the pair (E, M) up to isomorphism. Let $\mathbf{k}[\varepsilon]/(\varepsilon^2)$ be the exterior algebra on one generator. We say that a module invariant ν is *cone-stable* if $\nu(M) = \nu(M \otimes_{\mathbf{k}} \mathbf{k}[\varepsilon]/(\varepsilon^2))$.

Finally, ν is *weakly upper semicontinuous* if it is upper semicontinuous in flat families. More precisely, given a variety V over $\mathbf{k}$, let $\mathcal{M}$ be a flat module over the sheaf of algebras $\mathcal{O}_V \otimes_{\mathbf{k}} E$. We require that the function $x \mapsto \nu(\mathcal{M}_x)$ be upper semicontinuous; that is, for all n , $\{x \mid \nu(\mathcal{M}_x) \geq n\}$ is closed.

3. Parameter spaces and semicontinuity results

3.1 Parameter spaces

Since we are interested in the properties of cohomology tables on projective space, without a priori fixing the number of variables, we will introduce certain limiting topological spaces. This is similar to the setup in [ESS19b].

We fix for this section two column vectors $\mathbf{b} = (\mathbf{b}_0, \mathbf{b}_1, \dots)$ and $\mathbf{b}' = (\mathbf{b}'_0, \mathbf{b}'_1, \dots)$, and let s be the maximal index such that $\mathbf{b}_s$ or $\mathbf{b}'_s$ is nonzero. To begin, we consider $E = \bigwedge W_n$ for a fixed n . The two vectors determine free modules $F = \bigoplus_{i \geq 0} E(i)^{\mathbf{b}_i}$ and $F' = \bigoplus_{i \geq 0} E(i-1)^{\mathbf{b}'_i}$. A homogeneous map $F \rightarrow F'$ is given by a matrix ϕ with entries of specified degrees. We let $\mathbf{d}_i$ represent the number of entries of degree $-i$ that could appear in such a matrix, and we write $\mathbf{d} = (\mathbf{d}_1, \mathbf{d}_2, \dots)$. Note that $\mathbf{d}$ has only finitely many nonzero entries. Since the matrices we are interested in will arise from minimal free resolutions, there will be no nonzero entries of degree zero, and thus we do not include a $\mathbf{d}_0$. Let $X_{\mathbf{d},n} = (\bigwedge^1 W_n)^{\oplus \mathbf{d}_1} \oplus (\bigwedge^2 W_n)^{\oplus \mathbf{d}_2} \oplus \dots$ be the corresponding parameter space. We will identify points $x \in X_{\mathbf{d},n}$ with matrices $\phi: F \rightarrow F'$ and sometimes abuse notation by writing $\phi \in X_{\mathbf{d},n}$. We refer to such a matrix $\phi: F \rightarrow F'$ as a *matrix of type $(\mathbf{b}, \mathbf{b}')$* . An example will help clarify this setup.

Example 3.1. Following Example 2.2, let $\mathbf{b} = (1, 3, 0, \dots)$ and $\mathbf{b}' = (3, 1, 0, \dots)$. Then, with the notation of the previous paragraph, $F = E^1 \oplus E(1)^3$ and $F' = E(-1)^3 \oplus E^1$. A homogeneous map $F \rightarrow F'$ is thus determined by a block 4×4 matrix which consists of six entries of degree -1 (corresponding to submatrices $E^1 \rightarrow E(-1)^3$ and $E(1)^3 \rightarrow E^1$) and nine entries of degree -2 (corresponding to $E(1)^3 \rightarrow E(-1)^3$). So we set $\mathbf{d} = (6, 9, 0, 0, \dots)$. Based only on degree considerations, this 4×4 matrix also has a degree zero entry. However, the matrices that

arise in Tate resolutions never have units, and since we are parametrizing such matrices, we set all such entries to be zero. A matrix of type $(\mathbf{b}, \mathbf{b}')$ thus corresponds to a 4×4 matrix of the form

$$\begin{pmatrix} a_1 & b_{1,1} & b_{1,2} & b_{1,3} \\ a_2 & b_{2,1} & b_{2,2} & b_{2,3} \\ a_3 & b_{3,1} & b_{3,2} & b_{3,3} \\ 0 & a_4 & a_5 & a_6 \end{pmatrix},$$

where the a_i are elements in E of degree -1 and the $b_{j,k}$ are elements in E of degree -2 .

Now we take a limit as $n \rightarrow \infty$. Let W_∞ be the direct limit of the spaces W_n , under the standard inclusions, and let $E_\infty = \bigwedge W_\infty$. The inclusions $W_n \rightarrow W_{n+1}$ also induce inclusions $X_{\mathbf{d},n} \rightarrow X_{\mathbf{d},n+1}$, and we let $X_{\mathbf{d}}$ be the direct limit, equipped with the direct limit topology. Recall that a topological space Z with the action of a group G is said to be G -noetherian if every descending chain of G -invariant closed subspaces of Z stabilizes. Since each space $X_{\mathbf{d},n}$ has a natural action of $\mathbf{GL}_n$ on it, which is induced by the action of $\mathbf{GL}_n$ on W_n , we obtain an induced action of $\mathbf{GL} := \bigcup_n \mathbf{GL}_n$ on $X_{\mathbf{d}}$. We have that $X_{\mathbf{d}}$ is $\mathbf{GL}$ -noetherian by [Dra19] and [ESS19b, Corollary 2.8]. As above, we identify points $\phi \in X_{\mathbf{d}}$ with matrices, in this case with entries in E_∞ . We can thus view the space $X_{\mathbf{d}}$ as parametrizing matrices of type $(\mathbf{b}, \mathbf{b}')$.

This is not yet the parameter space that we want to work with. As noted in Lemma 2.3, the matrices that appear in a Tate resolution with specified columns will have built-in regularity bounds. We thus want to pass to the sublocus of $X_{\mathbf{d}}$ of matrices satisfying that same regularity bound, as this will allow us to parametrize coherent sheaves with specified columns in its cohomology table.

We will say that a coherent sheaf $\mathcal{E}$ on $\mathbf{P}^n$ has *type* $(\mathbf{b}, \mathbf{b}')$ if $\gamma^0(\mathcal{E}) = \mathbf{b}$ and $\gamma^1(\mathcal{E}) = \mathbf{b}'$. If some $\phi \in X_{\mathbf{d}}$ comes from a Tate resolution of a type $(\mathbf{b}, \mathbf{b}')$ sheaf $\mathcal{E}$, then by Lemma 2.3, it must be the case that the regularity of $\text{coker } \phi^\vee$ is zero. So we let

$$X_{\mathbf{d}}^0 = \{\phi \in X_{\mathbf{d}} \mid \text{reg coker } \phi^\vee = 0\}$$

be this subset. We endow $X_{\mathbf{d}}^0$ with the subspace topology, and thus $X_{\mathbf{d}}^0$ is also $\mathbf{GL}$ -noetherian. A Tate resolution for a coherent sheaf $\mathcal{E}$ of type $(\mathbf{b}, \mathbf{b}')$ thus induces an element $\phi \in X_{\mathbf{d}}^0$.

This construction is reversible as well. Assume that we are given some $x \in X_{\mathbf{d}}^0$, and choose n minimally such that $x \in X_{\mathbf{d},n}$. Let ϕ be the associated matrix of type $(\mathbf{b}, \mathbf{b}')$. Then, by Lemma 2.3, if $\mathbf{F}$ is the minimal free resolution of $\text{coker } \phi^\vee$, then $\mathbf{F}_j$ is generated in degree $-j$ for $j \geq k$ for some k . Writing $\mathbf{F}_{\geq k}^\vee$ for the dual complex, we let $M = \mathbf{L}(\mathbf{F}_{\geq k}^\vee)$ be the corresponding $\text{Sym}(V_n)$ -module, noting that, since $\mathbf{F}_{\geq k}^\vee$ is a linear complex, applying $\mathbf{L}$ yields a $\text{Sym}(V_n)$ -module, thought of as a complex of modules concentrated in a single homological degree [EFS03, Introduction]. Let $\mathcal{E}$ be the coherent sheaf on $\mathbf{P}^n$ corresponding to M . By construction, the truncated Tate resolution $\mathbf{T}(\mathcal{E})^{\geq k}$ is isomorphic to $\mathbf{F}_{\geq k}^\vee$ (see Section 2.2), and thus $\mathcal{E}$ has type $(\mathbf{b}, \mathbf{b}')$. In particular, by the uniqueness of minimal free resolutions, there is a Tate resolution for $\mathcal{E}$ where ϕ is the k th differential. Since $\mathcal{E}$ is determined by the point $x \in X_{\mathbf{d},n}$, we write $\mathcal{E}_x$ for the sheaf on $\mathbf{P}^n$ that results from this construction. We will see in Lemma 4.1 that we would obtain an isomorphic sheaf for any $m \geq n$, and thus any sheaf cohomology groups of the resulting sheaf $\mathcal{E}_x$ are cone-stable (see Section 2 for the definition of cone-stability and Remark 4.2 for more).

Remark 3.2. One potentially confusing point in this construction is that when we associate a matrix ϕ with a point x , the matrix ϕ represents the differential $\partial^0: \mathbf{T}^0(\mathcal{E}_x) \rightarrow \mathbf{T}^1(\mathcal{E}_x)$, but the regularity bound from Lemma 2.3 applies not to ϕ itself but to $\phi^\vee$. In fact, there is no a priori

bound on the regularity of coker ϕ . However, Theorem 1.3 implies the existence of such a bound.

3.2 Universal modules

Let $\mathcal{F} = \bigoplus_{i \geq 0} E_\infty(i)^{\mathbf{b}_i}$ and $\mathcal{F}' = \bigoplus_{i \geq 0} E_\infty(i-1)^{\mathbf{b}'_i}$ be the free E_∞ -modules corresponding to $\mathbf{b}$ and $\mathbf{b}'$. There is a map Φ of E_∞ -modules $\Phi: \mathcal{F} \rightarrow \mathcal{F}'$ which is a universal matrix of type $(\mathbf{b}, \mathbf{b}')$ in the sense that by specializing the coefficients of Φ , we can obtain any matrix $\phi \in X_{\mathbf{d}}$ of type $(\mathbf{b}, \mathbf{b}')$. Set $\mathcal{M} = \text{coker}(\Phi^\vee)$ to be the universal module. In particular, if x is a point of $X_{\mathbf{d}}$ which corresponds to the matrix $\phi: F \rightarrow F'$, then the specialization $\Phi|_x$ is precisely ϕ . Thus $\mathcal{M}_x \cong \text{coker}(\phi^\vee)$. These constructions also make sense at each finite level n , but when the context is clear, we will avoid subscripts to keep the notation manageable.

LEMMA 3.3. *Let $U_{\mathbf{d}} \subseteq X_{\mathbf{d}}$ be the locus of points x where the matrix $\phi^\vee$ corresponding to x is a minimal presentation of $\mathcal{M}_x$.*

- (a) *The subset $U_{\mathbf{d}}$ is Zariski open in $X_{\mathbf{d}}$.*
- (b) *For any point $x \notin U_{\mathbf{d}}$, there exists a matrix ψ of type $(\mathbf{b}, \mathbf{b}'')$ with $\mathbf{b}'' \leq \mathbf{b}'$ in the termwise partial order such that $\text{coker} \psi^\vee \cong \mathcal{M}_x$.*

Proof. Let $\phi \in X_{\mathbf{d}}$ be a matrix such that $\phi^\vee$ is not a minimal presentation of its cokernel. Since we have built into our construction that any such matrix has no entries that are units, this happens if and only if one of the columns of $\phi^\vee$ is in the image of the other columns. Let c be the number of columns of $\phi^\vee$; let $\phi_1^\vee, \dots, \phi_c^\vee$ be these columns, and suppose that we can write $\phi_c^\vee = \sum_{i=1}^{c-1} f_i \phi_i^\vee$. Let $\mathbf{d}'$ be the degree sequence corresponding to the entries of the matrix $(\phi_1^\vee \phi_2^\vee \cdots \phi_{c-1}^\vee)$, and let $\mathbf{d}''$ be the degree sequence corresponding to the exterior polynomials f_i . Then ϕ lies in the image of $X_{\mathbf{d}'} \times X_{\mathbf{d}''} \rightarrow X_{\mathbf{d}}$. In fact, this is equivalent to saying that $\phi_c^\vee$ is a polynomial combination of the other columns.

It follows that the complement of $U_{\mathbf{d}}$ is a finite union of such images. Part (b) is immediate from the construction, and $\mathcal{M}_x$ is just the cokernel of $(\phi_1^\vee \phi_2^\vee \cdots \phi_{c-1}^\vee)$. By [ESS19b, Proposition 2.8], to show that $U_{\mathbf{d}}$ is open, it suffices to show that $U_{\mathbf{d},n}$ is open for all n . The map $X_{\mathbf{d}',n} \times X_{\mathbf{d}'',n} \rightarrow X_{\mathbf{d},n}$ is invariant under scaling, so we may projectivize both the source and target. In that case, the map becomes proper, so the image is closed. This means that the image of the original map is an affine cone over a closed subset and hence is also closed. In particular, the complement of $U_{\mathbf{d},n}$ is a finite union of closed subsets, and hence $U_{\mathbf{d},n}$ is open. $\square$

3.3 Semicontinuity of Betti numbers

Using the universal module $\mathcal{M}$, we can obtain Zariski-closed subloci of $X_{\mathbf{d},n}$ via any semi-continuous module invariant. Here is an example.

LEMMA 3.4. *For any $n, r \in \mathbf{N}$ and $e \in \mathbf{Z}$, the locus of $x \in X_{\mathbf{d},n}$ such that $\dim(\mathcal{M}_x)_e \geq r$ is Zariski closed.*

Proof. As an $\mathcal{O}_{X_{\mathbf{d},n}}$ -module, $\mathcal{M}$ splits as $\mathcal{M} = \bigoplus_{d \in \mathbf{Z}} \mathcal{M}_d$. We thus have $(\mathcal{M}_x)_e \cong (\mathcal{M}_e)_x$. The statement now follows from standard semicontinuity results for coherent sheaves [Har77, Example III.12.7.2]. $\square$

Since Hilbert function values such as $\dim(\mathcal{M}_x)_e$ are not cone-stable, we will need to do a bit of extra work to obtain the required semicontinuity results on $X_{\mathbf{d}}$. For any finitely generated,

graded $\bigwedge W_n$ -module N , we define the invariant

$$\alpha_k(N) = \sum_{i=0}^{\infty} (-1)^i \beta_{i,k}(N).$$

Note that for fixed k , the Betti number $\beta_{i,k}(N)$ is nonzero for only finitely many i , so this is really a finite sum. Since Betti numbers are cone-stable, so are the invariants α_k . Since the notation in this subsection gets a bit heavy, we provide the following remark as an overview of the results.

Remark 3.5. We summarize the notation and results from this section with the following diagram. For any $x \in U_{\mathbf{d}}$ (see Lemma 3.3), the Betti table of $\mathcal{M}_x$ will look like (here s is the maximal integer such that either $\mathbf{b}_s$ or $\mathbf{b}'_s$ is nonzero)

$$\beta(\mathcal{M}_x) = \begin{array}{c|ccccccc} & 0 & 1 & 2 & \cdots & i & \cdots & s+2 \\ \hline s & \mathbf{b}_s & \mathbf{b}'_s & \beta_{2,s-2} & \cdots & \beta_{i,s-i} & \cdots & \beta_{s+2,-2} \\ \vdots & \vdots & \vdots & \vdots & \cdots & \cdots & & \\ 1 & \mathbf{b}_1 & \mathbf{b}'_1 & \beta_{2,-1} & \cdots & \beta_{i,1-i} & \cdots & \beta_{s+2,-s-1} \\ 0 & \mathbf{b}_0 & \mathbf{b}'_0 & \beta_{2,-2} & \cdots & \beta_{i,-i} & \cdots & \beta_{s+2,-s-2} \\ -1 & - & - & \vdots & \ddots & & \ddots & \end{array} \quad (3.5a)$$

If $x \notin U_{\mathbf{d}}$, then at least one of the $\mathbf{b}'_i$ can be reduced by 1.

There are no nonzero entries in row $s+1$ since this is the Betti table of a minimal free resolution. The invariants $\alpha_k(\mathcal{M}_x)$ are obtained by taking the alternating sum of the entries lying on slope 1 diagonals, starting with $\mathbf{b}_k$. The range $-2 \leq k \leq s$ that appears in Lemma 3.7 and the related results ensure that we contain every diagonal up to and including the one passing through $\beta_{2,-2}$. This plays a key role in the later application to sheaf cohomology tables, which have a built-in floor in the zeroth row (see Lemma 2.3), and the indexing is chosen to line up with that case.

Corollary 3.9 shows that for each $-2 \leq j \leq s$, there are finitely many possibilities for the Betti numbers $\beta_{2,j}(\mathcal{M}_x)$ as x varies in $X_{\mathbf{d}}$.

LEMMA 3.6. *Let N be a finitely generated, graded $\bigwedge W_n$ -module. The Hilbert function of N in degree e is determined by $\{\alpha_k(N)\}_{k \geq e}$ via the formula*

$$\dim N_e = \sum_{j=e}^{\infty} \alpha_j(N) \binom{n+1}{j-e}. \quad (3.6a)$$

Proof. Let $\mathbf{F}$ be a minimal free resolution of N . Write $E = \bigwedge W_n$. We have

$$\dim N_e = \sum_{i=0}^{\infty} (-1)^i \dim(\mathbf{F}_i)_e,$$

where the sum is finite because $\mathbf{F}$ is minimal. We can rewrite this as

$$\dim N_e = \sum_{i=0}^{\infty} \sum_{j \in \mathbf{Z}} (-1)^i \beta_{i,j}(N) \dim E(-j)_e = \sum_{j \in \mathbf{Z}} \alpha_j(N) \dim E(-j)_e.$$

However, $E(-j)_e = 0$ if $e > j$. If $j \geq e$, then $\dim E(-j)_e = \binom{n+1}{j-e}$. This yields $\dim N_e = \sum_{j=e}^{\infty} \alpha_j(N) \binom{n+1}{j-e}$. $\square$

It is well known that Betti numbers are weakly semicontinuous; that is, they are semicontinuous in families with a fixed Hilbert function. In fact, an elementary argument shows that if one

restricts to a family which has fixed Hilbert function values in degrees greater than k , then each Betti number $\beta_{i,j}$ with $j > k$ is semicontinuous. The following lemma shows that one can do even a bit better than this if one only cares about Betti numbers $\beta_{i,j}$ with $i \geq 2$. Note that, by the definition of $\mathbf{b}$, each Betti number $\beta_{i,j}(\mathcal{M}_x)$ will be zero for $j > s$ by Remark 3.5. In particular, $\alpha_j(\mathcal{M}_x) = 0$ for $j > s$, so there is no need to fix these Hilbert function values.

LEMMA 3.7. *For some $-2 \leq k \leq s$, fix a vector $\nu = (\nu_s, \nu_{s-1}, \dots, \nu_{k+1})$. Let*

$$X_{\mathbf{d}}^\nu = \{x \in X_{\mathbf{d}} \mid \alpha_j(\mathcal{M}_x) = \nu_j \text{ for all } j = k+1, \dots, s\}.$$

For any $i \geq 2$ and any $k \leq j \leq s$, the function $X_{\mathbf{d}}^\nu \rightarrow \mathbf{N}$ given by $x \mapsto \beta_{i,j}(\mathcal{M}_x)$ is upper semicontinuous.

Proof. Since Betti numbers are cone-stable, it suffices to prove that the function is upper semicontinuous on $X_{\mathbf{d},n}^\nu$. By our assumption on ν and Lemma 3.6, for $x \in X_{\mathbf{d},n}^\nu$, the module $\mathcal{M}_x$ has constant Hilbert function in degrees $k+1, \dots, s$. It follows that $\mathcal{M}_d$ is a finite-rank vector bundle for $d = k+1, \dots, s$.

By [Eis05, Theorem 7.8], the Betti number $\beta_{i,j}(\mathcal{M}_x)$ is the dimension of the homology of the complex

$$S_{i+1}W_n \otimes_{\mathbf{k}} (\mathcal{M}_x)_{j+i+1} \rightarrow S_i W_n \otimes_{\mathbf{k}} (\mathcal{M}_x)_{j+i} \rightarrow S_{i-1}W_n \otimes_{\mathbf{k}} (\mathcal{M}_x)_{j+i-1}.$$

Note that for $k \leq j \leq s$ and $i \geq 2$, all degrees in the above complex are in the range where $\mathcal{M}$ is a vector bundle. We also have $(\mathcal{M}_x)_d \cong (\mathcal{M}_d)_x$ for all degrees d . Thus, the locus of points x where $\beta_{i,j}(\mathcal{M}_x) \geq r$ is the locus where, after specializing to x , the homology of the following complex of finite-rank vector bundles has dimension at least r :

$$S_{i+1}W_n \otimes_{\mathbf{k}} \mathcal{M}_{j+i+1} \rightarrow S_i W_n \otimes_{\mathbf{k}} \mathcal{M}_{j+i} \rightarrow S_{i-1}W_n \otimes_{\mathbf{k}} \mathcal{M}_{j+i-1}.$$

Hence our desired result follows from standard upper semicontinuity properties of coherent sheaves [Har77, Example III.12.7.2]. $\square$

COROLLARY 3.8. *For any $-2 \leq k \leq s$, there are only finitely many distinct possibilities for $\alpha_k(\mathcal{M}_x)$ for $x \in X_{\mathbf{d}}$.*

Proof. Using Lemma 3.3, it suffices to understand the distinct possibilities for $x \in U_{\mathbf{d}'}$ for all $\mathbf{d}' \leq \mathbf{d}$. For any given $\mathbf{d}$, there are only finitely many possibilities for $\mathbf{d}'$, so without loss of generality, it is enough to consider $x \in U_{\mathbf{d}}$.

The case $k = s$ is immediate as $\alpha_s(\mathcal{M}_x) = \mathbf{b}_s$ by definition, and we proceed by descending induction on k . We assume that there are finitely many distinct possibilities for $\alpha_j(\mathcal{M}_x)$ for each $k+1 \leq j \leq s$. This implies that for $x \in U_{\mathbf{d}}$, there are only finitely many possibilities for the vector $\nu = (\alpha_{k+1}(\mathcal{M}_x), \dots, \alpha_s(\mathcal{M}_x))$. It thus suffices to prove the claim after restricting to one such ν .

For a fixed ν , we apply Lemma 3.7 to conclude, via **GL**-noetherianity, that there are only finitely many possibilities for each $\beta_{i,k}(\mathcal{M}_x)$ for each $i \geq 2$. It follows that, even taking the union over all ν , there are still only finitely many possibilities for each $\beta_{i,k}(\mathcal{M}_x)$ for each $i \geq 2$. We combine this with the following facts: $\beta_{0,k} = \mathbf{b}_k$ and $\beta_{1,k} = \mathbf{b}'_{k+1}$ (since $x \in U_{\mathbf{d}}$); $\alpha_k(\mathcal{M}_x) = \sum_{i=0}^{\infty} (-1)^i \beta_{i,k}(\mathcal{M}_x)$ by definition; and $\beta_{i,k}(\mathcal{M}_x) = 0$ for $i > s - k$ by Remark 3.5. We conclude that there are only finitely many possibilities for $\alpha_k(\mathcal{M}_x)$, as desired. $\square$

Combining Lemma 3.7 and Corollary 3.8 yields the following.

COROLLARY 3.9. *For any $-2 \leq k \leq s$, there are only finitely many distinct possibilities for $\beta_{2,k}(\mathcal{M}_x)$ for $x \in X_{\mathbf{d}}$.*

Proof. By Corollary 3.8, for $x \in X_{\mathbf{d}}$, there are only finitely many possibilities for $\alpha_j(\mathcal{M}_x)$ for each $-2 \leq j \leq s$. We may thus restrict to a locus in $X_{\mathbf{d}}$ where all of these values are constant. Then we can apply Lemma 3.7 to conclude that $\beta_{2,k}(\mathcal{M}_x)$ is upper semicontinuous, and by **GL**-noetherianity, this implies that there are only finitely many possible values. $\square$

The following corollary will also be a useful special case.

COROLLARY 3.10. *For $i \geq 2$, we let $Z_i \subseteq X_{\mathbf{d}}$ be the locus of x such that $\beta_{i,j}(\mathcal{M}_x) \neq 0$ for some $j > -i$. Then $Z_i \supseteq Z_{i+1}$ and each Z_i is Zariski closed.*

Remark. In reference to (3.5a), the condition for x to be in $Z_i \subseteq X_{\mathbf{d}}$ means that $\beta(\mathcal{M}_x)$ has a nonzero entry in column i and some row strictly above the zeroth row.

Proof. We have $Z_i \subseteq Z_{i-1}$ for degree reasons: if $x \notin Z_{i-1}$, then the $(i-1)$ st column of $\beta(\mathcal{M}_x)$ has no nonzero entries above the zeroth row, and the same must be true for the i th column, and thus $x \notin Z_i$.

The locus Z_2 is a finite union of Zariski-closed loci by Lemma 3.7, and hence it is Zariski closed. We will prove that Z_k is also a finite union of closed loci for any $k > 2$. Within Z_2 , there are only finitely many distinct possibilities for the Betti numbers $\beta_{2,k}(\mathcal{M}_x)$ with $-2 \leq k \leq s$. We restrict to one such possibility, which determines a new degree sequence, and we apply the same argument as we used for Z_2 . Iterating the argument yields the desired result for Z_i . $\square$

4. Proof of Theorem 1.3

Tate resolutions are stable under pushforward via linear embeddings $\mathbf{P}^n \rightarrow \mathbf{P}^{n+1}$. This is like a natural dual of the notion of cone-stability introduced in [ESS19b].

LEMMA 4.1. *Let $\mathcal{E}$ be a coherent sheaf on $\mathbf{P}^n$, and let $\iota: \mathbf{P}^n \rightarrow \mathbf{P}^{n+1}$ be a linear embedding. The Tate resolution $\mathbf{T}(\mathcal{E})$ is stable under pushforward by ι , namely: $\mathbf{T}(\mathcal{E}) \otimes_{\bigwedge W_n} \bigwedge W_{n+1}$ is isomorphic to $\mathbf{T}(\iota_* \mathcal{E})$.*

Proof. We consider $V_n = \langle x_0, \dots, x_n \rangle$ as a direct summand of $V_{n+1} = \langle x_0, \dots, x_{n+1} \rangle = V_n \oplus V'$. Since the Tate resolution is everywhere exact, and since minimal free resolutions over $\bigwedge W_{n+1}$ are unique up to isomorphism, it is enough to check the lemma for a single differential in the complex. Let $k \geq \text{reg}(\mathcal{E})$. Define the $\text{Sym}(V_n)$ -module $M = \bigoplus_{j \geq k} H^0(\mathbf{P}^n, \mathcal{E}(j))$ and the $\text{Sym}(V_{n+1})$ -module $M' = \bigoplus_{j \geq k} H^0(\mathbf{P}^{n+1}, \iota_* \mathcal{E}(j))$. By definition, $M_k \cong M'_k$ for all k , and x_{n+1} annihilates M' .

The differential $\rho^k: \mathbf{T}(\iota_* \mathcal{E})^k \rightarrow \mathbf{T}(\iota_* \mathcal{E})^{k+1}$ is defined by $M'_k \rightarrow M'_{k+1} \otimes W_{n+1}$, which is the adjoint of the map $M'_k \otimes V_{n+1} \rightarrow M'_{k+1}$. However, since x_{n+1} annihilates M' , we have

$$\begin{array}{ccc} M'_k \otimes V_{n+1} & \longrightarrow & M'_{k+1} \\ \downarrow & \nearrow & \\ M'_k \otimes V_n & & \end{array}$$

where the vertical arrow is induced by the projection $V_{n+1} \rightarrow V_n$. It follows that the image of $M'_k \rightarrow M'_{k+1} \otimes W_{n+1}$ lies in $M'_{k+1} \otimes W_n$. In particular, ρ^k only involves linear forms from $\bigwedge W_n$.

Now, the differential $\partial^k: \mathbf{T}(\mathcal{E})^k \rightarrow \mathbf{T}(\mathcal{E})^{k+1}$ is defined by $M_k \rightarrow M_{k+1} \otimes W_n$, which is the adjoint of the map $M_k \otimes V_n \rightarrow M_{k+1}$. Since we have natural isomorphisms $M_k \cong M'_k$ and $M_{k+1} \cong M'_{k+1}$, we can identify the map $M_k \rightarrow M_{k+1} \otimes W_n$ with $M'_k \rightarrow M'_{k+1} \otimes W_n$. This

identifies the matrix ϕ with the matrix ϕ' . By the exactness and uniqueness of Tate resolutions, this implies the desired statement. $\square$

Remark 4.2. Let $x \in X_{\mathbf{d}}^0 \cap X_{\mathbf{d},n}$ and $\mathcal{E}_x$ be the corresponding coherent sheaf on $\mathbf{P}^n$. Lemma 4.1 implies that if we then consider x as a point $\phi \in X_{\mathbf{d}}^0 \cap X_{\mathbf{d},n+1}$, the corresponding sheaf would simply be $\iota_*(\mathcal{E}_x)$, where $\iota: \mathbf{P}^n \rightarrow \mathbf{P}^{n+1}$. Since cohomology groups are invariant under closed immersions, it follows that the function which sends x to the cohomology table of $\mathcal{E}_x$ is well-defined on $X_{\mathbf{d}}^0$. We may henceforth refer to the cohomology table of $\mathcal{E}_x$ without reference to n .

LEMMA 4.3. *Let $\mathcal{E}$ be a coherent sheaf on $\mathbf{P}^n$. If there exists an $m < n$ such that the matrices in the Tate resolution $\mathbf{T}(\mathcal{E})$ only involve the variables $e_0, \dots, e_m$, then there exist a coherent sheaf $\mathcal{E}'$ on $\mathbf{P}^m$ and a linear embedding $\iota': \mathbf{P}^m \rightarrow \mathbf{P}^n$ such that $\iota'_*\mathcal{E}' \cong \mathcal{E}$.*

Proof. Let $V_m = \langle x_0, \dots, x_m \rangle$, considered as a split summand of V_n . Let $k \geq \text{reg}(\mathcal{E})$, and let $M = \bigoplus_{j \geq k} H^0(\mathbf{P}^n, \mathcal{E}(j))$. For each $j \geq k$, the map $M_j \rightarrow M_{j+1} \otimes V_n^*$ factors through $M_j \rightarrow M_{j+1} \otimes V_m^*$. Writing $V_n = V_m \oplus V''$, with $V'' = \langle x_{m+1}, \dots, x_n \rangle$ the vector space on $x_{m+1}, \dots, x_n$, this implies that for the multiplication map $M_j \otimes (V_m \oplus V'') \rightarrow M_{j+1}$, the map $M_j \otimes V'' \rightarrow M_{j+1}$ is zero. Since this holds for all $j \geq k$, this implies that M is annihilated by $x_{m+1}, \dots, x_n$. In particular, M can be viewed as a module over $\mathbf{k}[x_0, \dots, x_m]$, and the statement follows. $\square$

The following lemma is a special case of Theorem 1.3 and Corollary 1.5, where $\mathbf{b}_i$ and $\mathbf{b}'_i$ are zero for $i > 1$. We will eventually use noetherian induction to reduce to this case.

LEMMA 4.4. *Let $\mathbf{b} = (\mathbf{b}_0, 0, 0, \dots)$ and $\mathbf{b}' = (\mathbf{b}'_0, 0, 0, \dots)$ be column vectors. For any coherent sheaf $\mathcal{E}$ on $\mathbf{P}^n$ with $\gamma^0(\mathcal{E}) = \mathbf{b}$ and $\gamma^1(\mathcal{E}) = \mathbf{b}'$,*

- (a) $\mathcal{E}$ is 0-regular, and
- (b) if $n_0 = \mathbf{b}_0 \mathbf{b}'_0 - 1$, then there exist a coherent sheaf $\mathcal{E}'$ on $\mathbf{P}^{n_0}$ and a linear map $\iota: \mathbf{P}^{n_0} \rightarrow \mathbf{P}^n$ such that $\mathcal{E} \cong \iota_* \mathcal{E}'$.

Proof. The fact that $\mathcal{E}$ is 0-regular follows from Remark 3.5: the shape of the Betti table implies that $\mathcal{E}(j)$ does not have higher cohomology for $j \geq 0$. Choosing a Tate resolution $\mathbf{T}(\mathcal{E})$ for $\mathcal{E}$, we have $\mathbf{T}(\mathcal{E})^0 = E^{\mathbf{b}_0}$ and $\mathbf{T}(\mathcal{E})^1 = E(-1)^{\mathbf{b}'_0}$. The matrix $\phi: E^{\mathbf{b}_0} \rightarrow E(-1)^{\mathbf{b}'_0}$ involves at most $\mathbf{b}_0 \mathbf{b}'_1$ many distinct linear forms from E . In particular, if we set $n_0 = \mathbf{b}_0 \mathbf{b}'_0 - 1$, then after a change of coordinates, we can assume that ϕ only involves the variables $e_0, \dots, e_{n_0}$. Since $\mathbf{T}(\mathcal{E})$ is doubly exact, it follows that the entire Tate resolution $\mathbf{T}(\mathcal{E})$ can be defined over the subring in the variables $e_0, \dots, e_{n_0}$. The statement now follows from Lemma 4.3. $\square$

Proof of Corollary 1.5. For each $x \in X_{\mathbf{d}}^0$, we write $\mathcal{E}_x$ (as in Remark 4.2) for the corresponding coherent sheaf. Choose s such that $\mathbf{b}_i = 0$ and $\mathbf{b}'_i = 0$ for all $i > s$. The right side of the cohomology table of $\mathcal{E}_x$ will look like

$$\gamma(\mathcal{E}_x) = \begin{array}{c|cccccc} & & \gamma^0 & \gamma^1 & \gamma^2 & \dots & \gamma^k & \dots \\ \hline \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ s+1 & \dots & 0 & 0 & 0 & \dots & 0 & \dots \\ s & \dots & \mathbf{b}_s & \mathbf{b}'_s & \gamma_{s,2-s} & & \gamma_{s,k-s} & \dots \\ \vdots & \dots & \vdots & \vdots & \vdots & \ddots & \vdots & \dots \\ 1 & \dots & \mathbf{b}_1 & \mathbf{b}'_1 & \gamma_{1,1} & \dots & \gamma_{1,k-1} & \dots \\ 0 & \dots & \mathbf{b}_0 & \mathbf{b}'_0 & \gamma_{0,2} & \dots & \gamma_{0,k} & \dots \end{array}$$

In particular, for $k \geq 0$, the column γ^k can only have nonzero entries in rows $0, 1, \dots, s$.

We can rephrase this in terms of the Betti table of $\mathcal{M}_x$. Namely, since $\mathcal{M}_x$ has regularity zero by Lemma 2.3, the Betti table of $\mathcal{M}_x$ will look like

$$\beta(\mathcal{M}_x) = \begin{array}{c|cccccc} & 0 & 1 & 2 & \cdots & i & \cdots \\ \hline s & \mathbf{b}_s & \mathbf{b}'_s & \beta_{2,s-2} & \cdots & \beta_{i,s-i} & \\ \vdots & \vdots & \vdots & \vdots & \cdots & \vdots & \cdots \\ 1 & \mathbf{b}_1 & \mathbf{b}'_1 & \beta_{2,-1} & \cdots & \beta_{i,1-i} & \\ 0 & \mathbf{b}_0 & \mathbf{b}'_0 & \beta_{2,-2} & \cdots & \beta_{i,-i} & \end{array}$$

where $\gamma_{i,j} = \beta_{i+j,-j}$.

For $i \geq 2$, we define $Z_i \subseteq X_{\mathbf{d}}^0$ as in Corollary 3.10, but restricted to $X_{\mathbf{d}}^0$. Namely: some $x \in X_{\mathbf{d}}^0$ lies in Z_i if and only if $\beta_{i,j}(\mathcal{M}_x) \neq 0$ for some $j > -i$, which is in turn equivalent to the condition $\gamma_{k,i-k}(\mathcal{E}_x) \neq 0$ for some $k > 0$. By Corollary 3.10, the Z_i form a descending chain of closed loci in $X_{\mathbf{d}}^0$. These thus eventually stabilize by **GL**-noetherianity. Since every element of $X_{\mathbf{d}}^0$ has a cohomology table which is eventually entirely supported on the zeroth row (see Lemma 2.3), this must stabilize to the empty set. Thus, there is some k_0 such that each $\mathcal{E}$ has regularity at most k_0 . $\square$

Proof of Theorem 1.3. We have now shown that for every $x \in X_{\mathbf{d}}^0$, the corresponding sheaf $\mathcal{E}_x$ has a cohomology table that looks like

$$\gamma(\mathcal{E}_x) = \begin{array}{c|cccccccccc} & \cdots & \gamma^0 & \gamma^1 & \gamma^2 & \cdots & \gamma^{k_0-1} & \gamma^{k_0} & \gamma^{k_0+1} & \cdots \\ \hline s & \cdots & \mathbf{b}_s & \mathbf{b}'_s & \gamma_{s,2-s} & \cdots & \gamma_{s,k_0-1-s} & 0 & 0 & \cdots \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \cdots \\ 1 & \cdots & \mathbf{b}_1 & \mathbf{b}'_1 & \gamma_{1,1} & \cdots & \gamma_{1,k_0-2} & 0 & 0 & \cdots \\ 0 & \cdots & \mathbf{b}_0 & \mathbf{b}'_0 & \gamma_{0,2} & \cdots & \gamma_{0,k_0-1} & \gamma_{0,k_0} & \gamma_{0,k_0+1} & \cdots \end{array}$$

We apply Lemma 3.9 to conclude that there are only finitely many possibilities for $\gamma^2(\mathcal{E})$. For each such possibility, we again apply Lemma 3.9 to ensure that there are only finitely many possibilities for $\gamma^3(\mathcal{E})$. Iterating this argument, we eventually conclude that there are only finitely many possibilities for γ^{k_0} and γ^{k_0+1} .

Fix some such pair $(\gamma_{0,k_0}, \gamma_{0,k_0+1})$. Replacing $\mathcal{E}$ with $\mathcal{E}(k_0)$, we can apply Lemma 4.4 to see that $\mathcal{E}$ is defined on a subprojective space of dimension $\gamma_{0,k_0}\gamma_{0,k_0+1} - 1$. We now cycle over all of the finitely many possibilities for this pair $(\gamma_{0,k_0}, \gamma_{0,k_0+1})$. Let n_0 be the maximal dimension of any of the resulting subprojective spaces. $\square$

Proof of Corollary 1.4. By Theorem 1.3, it suffices to prove the corollary for a fixed n . Since $X_{\mathbf{d},n}$ is reduced and noetherian, this follows from Corollary 4.6 below. $\square$

LEMMA 4.5. *Let V be a reduced noetherian scheme, and let $\mathcal{F}$ be a coherent sheaf on $V \times \mathbf{P}^n$. For a point $x \in V$, we write $\mathcal{F}_x$ for the pullback to $\{x\} \times \mathbf{P}^n$.*

- (a) *There exist a nonempty open set $U \subseteq V$ and an integer j_0 such that for all $x \in U, j \geq j_0$, and all $i = 0, 1, \dots, n$, the dimension of $H^i(\mathcal{F}_x(j))$ does not depend on x .*
- (b) *Write ω for $\mathcal{O}_V \boxtimes \omega_{\mathbf{P}^n}$ on $V \times \mathbf{P}^n$. For any i , there is a nonempty open subset $U \subseteq V$ such that $\mathcal{E}xt^i(\mathcal{F}, \omega)_x \cong \mathcal{E}xt^i(\mathcal{F}_x, \omega_{\mathbf{P}^n})$ for all $x \in U$.*
- (c) *There exist a nonempty open set $U \subseteq V$ and integers j_- and j_+ such that $\gamma^j(\mathcal{F}_x)$ does not depend on x for all $x \in U$ and all $j \geq j_+$ or $j \leq j_-$.*
- (d) *There is a nonempty open set $U \subseteq V$ such that the entire cohomology table of $\mathcal{F}_x$ does not depend on x for all $x \in U$.*

Proof. Without loss of generality, we can assume throughout that V is integral.

(a) By generic flatness, we can assume that $\mathcal{F}$ is flat over an open set U . Since the regularity of a sheaf is semicontinuous in a flat family, we can moreover assume that $\mathcal{F}_x$ is j_0 -regular for all x in some open set U . It follows that the Hilbert polynomial of $\mathcal{F}_x$ is constant for $x \in U$. In addition, for each $j \geq j_0$ and $x \in U$, the dimension $\dim H^0(\mathcal{F}_x(j))$ equals the value of that constant Hilbert polynomial evaluated at j , and the higher cohomology groups vanish.

(b) Each of the Ext-sheaves is computed as the i th homology of a certain complex of coherent sheaves on $V \times \mathbf{P}^n$. By generic flatness, we can find an open set U where taking homology commutes with specialization to points $x \in U$, and the statement follows.

(c) Fix some i between 0 and n . We assume that the conclusion of part (b) holds and, by generic flatness, that $\mathcal{E}xt^i(\mathcal{F}, \omega)$ is flat over U . We may further find some j_i such that for all $x \in U$, the sheaf $\mathcal{E}xt^i(\mathcal{F}_x, \omega_{\mathbf{P}^n})(j)$ has no higher cohomology for $j \geq j_i$. We compute

$$H^0(\mathcal{E}xt^i(\mathcal{F}_x, \omega_{\mathbf{P}^n})(j)) = H^0(\mathcal{E}xt^i(\mathcal{F}_x(-j), \omega_{\mathbf{P}^n})) = \text{Ext}^i(\mathcal{F}_x(-j), \omega_{\mathbf{P}^n}) = H^{n-i}(\mathcal{F}_x(-j)),$$

where the first equality is [Har77, Proposition III.6.7], the second equality follows from the local-to-global spectral sequence for Ext, and the third equality is Serre duality.

By part (a), perhaps after increasing j_i or shrinking to a smaller open set, we can obtain that the dimension of $H^0(\mathcal{E}xt^i(\mathcal{F}_x, \omega_{\mathbf{P}^n})(j))$ does not depend on x for all $j \geq j_i$, and thus $\dim H^{n-i}(\mathcal{F}_x(-j))$ does not depend on x for all $j \geq j_i$. It follows that, in sufficiently negative degrees, there is an open set U where column $\gamma^j(\mathcal{F}_x)$ does not depend on x for all $x \in U$ and all $j \ll 0$. A similar statement holds for $j \gg 0$ by part (a).

(d) Let U be the open set from part (c), and let $\eta \in U$ its generic point. Since each individual entry of the cohomology table is semicontinuous, part (c) shows that there are a finite number of Zariski-open conditions which define the locus in U where the cohomology of $\mathcal{F}_x$ agrees with the cohomology table of $\mathcal{F}_\eta$. $\square$

COROLLARY 4.6. *Let V be a reduced noetherian scheme, and let $\mathcal{F}$ be a coherent sheaf on $V \times \mathbf{P}^n$. Then the fibers $\mathcal{F}_x$ with $x \in V$ attain only finitely many distinct cohomology tables.*

Proof. By Lemma 4.5(d), there is a Zariski-open set $U \subseteq V$ consisting of points with a constant cohomology table. The statement then follows by noetherian induction. $\square$

Remark 4.7. Corollary 1.4 is a finiteness result based on fixing *any* two consecutive columns of the cohomology table. By contrast, similar results about Betti tables, like [AH20, Theorem D], require one to fix the first two columns of the Betti table.

For Betti tables, this turns out to be necessary. For instance, if $S = \mathbf{k}[x, y]$ then for $I = (x^n, x^{n-1}y)$, the Betti table of $S(n-1)/I$ is

$$\begin{array}{ccc} -n & \begin{pmatrix} 1 & - & - \\ \vdots & \vdots & \vdots \\ 0 & - & 2 & 1 \end{pmatrix} \end{array}.$$

Thus, there are infinitely many different Betti table with $\beta_{1,1} = 2$ and $\beta_{2,2} = 1$ and all other entries in columns 1 and 2 equal to zero.

However, if we fix one additional piece of data—the degree in which the module is generated—then we can obtain a finiteness result. This is a consequence of Stillman’s conjecture (which bounds the tail of such a Betti table) and Boij–Söderberg theory (which bounds the rest of the entries).

REFERENCES

- AH20 T. Ananyan and M. Hochster, *Small subalgebras of polynomial rings and Stillman’s conjecture*, J. Amer. Math. Soc. **33** (2020), no. 1, 291–309; [doi:10.1090/jams/932](#).
- Bei78 A. A. Beilinson, *Coherent sheaves on $\mathbf{P}^n$ and problems in linear algebra*, Funct. Anal. Appl. **12** (1978), no. 3, 214–216; [doi:10.1007/BF01681436](#).
- BGG78 I. N. Bernshtein, I. M. Gel’fand, and S. I. Gel’fand, *Algebraic vector bundles on $\mathbf{P}^n$ and problems of linear algebra*, Funct. Anal. Appl. **12** (1978), no. 3, 212–214; [doi:10.1007/BF01681435](#).
- Dra19 J. Draisma, *Topological Noetherianity of polynomial functors*, J. Amer. Math. Soc. **32** (2019), no. 3, 691–707; [doi:10.1090/jams/923](#).
- EE17 D. Eisenbud and D. Erman, *Categorified duality in Boij–Söderberg theory and invariants of free complexes*, J. Eur. Math. Soc. **19** (2017), no. 9, 2657–2695; [doi:10.4171/JEMS/725](#).
- EFS03 D. Eisenbud, G. Fløystad, and F.-O. Schreyer, *Sheaf cohomology and free resolutions over exterior algebras*, Trans. Amer. Math. Soc. **355** (2003), no. 11, 4397–4426; [doi:10.1090/S0002-9947-03-03291-4](#).
- Eis05 D. Eisenbud, *The geometry of syzygies: a second course in commutative algebra and algebraic geometry*, Grad. Texts in Math., vol. 229 (Springer-Verlag, New York, 2005); [doi:10.1007/b137572](#).
- ES09 D. Eisenbud and F.-O. Schreyer, *Betti numbers of graded modules and cohomology of vector bundles*, J. Amer. Math. Soc. **22** (2009), no. 3, 859–888; [doi:10.1090/S0894-0347-08-00620-6](#).
- ESS18 D. Erman, S. V Sam, and A. Snowden, *Big polynomial rings with imperfect coefficient fields*, 2018, [arXiv:1806.04208](#).
- ESS19a ———, *Big polynomial rings and Stillman’s conjecture*, Invent. Math. **218** (2019), no. 2, 413–439; [doi:10.1007/s00222-019-00889-y](#).
- ESS19b ———, *Generalizations of Stillman’s conjecture via twisted commutative algebras*, Int. Math. Res. Not., published online on 19 August 2019, [doi:10.1093/imrn/rnz123](#), to appear in print.
- Har77 R. Hartshorne, *Algebraic geometry*, Grad. Texts in Math., vol. 52 (Springer-Verlag, New York – Heidelberg, 1977); [doi:10.1007/978-1-4757-3849-0](#).
- McC19 J. McCullough, *Stillman’s question for exterior algebras and Herzog’s conjecture on Betti numbers of syzygy modules*, J. Pure Appl. Algebra **223** (2019), no. 2, 634–640; [doi:10.1016/j.jpaa.2018.04.012](#).
- SE10 F. O. Schreyer and D. Eisenbud, *Betti numbers of syzygies and cohomology of coherent sheaves*, Proceedings of the ICM 2010, vol. 2 (Hindustan Book Agency, New Delhi, 2010), 586–602.

Daniel Erman derman@math.wisc.edu

Department of Mathematics, University of Wisconsin, 480 Lincoln Drive, Madison, WI 53706, USA

Steven V Sam ssam@ucsd.edu

Department of Mathematics, University of California, San Diego, 9500 Gilman Dr., La Jolla, CA 92093, USA

Andrew Snowden asnowden@umich.edu

Department of Mathematics, University of Michigan, 530 Church Street, Ann Arbor, MI 48109, USA