SMALL PROJECTIVE SPACES AND STILLMAN UNIFORMITY FOR SHEAVES

DANIEL ERMAN, STEVEN V SAM, AND ANDREW SNOWDEN

ABSTRACT. We prove an analogue 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, ES10]. 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 $\geq d$. Define $\beta_{i,j}(M) = \dim_{\mathbf{k}} \operatorname{Tor}_i^R(M, \mathbf{k})_j$, where the subscript here indicates the jth 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 non-zero entry, that is by $\max\{j - i \mid \beta_{i,j} \neq 0\}$; and the projective 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 $R \to S$ 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:

2010 *Mathematics Subject Classification.* 14F05, 13D02.

Key words and phrases. Stillman's Conjecture, sheaf cohomology.

DE was partially supported by NSF DMS-1601619. SS was partially supported by NSF DMS-1849173 and a Sloan Fellowship. AS was supported by NSF DMS-1453893.

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* 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 exists a graded polynomial algebra* S_0*, generated by at most* n *variables of degrees at most* d*, a flat* $\mathbf{k}$*-algebra homomorphism* $S_0 \to S$*, 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 [ESS19, Theorem 5.14]. The precise statement is:

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 the theorem 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, ES10, 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 $\gamma_{i,j}(\mathcal{E}) = \dim_{\mathbf{k}} \mathrm{H}^i(X, \mathcal{E}(j))$. These numbers are arranged into a table—the **cohomology table** of $\mathcal{E}$—as follows:

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

Unlike Betti tables, cohomology tables extend infinitely to the left and right. The number of non-zero rows is bounded by the dimension of the support of $\mathcal{E}$. We let $\gamma^k(\mathcal{E})$ denote the kth column of $\gamma(\mathcal{E})$, i.e., 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 $i\colon X \to Y$ 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 exists a linear embedding* $i\colon \mathbf{P}^{n_0} \to \mathbf{P}^n$ *with* $n_0 \le N$ *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 the theorem 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 the theorem then the Castelnuovo–Mumford regularity of* $\mathcal{E}$ *is at most* r*.*

In fact, we prove the above corollary first, and use it in our proof of Theorem 1.3.

Remark 1.6. In Theorem 1.3 we fixed the 0th and 1st 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 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 §2.

The initial data $\mathbf{b}, \mathbf{b}'$ determines the degrees of the entries of the 0th differential of the Tate resolution. We use this in §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 regularity of 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 $Z_r = \varnothing$ 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 §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, e.g. 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 [ESS19]? 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 §2 we review the necessary material, with a specific emphasis on the BGG correspondence and on properties of Tate resolutions. In §3, we construct the parameter spaces $X^0_{\mathbf{d}}$ 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, §4 contains the proofs of Theorem 1.3 and Corollary 1.4.

Acknowledgments. We thank Jason McCullough for useful conversations.

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 [Beĭ78, 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 $\partial^i \colon M_i \otimes_{\mathbf{k}} E(-i) \to M_{i+1} \otimes_{\mathbf{k}} E(-i-1)$ is determined by the image of the generators $M_i \to M_{i+1} \otimes W_n$ via the formula $\partial(m) = \sum_{i=0}^n (-1)^i x_i m \otimes e_i$. Alternately, one may view the map $M_i \to M_{i+1} \otimes W_n$ as the adjoint of the multiplication map $M_i \otimes V_n \to M_{i+1}$ 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 $\mathbf{P}^n = \mathrm{Proj}(\mathrm{Sym}(V_n))$. 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 $\mathbf{T}(\mathcal{E})$ [EFS03, Theorem 4.1]. To be more specific, if

$\mathbf{T}(\mathcal{E}) = [\cdots \to \mathbf{T}(\mathcal{E})^k \to \mathbf{T}(\mathcal{E})^{k+1} \to \cdots]$ is a Tate resolution for $\mathcal{E}$, then in cohomological degree k we have the free module

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

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

$$\gamma(\mathcal{O}_C) = \begin{array}{c|ccccccc} & & \gamma^{-2}(\mathcal{O}_C) & \gamma^{-1}(\mathcal{O}_C) & \gamma^0(\mathcal{O}_C) & \gamma^1(\mathcal{O}_C) & \gamma^2(\mathcal{O}_C) & \\ \hline 2 & \cdots & 0 & 0 & 0 & 0 & 0 & \cdots \\ 1 & \cdots & 9 & 6 & 3 & 1 & 0 & \cdots \\ 0 & \cdots & 0 & 0 & 1 & 3 & 6 & \cdots \end{array}.$$

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

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

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 $\operatorname{coker} \varphi$ solely in terms of the degrees of the entries of φ (i.e., 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 $\partial^k \colon \mathbf{T}(\mathcal{E})^k \to \mathbf{T}(\mathcal{E})^{k+1}$ be one of the differentials in a Tate resolution, and let $\varphi^\vee$ be a matrix representing the dual of ∂^k, where the duality functor is $\operatorname{Hom}_E(-, E)$. As the following lemma shows, it is always the case $\operatorname{coker} \varphi^\vee$ 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 $\operatorname{coker} \varphi^\vee$ 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 $\mathbf{T}(\mathcal{E})^k \to \mathbf{T}(\mathcal{E})^{k+1}$ and let $\varphi^\vee \colon (\mathbf{T}(\mathcal{E})^{k+1})^* \to (\mathbf{T}(\mathcal{E})^k)^*$ be a matrix representing the dual of ∂^k. The regularity of $\operatorname{coker}(\varphi^\vee)$ as a graded E-module is $-k$.*

In particular, if $\mathbf{G}$ is a minimal free resolution of $\operatorname{coker}(\varphi^\vee)$ 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 $\operatorname{coker}(\varphi^\vee)$ 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 $\nu(M)$ only depends on the pair (E, M) up to isomorphism. Let $\mathbf{k}[\varepsilon]/(\varepsilon^2)$ be the exterior algebra on 1 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 semi-continuous** if it is upper semi-continuous 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 $x \mapsto \nu(\mathcal{M}_x)$ is upper semi-continuous, i.e., 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 [ESS].

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 non-zero. 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 $F \to F'$ 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 will arise from minimal free resolutions, there will be no nonzero entries of degree 0, and thus we do not include a $\mathbf{d}_0$. Let $X_{\mathbf{d},n} = (\bigwedge^1 W_n)^{\oplus \mathbf{d}_1} \oplus (\bigwedge^2 W_n)^{\oplus \mathbf{d}_2} \oplus \cdots$ be the corresponding parameter space. We will identify points $x \in X_{\mathbf{d},n}$ with matrices $\varphi\colon F \to F'$ and sometimes abuse notation by writing $\varphi \in X_{\mathbf{d},n}$. We refer to such a matrix $\varphi\colon F \to F'$ 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 $F \to F'$ is thus determined by a block 4×4 matrix which consists of 6 entries of degree -1 (corresponding to submatrices $E^1 \to E(-1)^3$ and $E(1)^3 \to E^1$) and 9 entries of degree -2 (corresponding to $E(1)^3 \to E(-1)^3$). 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 $n \to \infty$. Let W_∞ be the direct limit of the spaces W_n, under the standard inclusions, and let $E_\infty = \bigwedge W_\infty$. The inclusions $W_n \to W_{n+1}$ also induce inclusions $X_{\mathbf{d},n} \to X_{\mathbf{d},n+1}$, 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 $\mathbf{GL} := \cup_n \mathbf{GL}_n$ on $X_{\mathbf{d}}$. We have that $X_{\mathbf{d}}$ is $\mathbf{GL}$-noetherian by [Dra19] and [ESS, Corollary 2.8]. As above, we identify points $\varphi \in X_{\mathbf{d}}$ 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 $\varphi \in X_{\mathbf{d}}$ 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 $\operatorname{coker} \varphi^\vee$ is zero. So we let

$$X^0_{\mathbf{d}} = \{\varphi \in X_{\mathbf{d}} \mid \operatorname{reg} \operatorname{coker} \varphi^\vee = 0\}$$

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

This construction is reversible as well. Assume that we are given some $x \in X^0_{\mathbf{d}}$ and choose n minimally so 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 $\operatorname{coker} \varphi^\vee$, then $\mathbf{F}_j$ is generated in degree $-j$ for $j \geq k$ for some k. Writing $\mathbf{F}^\vee_{\geq k}$ for the dual complex, we let $M = \mathbf{L}(\mathbf{F}^\vee_{\geq k})$ be the corresponding $\operatorname{Sym}(V_n)$-module, noting that, since $\mathbf{F}^\vee_{\geq k}$ is a linear complex, applying $\mathbf{L}$ yields a $\operatorname{Sym}(V_n)$-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 $\mathbf{F}^\vee_{\geq k}$ (see §2.2) and thus $\mathcal{E}$ has type $(\mathbf{b}, \mathbf{b}')$. In particular, by uniqueness of minimal free resolutions, there is a Tate resolution for $\mathcal{E}$ where φ is the kth 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 §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 associating a matrix φ to a point x, the matrix φ represents the differential $\partial^0 \colon \mathbf{T}^0(\mathcal{E}_x) \to \mathbf{T}^1(\mathcal{E}_x)$, but the regularity bound from Lemma 2.3 applies not to φ itself but $\varphi^\vee$. In fact, there is no a priori bound on the regularity of $\operatorname{coker} \varphi$. 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 $\Phi \colon \mathcal{F} \to \mathcal{F}'$ 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 $\varphi \in X_{\mathbf{d}}$ of type $(\mathbf{b}, \mathbf{b}')$. Set $\mathcal{M} = \operatorname{coker}(\Phi^\vee)$ to be the universal module. In particular, if x is a point of $X_{\mathbf{d}}$ which corresponds to the matrix $\varphi \colon F \to F'$, then the specialization $\Phi|_x$ is precisely φ. Thus $\mathcal{M}_x \cong \operatorname{coker}(\varphi^\vee)$. 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 $\varphi^\vee$ corresponding to x is a minimal presentation of $\mathcal{M}_x$.*

(a) *$U_{\mathbf{d}}$ 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 $\operatorname{coker} \psi^\vee \cong \mathcal{M}_x$.*

Proof. Let $\varphi \in X_{\mathbf{d}}$ be a matrix such that $\varphi^\vee$ 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 $\varphi^\vee$ is in the image of the other columns. Let c be the number of columns of $\varphi^\vee$; let $\varphi_1^\vee, \dots, \varphi_c^\vee$ be these columns and suppose that we can write $\varphi_c^\vee = \sum_{i=1}^{c-1} f_i \varphi_i^\vee$. Let $\mathbf{d}'$ be the degree sequence corresponding to the entries of the matrix $(\varphi_1^\vee \varphi_2^\vee \cdots \varphi_{c-1}^\vee)$ and let $\mathbf{d}''$ be the degree sequence corresponding to the exterior polynomials f_i. Then φ lies in the image of $X_{\mathbf{d}'} \times X_{\mathbf{d}''} \to X_{\mathbf{d}}$. In fact, this is equivalent to saying that $\varphi_c^\vee$ 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 $(\varphi_1^\vee \varphi_2^\vee \cdots \varphi_{c-1}^\vee)$. By [ESS, 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 $X_{\mathbf{d}',n} \times X_{\mathbf{d}'',n} \to X_{\mathbf{d},n}$ 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 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. □

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. For instance:

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 semi-continuity results for coherent sheaves [Har77, Example III.12.7.2]. □

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 semi-continuity 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, $\beta_{i,k}(N)$ 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):

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

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

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*

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

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

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

However $E(-j)_e = 0$ if $e > j$. If $j \geq e$ then $\dim E(-j)_e = \binom{n+1}{j-e}$. This yields:

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

It is well-known that Betti numbers are weakly semi-continuous, i.e., they are semi-continuous 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 $> k$, then each Betti number $\beta_{i,j}$ with $j > k$ is semi-continuous. 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 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 $-2 \le k \le s$ fix a vector $\nu = (\nu_s, \nu_{s-1}, \dots, \nu_{k+1})$. Let*

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

For any $i \ge 2$ and any $k \le j \le s$, the function $X^{\nu}_{\mathbf{d}} \to \mathbf{N}$ given by $x \mapsto \beta_{i,j}(\mathcal{M}_x)$ is upper semi-continuous.

Proof. Since Betti numbers are cone-stable, it suffices to prove that the function is upper semi-continuous on $X^{\nu}_{\mathbf{d},n}$. By our assumption on ν and Lemma 3.6, for $x \in X^{\nu}_{\mathbf{d},n}$, 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], $\beta_{i,j}(\mathcal{M}_x)$ is the dimension of the homology of the complex:

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

Note that, for $k \le j \le s$ and $i \ge 2$, 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 $\beta_{i,j}(\mathcal{M}_x) \ge r$ is the locus where, after specializing to x, the homology of the following complex of finite rank vector bundles has dimension $\ge r$:

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

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

Corollary 3.8. *For any $-2 \le k \le s$, 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 $\mathbf{d}' \le \mathbf{d}$. 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 $k+1 \le j \le s$. 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 $i \ge 2$. 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 $i \ge 2$. 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. □

Combining Lemma 3.7 and Corollary 3.8 yields:

Corollary 3.9. *For any $-2 \le k \le s$, 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 $-2 \le j \le s$. 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 semi-continuous and by **GL**-noetherianity, this implies that there are only finitely many possible values. □

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 $j > -i$.*[1] *Then $Z_i \supseteq Z_{i+1}$ and each Z_i is Zariski closed.*

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 ith 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. □

4. Proof of Theorem 1.3

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

Lemma 4.1. *Let $\mathcal{E}$ be a coherent sheaf on $\mathbf{P}^n$ and let $\iota\colon \mathbf{P}^n \to \mathbf{P}^{n+1}$ 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 $k \geq \operatorname{reg}(\mathcal{E})$. Define the $\operatorname{Sym}(V_n)$-module $M = \bigoplus_{j\geq k} \mathrm{H}^0(\mathbf{P}^n, \mathcal{E}(j))$ and the $\operatorname{Sym}(V_{n+1})$-module $M' = \bigoplus_{j\geq k} \mathrm{H}^0(\mathbf{P}^{n+1}, \iota_*\mathcal{E}(j))$. By definition, $M_k \cong M'_k$ for all k, and x_{n+1} annihilates M'.

The differential $\rho^k\colon \mathbf{T}(\iota_*\mathcal{E})^k \to \mathbf{T}(\iota_*\mathcal{E})^{k+1}$ is defined by $M'_k \to M'_{k+1} \otimes W_{n+1}$ which is the adjoint of the map $M'_k \otimes V_{n+1} \to M'_{k+1}$. 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 $V_{n+1} \to V_n$. It follows that the image of $M'_k \to M'_{k+1} \otimes W_{n+1}$ lies in $M'_{k+1} \otimes W_n$. In particular, ρ^k only involves linear forms from $\bigwedge W_n$.

Now, the differential $\partial^k\colon \mathbf{T}(\mathcal{E})^k \to \mathbf{T}(\mathcal{E})^{k+1}$ is defined by $M_k \to M_{k+1} \otimes W_n$, which is the adjoint of the map $M_k \otimes V_n \to M_{k+1}$. Since we have natural isomorphisms $M_k \cong M'_k$ and $M_{k+1} \cong M'_{k+1}$, we can identify the map $M_k \to M_{k+1} \otimes W_n$ with $M'_k \to M'_{k+1} \otimes W_n$. This identifies the matrix φ with the matrix φ'. By exactness and uniqueness of Tate resolutions, this implies the desired statement. □

[1] In reference to (3.5a), this means that $\beta(\mathcal{M}_x)$ has a nonzero entry in column i and some row strictly above the zeroth row.

Remark 4.2. Let $x \in X^0_{\mathbf{d}} \cap X_{\mathbf{d},n}$ 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 $\varphi \in X^0_{\mathbf{d}} \cap X_{\mathbf{d},n+1}$, the corresponding sheaf would simply be $\iota_*(\mathcal{E}_x)$ where $\iota\colon \mathbf{P}^n \to \mathbf{P}^{n+1}$. 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 $X^0_{\mathbf{d}}$. 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 $m < n$ such that the matrices in the Tate resolution $\mathbf{T}(\mathcal{E})$ only involve the variables $e_0, \dots, e_m$, then there exists a coherent sheaf $\mathcal{E}'$ on $\mathbf{P}^m$ and a linear embedding $\iota'\colon \mathbf{P}^m \to \mathbf{P}^n$ 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 $k \geq \operatorname{reg}(\mathcal{E})$ and let $M = \bigoplus_{j \geq k} \mathrm{H}^0(\mathbf{P}^n, \mathcal{E}(j))$. For each $j \geq k$, the map $M_j \to M_{j+1} \otimes V_n^*$ factors through $M_j \to M_{j+1} \otimes V_m^*$. 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 $M_j \otimes (V_m \oplus V'') \to M_{j+1}$, the map $M_j \otimes V'' \to M_{j+1}$ 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. □

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 $\gamma^1(\mathcal{E}) = \mathbf{b}'$ we have:*

(a) *$\mathcal{E}$ is 0-regular, and*
(b) *If $n_0 = \mathbf{b}_0\mathbf{b}'_0 - 1$ then there exists a coherent sheaf $\mathcal{E}'$ on $\mathbf{P}^{n_0}$ and a linear map $\iota\colon \mathbf{P}^{n_0} \to \mathbf{P}^n$ 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 $\varphi\colon E^{\mathbf{b}_0} \to E(-1)^{\mathbf{b}'_0}$ 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. □

Proof of Corollary 1.5. For each $x \in X^0_{\mathbf{d}}$ we write $\mathcal{E}_x$ (as in Remark 4.2) for the corresponding coherent sheaf. Choose s so 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:

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

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

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

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

For $i \geq 2$ we define $Z_i \subseteq X^0_{\mathbf{d}}$ as in Corollary 3.10, but restricted to $X^0_{\mathbf{d}}$. Namely: some $x \in X^0_{\mathbf{d}}$ 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 $X^0_{\mathbf{d}}$. These thus eventually stabilize by **GL**-noetherianity. Since every element of $X^0_{\mathbf{d}}$ 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. □

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

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

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 $\mathcal{E}$ by $\mathcal{E}(k_0)$, 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. □

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. □

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 is 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 $i = 0, 1, \ldots, n$, the dimension of $\mathrm{H}^i(\mathcal{F}_x(j))$ 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 is a nonempty open set $U \subseteq V$ and integers j_- and j_+ where $\gamma^j(\mathcal{F}_x)$ 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 $U \subseteq V$ where 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 $x \in U$, $\dim \mathrm{H}^0(\mathcal{F}_x(j))$ 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 ith 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 (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 $x \in U$, $\mathcal{E}xt^i(\mathcal{F}_x, \omega_{\mathbf{P}^n})(j)$ has no higher cohomology for $j \geq j_i$. We compute:

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

Where the first equality is [Har77, Prop 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 $\mathrm{H}^0(\mathcal{E}xt^i(\mathcal{F}_x, \omega_{\mathbf{P}^n})(j))$ does not depend on x for all $j \geq j_i$, and thus $\dim \mathrm{H}^{n-i}(\mathcal{F}_x(-j))$ 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 $\mathcal{F}_\eta$. □

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. □

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 $S = \mathbf{k}[x, y]$ and let $I = (x^n, x^{n-1}y)$, then the Betti table of $S(n-1)/I$ is:

$$\begin{matrix} -n \\ \vdots \\ 0 \end{matrix} \begin{pmatrix} 1 & - & - \\ \vdots & \vdots & \vdots \\ - & 2 & 1 \end{pmatrix}$$

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] Tigran Ananyan and Melvin Hochster, *Small subalgebras of polynomial rings and Stillman's Conjecture*, J. Amer. Math. Soc. **33** (2020), no. 1, 291–309. ↑1, 2, 14

[Beĭ78] A. A. Beĭlinson, *Coherent sheaves on* $\mathbf{P}^n$ *and problems in linear algebra*, Funktsional. Anal. i Prilozhen. **12** (1978), no. 3, 68–69 (Russian). ↑4

[BGG78] I. N. Bernšteĭn, I. M. Gelfand, and S. I. Gelfand, *Algebraic vector bundles on* $\mathbf{P}^n$ *and problems of linear algebra*, Funktsional. Anal. i Prilozhen. **12** (1978), no. 3, 66–67 (Russian). ↑4

[Dra19] Jan Draisma, *Topological noetherianity for polynomial functors*, J. Amer. Math. Soc. **32** (2019), no. 3, 691–707. arXiv:1705.01419v4. ↑3, 7

[Eis05] David Eisenbud, *The geometry of syzygies*, Graduate Texts in Mathematics, vol. 229, Springer-Verlag, New York, 2005. A second course in commutative algebra and algebraic geometry. ↑10

[EE17] David Eisenbud and Daniel Erman, *Categorified duality in Boij-Söderberg theory and invariants of free complexes*, J. Eur. Math. Soc. (JEMS) **19** (2017), no. 9, 2657–2695. ↑2

[EFS03] David Eisenbud, Gunnar Fløystad, and Frank-Olaf Schreyer, *Sheaf cohomology and free resolutions over exterior algebras*, Trans. Amer. Math. Soc. **355** (2003), no. 11, 4397–4426. ↑3, 4, 5, 7

[ES10] David Eisenbud and Frank-Olaf Schreyer, *Betti numbers of syzygies and cohomology of coherent sheaves*, Proceedings of the International Congress of Mathematicians. Volume II, Hindustan Book Agency, New Delhi, 2010, pp. 586–602. ↑1, 2

[ES09] David Eisenbud and Frank-Olaf Schreyer, *Betti numbers of graded modules and cohomology of vector bundles*, J. Amer. Math. Soc. **22** (2009), no. 3, 859–888. ↑1, 2

[ESS] Daniel Erman, Steven V Sam, and Andrew Snowden, *Generalizations of Stillman's Conjecture via twisted commutative algebras*, Int. Math. Res. Not. IMRN (to appear). arXiv:1804.09807v1. ↑6, 7, 8, 11

[ESS19] Daniel Erman, Steven V. Sam, and Andrew Snowden, *Big polynomial rings and Stillman's conjecture*, Invent. Math. **218** (2019), no. 2, 413–439. ↑2, 4

[ESS] Daniel Erman, Steven V Sam, and Andrew Snowden, *Big polynomial rings with imperfect coefficient fields*, Michigan Math. J. (to appear). arXiv:1806.04208v1. ↑2

[Har77] Robin Hartshorne, *Algebraic geometry*, Springer-Verlag, New York-Heidelberg, 1977. Graduate Texts in Mathematics, No. 52. ↑8, 10, 14

[McC19] Jason 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. arXiv:1307.8162v2. ↑2, 3, 5

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

Email address: derman@math.wisc.edu

URL: http://math.wisc.edu/~derman/

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

Email address: ssam@ucsd.edu

URL: http://math.ucsd.edu/~ssam/

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

Email address: asnowden@umich.edu

URL: http://www-personal.umich.edu/~asnowden/