

FINITE GRÖBNER BASES IN INFINITE DIMENSIONAL POLYNOMIAL RINGS AND APPLICATIONS

CHRISTOPHER J. HILLAR AND SETH SULLIVANT

ABSTRACT. We introduce the theory of monoidal Gröbner bases, a concept which generalizes the familiar notion in a polynomial ring and allows for a description of Gröbner bases of ideals that are stable under the action of a monoid. The main motivation for developing this theory is to prove finiteness results in commutative algebra and applications. A basic theorem of this type is that ideals in infinitely many indeterminates stable under the action of the symmetric group are finitely generated up to symmetry. Using this machinery, we give new streamlined proofs of some classical finiteness theorems in algebraic statistics as well as a proof of the independent set conjecture of Hoşten and the second author.

1. INTRODUCTION

In commutative algebra and its applications, one is frequently presented with a family of ideals in increasingly larger polynomial rings, and often it is observed that, up to some natural symmetry of the ideals, there exists a finite set of polynomials generating all of them. Such situations arise in universal algebra and group theory [4, 10], algebraic statistics [1, 19, 11, 14, 9, 3], algebraic problems in chemistry [18, 2, 9], and in classical results from combinatorial commutative algebra (for instance, that the $k \times k$ minors of a generic matrix form a Gröbner basis for the ideal they generate [20]). The particular form of one of these finiteness results typically depends on the specifics of the family of ideals. However, one wonders if there is a general principle at work that can explain a large portion of these phenomena.

We propose a general framework for proving finiteness theorems in rings with a monoid action. In this setting, a finiteness theorem takes one of two forms: (1) that a certain module over a noncommutative ring is Noetherian or (2) that a chain of ideals involving a monoidal filtration stabilizes. Although the precise formulation of our theory requires the setup found in Section 2, a typical result of the first type has the following flavor:

Theorem 1.1. *The polynomial ring $\mathbb{K}[X_{[r] \times \mathbb{P}}]$ is a Noetherian $\mathbb{K}[X_{[r] \times \mathbb{P}}] * \mathfrak{S}_{\mathbb{P}}$ -module.*

Here, $\mathbb{K}[X_{[r] \times \mathbb{P}}]$ is a polynomial ring over a field \mathbb{K} in the indeterminates $x_{i,j}$ with $i \in [r] := \{1, 2, \dots, r\}$ and $j \in \mathbb{P} := \{1, 2, 3, \dots\}$, the set of positive integers. Also, $\mathfrak{S}_{\mathbb{P}}$ is the set of permutations of \mathbb{P} , acting on $\mathbb{K}[X_{[r] \times \mathbb{P}}]$ by way of $\sigma \cdot x_{i,j} = x_{i,\sigma(j)}$, and the ring $\mathbb{K}[X_{[r] \times \mathbb{P}}] * \mathfrak{S}_{\mathbb{P}}$

Date: September 10, 2011.

Hillar was partially supported by an NSA Young Investigator Grant and an NSF All-Institutes Postdoctoral Fellowship administered by the Mathematical Sciences Research Institute through its core grant DMS-0441170. Sullivant is partially supported by NSF grant DMS-0840795. Part of this research was carried out during visits to SAMSI.

is the skew-monoid ring associated to $\mathbb{K}[X_{[r] \times \mathbb{P}}]$ and $\mathfrak{S}_{\mathbb{P}}$ (see Section 2 for more details). Stated simply, Theorem 1.1 says that every ideal in $\mathbb{K}[X_{[r] \times \mathbb{P}}]$ that is stable under the action of $\mathfrak{S}_{\mathbb{P}}$ has a finite generating set up to $\mathfrak{S}_{\mathbb{P}}$ symmetry.

A version of Theorem 1.1 was first proved by Cohen [4] in an application to the theory of free metabelian groups, and then rediscovered much later in the study of some polynomial finiteness questions inspired by chemistry [2] (see also [15] for another recent proof). Here, we study its application to algebraic statistics, and in particular its uniform treatment of some classical results in that field [14, 19]. Recent work by Draisma on finiteness problems for the factor analysis model [9] also depends on Theorem 1.1.

To prove Theorem 1.1 and similar results, we shall develop a suitable theory of Gröbner bases for certain modules over (noncommutative) rings. Section 2 contains this general theory of monoidal Gröbner bases and is the technical heart of the paper. In this framework, we have a monoid P of endomorphisms acting on a semigroup ring $\mathbb{K}[Q]$ (over a field \mathbb{K}), and a partial order (called the *P -divisibility order*) on the monomials of $\mathbb{K}[Q]$ that respects this action. Theorem 2.12, the main result in Section 2, is then the statement that finite Gröbner bases exist with respect to the monoid P if and only if P -divisibility is a well-partial-ordering. In many cases of interest (such as in our applications to algebraic statistics), this order condition is straightforward to check, leading directly to finite generation of ideals up to P -action. For instance, in the particular case of Theorem 1.1, the condition reduces to a classical lemma of Higman [13] in the order theory of words. Not surprisingly, all known proofs of Theorem 1.1 use Higman's lemma in an essential way.

We also introduce in Section 2 the concept of a filtration for a chain of ideals subject to the action of the monoid P (Definition 2.15). This notion allows us to pass from ideals in finitely many variables to ideals in infinitely many variables, and it can be used to formulate and prove finiteness theorems for P -invariant chains of ideals. Our main result in this regard is Theorem 2.19; it says that a P -invariant chain stabilizes with respect to a filtration (also) when P -divisibility is a well-partial-order.

Section 3 is concerned with the major implications of the theory contained in Section 2 and, in particular, a proof of Theorem 1.1. Beyond this result, we also provide a strategy using quotient modules for proving finite generation theorems for special ideals in rings (such as $\mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}]$) that are not Noetherian modules over skew-group rings (such as $\mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}] * (\mathfrak{S}_{\mathbb{P}} \times \mathfrak{S}_{\mathbb{P}})$). Section 4 contains our application of these ideas to finiteness theorems for Markov bases in algebraic statistics, including new proofs of the main results in [14, 19] as well as a proof of the independent set conjecture [14, Conj. 4.6]. The latter result, stated as Theorem 4.7 below and proved using filtrations, gives a finiteness property for Markov bases in models that have independent vertex sets.

Finally, Section 5 is devoted to a discussion of questions and problems left unresolved by this paper. In particular, the computational consequences of our work remain open.

2. MONOIDAL GRÖBNER BASES

In this section we develop our most important basic tools: finiteness theorems for invariant ideals of monoidal rings. These ideas generalize those of Aschenbrenner and the first author [2], and the proofs use similar ideas. The importance of our generalization comes both from

its usefulness, which will be illustrated throughout the paper, and from our distillation and simplification of the main techniques from [2], which might be of independent interest.

The main results of this section are Theorem 2.12 and Theorem 2.19. Theorem 2.12 gives a finiteness criterion for monoidal Gröbner bases which we will combine with Higman’s lemma in Section 3 to prove Theorem 1.1 from the introduction. Our other main result, Theorem 2.19, gives the same criterion for chains of ideals to stabilize, and we will use it to prove the independent set conjecture [14] in algebraic statistics (Theorem 4.7 below).

We begin with an abstract setting. Let \mathbb{K} be a field, let Q be a (possibly noncommutative) semigroup with identity (also called a *monoid*), and let $\mathbb{K}[Q]$ be the semigroup ring associated to Q (over \mathbb{K}). We call the elements of Q the *monomials* of $\mathbb{K}[Q]$. Let P be a monoid of \mathbb{K} -algebra endomorphisms of $\mathbb{K}[Q]$ (with multiplication in P given by composition).

Associated to $\mathbb{K}[Q]$ and P is the *skew-monoid ring* $\mathbb{K}[Q] * P$, which is formally the set of all linear combinations,

$$\mathbb{K}[Q] * P = \left\{ \sum_{i=1}^k c_i q_i p_i \ : \ c_i \in \mathbb{K}, q_i \in Q, p_i \in P \right\}.$$

Multiplication of monomials in the ring $\mathbb{K}[Q] * P$ is given by

$$q_1 p_1 \cdot q_2 p_2 = q_1 (p_1 q_2) (p_1 p_2),$$

and extended by distributivity to the whole ring. Note that $p_1 q_2$ in this expression denotes the result of applying the endomorphism p_1 to q_2 which is in $\mathbb{K}[Q]$ but is not necessarily a monomial. The natural (left) action of the skew-monoid ring on $\mathbb{K}[Q]$ makes $\mathbb{K}[Q]$ into a (left) module over $\mathbb{K}[Q] * P$ as one can readily verify.¹

We say that a (left) ideal $I \subseteq \mathbb{K}[Q]$ is *P-invariant* if

$$PI := \{pn : p \in P, n \in I\} = I.$$

Stated another way, a *P*-invariant ideal is simply a $\mathbb{K}[Q] * P$ -submodule of $\mathbb{K}[Q]$. We want to provide a general setting for defining what it means for a *P*-invariant ideal I of $\mathbb{K}[Q]$ to have a *P*-Gröbner basis. Of specific interest for applications is when I has a finite *P*-Gröbner basis, and our main contribution is a sufficient condition on P and Q under which this happens (see Theorem 2.12). The examples found in the next section will illustrate the usefulness of our general framework.

Remark 2.1. In many of our applications, Q will be a subsemigroup of the semigroup of natural number sequences with finite support (so that $\mathbb{K}[Q]$ is a subring of a polynomial ring), and P will be defined using maps on the indices of the indeterminates in that polynomial ring. When $P = \{\mathbf{1}\}$ consists of only the identity and $\mathbb{K}[Q]$ is a polynomial ring in a finite number of variables, we recover the classical formulation of Gröbner bases (see e.g. [6, Ch. 2]).

If we have a total ordering \preceq of Q , we can speak of the *initial monomial* or *leading monomial* $q = \text{in}_{\prec}(f)$ of any nonzero $f \in \mathbb{K}[Q]$, which is the largest element $q \in Q$ with respect to \preceq appearing with nonzero coefficient in f . For notational convenience, we set

¹We must use the skew-monoid ring $\mathbb{K}[Q] * P$ instead of the *monoid ring* $\mathbb{K}[Q][P]$ to ensure that $\mathbb{K}[Q]$ is a module. The authors of [2] made such a mistake although none of the results there need to be modified except to make this adjustment (the ring structure of $\mathbb{K}[Q] * P$ was not used in their proofs).

$\text{in}_<(f) = 0$ whenever $f = 0$, and also $0 < q$ for all $q \in Q$. We are interested in those orderings which are naturally compatible with the linear action of $\mathbb{K}[Q] * P$.

Definition 2.2 (*P-orders*). A well-ordering \preceq of Q is called a *P-order* on $\mathbb{K}[Q]$ if for all $q \in Q$, $p \in P$, and $f \in \mathbb{K}[Q]$, we have

$$\text{in}_<(qp \cdot f) = \text{in}_<(qp \cdot \text{in}_<(f)).$$

In the next section, we shall provide examples of *P-orders*. The most important example of a *P-order* for us will be the *shift order* on monomials (see Theorem 3.1).

Some basic facts about *P-orders* are collected in the following lemma. Note that when $P = \{\mathbf{1}\}$, a *P-order* is simply a *term order* on monomials. For a useful characterization of *P-orders*, see Proposition 2.4 below.

Lemma 2.3. *Suppose that \preceq is a P-order on $\mathbb{K}[Q]$. Then the following hold:*

- (1) *For all $q \in Q$, $p \in P$, and $q_1, q_2 \in Q$, we have $q_1 < q_2 \implies \text{in}_<(qpq_1) \preceq \text{in}_<(qpq_2)$.*
- (2) *If $\text{in}_<(qpf) = \text{in}_<(qpg)$ for some $q \in Q$, $p \in P$ and $f, g \in \mathbb{K}[Q]$, then either $\text{in}_<(f) = \text{in}_<(g)$ or $qpf = qpg = 0$.*
- (3) *Q is left-cancellative: for all $q, q_1, q_2 \in Q$, we have $qq_1 = qq_2 \implies q_1 = q_2$.*
- (4) *$q_2 \preceq q_1q_2$ for all $q_1, q_2 \in Q$ (in particular, $\mathbf{1}$ is the smallest monomial).*
- (5) *All endomorphisms in P are injective.*
- (6) *For all $q \in Q$ and $p \in P$, we have $q \preceq \text{in}_<(pq)$.*

Proof. (1): If $\text{in}_<(qpq_1) \neq \text{in}_<(qpq_2)$, then

$$\max\{\text{in}_<(qpq_1), \text{in}_<(qpq_2)\} = \text{in}_<(qpq_1 + qpq_2) = \text{in}_<(qp \cdot \text{in}_<(q_1 + q_2)) = \text{in}_<(qpq_2),$$

and the claim follows.

(2): If $\text{in}_<(qpg) = 0$, then $qpg = 0$, so assume that $\text{in}_<(qpf) = \text{in}_<(qpg) \neq 0$. If $\text{in}_<(f) < \text{in}_<(g)$, there exists $c \in \mathbb{K}$ such that the leading terms of qpg and $cqpf$ are the same. This implies

$$\text{in}_<(qp \cdot \text{in}_<(g)) = \text{in}_<(qp \cdot \text{in}_<(g - cf)) = \text{in}_<(qpg - cqpf) < \text{in}_<(qpg) = \text{in}_<(qp \cdot \text{in}_<(g)),$$

which is a contradiction. The first equality follows since $\text{in}_<(f) < \text{in}_<(g)$, the second since \preceq is a *P-order*, the middle inequality since the leading terms of qpg and $cqpf$ are the same, and the final equality follows again since \preceq is a *P-order*. Switching the roles of f and g , we therefore have $\text{in}_<(f) = \text{in}_<(g)$.

(3): Follows directly from (2) with $p = \mathbf{1}$, $f = q_1$, and $g = q_2$.

(4): Suppose that $q_1q_2 \preceq q_2$ for some $q_1, q_2 \in Q$. Since \preceq is a well-order, the infinite decreasing sequence obtained by using (1) repeatedly:

$$\cdots \preceq q_1^3q_2 \preceq q_1^2q_2 \preceq q_1q_2 \preceq q_2,$$

must terminate; in this case, we have $q_1^{k+1}q_2 = q_1^kq_2$ for some $k \in \mathbb{N}$. It follows that $q_1q_2 = q_2$ from (3), which proves (4).

(5): Let $p \in P$ and let $0 \neq f \in \mathbb{K}[Q]$. From (1) and (4) and the fact that p is a ring homomorphism, it follows that $1 = \text{in}_<(p \cdot \mathbf{1}) \preceq \text{in}_<(p \cdot \text{in}_<(f)) = \text{in}_<(pf)$. Thus, pf is nonzero for all $f \neq 0$, so p is injective.

(6): Finally, suppose that $\text{in}_\prec(pq) \preceq q$ for some $q \in Q$ and $p \in P$. This gives us an infinite decreasing sequence,

$$\cdots \preceq \text{in}_\prec(p^3q) \preceq \text{in}_\prec(p^2q) \preceq \text{in}_\prec(pq) \preceq q.$$

Since \preceq is a well-ordering, we must have $\text{in}_\prec(p^{k+1}q) = \text{in}_\prec(p^kq)$ for some $k \in \mathbb{N}$. Using (2) and (5) in conjunction, it follows that $\text{in}_\prec(pq) = q$, thereby proving (6). \square

It turns out that properties (1) and (2) in Lemma 2.3 characterize when P -orders exist (the others follow from these). As the following proposition demonstrates, we may further reduce the number of axioms to one. This will be useful in proving that certain well-orderings on Q are P -orders.

Proposition 2.4 (Characterization of P -orders). *Let Q be a monoid and let P be a monoid of \mathbb{K} -algebra endomorphisms of $\mathbb{K}[Q]$. Then a well-ordering \preceq of Q is a P -order if and only if for all $q \in Q$, $p \in P$, and $q_1, q_2 \in Q$, we have*

$$q_1 \prec q_2 \implies \text{in}_\prec(qpq_1) \prec \text{in}_\prec(qpq_2).$$

Proof. Suppose first that \preceq is a P -order. By Lemma 2.3 part (1) we know that $q_1 \prec q_2$ implies that $\text{in}_\prec(qpq_1) \preceq \text{in}_\prec(qpq_2)$. If $\text{in}_\prec(qpq_1) = \text{in}_\prec(qpq_2)$ for some $q \in Q$, $p \in P$, and $q_1, q_2 \in Q$, then Lemma 2.3 part (2) implies that $q_1 = \text{in}_\prec(q_1) = \text{in}_\prec(q_2) = q_2$ or $qpq_1 = qpq_2 = 0$, and part (5) implies that the second option is not possible. This proves the only-if direction.

Conversely, suppose that \preceq is a well-ordering of Q satisfying the hypothesis of the proposition. Let $q \in Q$, $p \in P$, and $0 \neq f \in \mathbb{K}[Q]$; we shall verify that $\text{in}_\prec(qpf) = \text{in}_\prec(qp \cdot \text{in}_\prec(f))$. Order the monomials $q_1 \prec \cdots \prec q_k$ appearing in f with nonzero coefficient. By assumption, we have $\text{in}_\prec(qpq_i) \prec \text{in}_\prec(qpq_{i+1})$ for all i . It follows that $\text{in}_\prec(qpf) = \text{in}_\prec(qp \cdot \text{in}_\prec(f))$ as desired. \square

Having a P -order is quite restrictive as the following example demonstrates.

Example 2.5 (Semigroup ring without a P -order). Let $\mathbb{K}[Q] = \mathbb{K}[X_{\mathbb{P}}]$ be the polynomial ring in infinitely many variables $X_{\mathbb{P}} = \{x_i : i \in \mathbb{P}\}$. Also, let $P = \mathfrak{S}_{\mathbb{P}}$ be the permutations of the positive integers \mathbb{P} , and let $\mathfrak{S}_{\mathbb{P}}$ act on $\mathbb{K}[X_{\mathbb{P}}]$ by permuting indices. Then there is no $\mathfrak{S}_{\mathbb{P}}$ -order on $\mathbb{K}[X_{\mathbb{P}}]$. To see this, let $g = x_1 + x_2$, and suppose (without loss of generality) that a P -order makes $\text{in}_\prec(g) = x_1$. Then if $p = (12)$, we have $\text{in}_\prec(p \cdot g) = \text{in}_\prec(g) = x_1$, while $\text{in}_\prec(p \cdot \text{in}_\prec(g)) = \text{in}_\prec(p \cdot x_1) = x_2$.

More generally, if $R = \mathbb{K}[Q] * P$ where P is a nontrivial group acting by permutations on Q , then there cannot exist a P -order on $\mathbb{K}[Q]$. This will necessitate our study of special classes of monoids P . \square

Before formulating a theory of Gröbner bases in this setting, we shall also need a relation (refining monomial divisibility) that is compatible with the canceling of leading monomials.

Definition 2.6 (The P -divisibility relation). Given monomials $q_1, q_2 \in Q$, we say that $q_1 \mid_P q_2$ if there exists $p \in P$ and $q \in Q$ such that $q_2 = q \cdot \text{in}_\prec(pq_1)$. Such a p is called a *witness* for the relation $q_1 \mid_P q_2$.

Proposition 2.7. *If \preceq is a P -order on Q , then P -divisibility \mid_P is a partial order on Q that is a coarsening of \preceq (i.e., $q_1 \mid_P q_2 \implies q_1 \preceq q_2$).*

Proof. First of all, it is clear that $|_P$ is reflexive. To prove transitivity, suppose that $q_2 = m_1 \cdot \text{in}_{\prec}(p_1 q_1)$ and $q_3 = m_2 \cdot \text{in}_{\prec}(p_2 q_2)$ for monomials $m_1, m_2 \in Q$ and $p_1, p_2 \in P$. Using the fact that p_2 is a ring homomorphism and (repeatedly) the defining property of P -orders, we have,

$$\begin{aligned} q_3 &= m_2 \cdot \text{in}_{\prec}(p_2 m_1 \cdot (p_2 \cdot \text{in}_{\prec}(p_1 q_1))) \\ &= m_2 \cdot \text{in}_{\prec}(p_2 m_1 \cdot \text{in}_{\prec}(p_2 \cdot \text{in}_{\prec}(p_1 q_1))) \\ &= m_2 \cdot \text{in}_{\prec}(p_2 m_1 \cdot \text{in}_{\prec}(p_2 p_1 q_1)). \end{aligned}$$

Since $\text{in}_{\prec}(p_2 m_1 \cdot \text{in}_{\prec}(p_2 p_1 q_1)) \neq 0$, it must be of the form $q \cdot \text{in}_{\prec}(p_2 p_1 q_1)$ for some $q \in Q$. It follows that $q_1 |_P q_3$ with witness $p = p_2 p_1$.

Finally, to prove antisymmetry, it is enough to verify that P -divisibility is a coarsening of \preceq . If $q_1 |_P q_2$, then for some $p \in P$ and $q \in Q$, we have $q_2 = q \cdot \text{in}_{\prec}(p q_1)$. Thus, by properties (4) and (6) in Lemma 2.3, we have $q_1 \preceq \text{in}_{\prec}(p q_1) \preceq q \cdot \text{in}_{\prec}(p q_1) = q_2$ as desired. \square

If \preceq is a P -order, then we may compute the *initial final segment* with respect to the P -divisibility partial order of any subset $G \subseteq \mathbb{K}[Q]$:

$$\text{in}_{\prec}(G) := \{q : \text{in}_{\prec}(g) |_P q \text{ for some } g \in G \setminus \{0\}\}.$$

It is clear that the set $\text{in}_{\prec}(G)$ contains all the initial monomials of G . Moreover, when $I \subseteq \mathbb{K}[Q]$ is a P -invariant ideal, it is straightforward to check that it contains no other ones:

$$\text{in}_{\prec}(I) = \{\text{in}_{\prec}(f) : f \in I \setminus \{0\}\}.$$

Remark 2.8. The schizophrenic terminology *initial final segment* comes from the combination of two mathematical traditions. From order theory, we have an upward closed subset of a partially ordered set, which is a final segment. On the other hand, we have constructed this set by taking initial or leading terms of polynomials.

Note that the initial final segment is not an ideal (or initial segment) in the sense of order theory (as it is not closed downward). Furthermore, it cannot, in general, be made into a monomial ideal of $\mathbb{K}[Q]$, as is typically done in commutative algebra, because P does not necessarily act by maps that send Q to itself.

We now arrive at our definition of Gröbner bases for invariant ideals with respect to a given P -order. We remark that a similar definition appears in [3], where they are given the name *equivariant Gröbner bases*, and [10] contains related work in the noncommutative case (but without the assumption that the term order is compatible with the monoid actions).

Definition 2.9. A set $G \subseteq I \subseteq \mathbb{K}[Q]$ is a P -Gröbner basis for a P -invariant ideal I (with respect to the P -order \preceq) if and only if

$$\text{in}_{\prec}(I) = \text{in}_{\prec}(G).$$

Of course, the set I can itself be considered a Gröbner basis for the ideal I , so the interest theoretically and computationally is when a finite Gröbner basis exists. One goal of this section is to arrive at a criterion for \preceq guaranteeing that finite P -Gröbner bases exist for all P -invariant I .

In analogy with the classical case, a P -Gröbner basis generates the ideal up to the action of P . Here, for an R -module M and a subset $G \subseteq M$, the submodule $\langle G \rangle_R \subseteq M$ is the R -module generated by G .

Proposition 2.10. *If G is a P -Gröbner basis for a P -invariant ideal $I \subseteq \mathbb{K}[Q]$, then*

$$I = \langle G \rangle_{\mathbb{K}[Q]*P}.$$

Proof. Since I is P -invariant, we have $\langle G \rangle_{\mathbb{K}[Q]*P} \subseteq I$. Conversely, given $f_1 \in I$, we shall prove $f \in \langle G \rangle_{\mathbb{K}[Q]*P}$. Since $\text{in}_\prec(f_1) \in \text{in}_\prec(I) = \text{in}_\prec(G)$, there exist $q_1 \in Q$, $p_1 \in P$, and $g_1 \in G$ such that $\text{in}_\prec(f_1) = \text{in}_\prec(q_1 p_1 g_1)$. Thus, for some $c_1 \in \mathbb{K}$, the element

$$f_2 := f_1 - c_1 q_1 p_1 g_1$$

is either zero or has a smaller initial monomial than $\text{in}_\prec(f_1)$. Also, $f_2 \in I$, so there are $q_2 \in Q$, $p_2 \in P$, and $g_2 \in G$ such that $\text{in}_\prec(f_2) = \text{in}_\prec(q_2 p_2 g_2)$. As before, we define a new polynomial $f_3 := f_2 - c_2 q_2 p_2 g_2$, which again is zero or has a smaller initial term. Continuing in this way, we produce a sequence of polynomials $f_1, f_2, f_3, \dots \in I$ all of whose initial terms form an infinite decreasing sequence. Since \preceq is a well-order, this sequence must terminate in a finite number of steps with some $f_{k+1} = 0$. But then we have that $f_1 = \sum_{i=1}^k c_i q_i p_i g_i$ with the $g_i \in G$. This proves the proposition. \square

If P -divisibility $|_P$ generates enough relations between elements of Q , then finite Gröbner bases for P -invariant ideals always exist. To state this result precisely, however, we need to introduce some basic definitions from order theory.

Recall that a *well-partial-ordering* \leq on a set S is a partial order such that (1) there are no infinite collections of pairwise incomparable elements (i.e., *antichains*) and (2) there are no infinite strictly decreasing sequences. This definition is a natural generalization of the notion of “well-ordering” when \leq is not total. A *final segment* is a subset $F \subseteq S$ which is closed upwards: $s \leq t$ and $s \in F \Rightarrow t \in F$ for all $s, t \in S$. Given a subset $B \subseteq S$, the set

$$\mathcal{F}(B) := \{t \in S : b \leq t \text{ for some } b \in B\}$$

is a final segment of S , the *final segment generated by B* . For example, with P -divisibility $|_P$ as the partial order, the set of monomials $\text{in}_\prec(G)$ is a final segment generated by the initial monomials of G . Thus, another way to state Definition 2.9 is to say that a subset $G \subseteq I$ is a P -Gröbner basis of I if and only if the final segment generated by the leading monomials of G contains all the leading monomials of I .

Continuing further with order terminology, let us call an infinite sequence s_1, s_2, \dots in S *good* if $s_i \leq s_j$ for some indices $i < j$, and *bad* otherwise. The following elementary characterization of well-partial-orderings is classical [16].

Proposition 2.11. *The following are equivalent for a partial order \leq on a set S :*

- (1) S is well-partially-ordered.
- (2) Every infinite sequence in S is good.
- (3) Every infinite sequence in S contains an infinite increasing subsequence.
- (4) Any final segment of S is finitely generated.
- (5) The ascending chain condition holds for final segments of S .

We now have all the ingredients to prove that finite P -Gröbner bases exist when P -divisibility is a well-partial-ordering (our finiteness criterion). In the case that $Q = \mathbb{N}^k$, $P = \{\mathbf{1}\}$, and \preceq is any term order on Q , the theorem says that a finite Gröbner basis exists if monomial divisibility is a well-partial-order. As this is the basic content of Dickson's Lemma, we recover the classical finiteness result for Gröbner bases in polynomial rings with a finite number of variables.

Theorem 2.12. *Let \preceq be a P -order. If P -divisibility $|_P$ is a well-partial-ordering, then every P -invariant ideal $I \subseteq \mathbb{K}[Q]$ has a finite P -Gröbner basis with respect to \preceq . Moreover, if elements of P send monomials to scalar multiples of monomials, the converse holds.*

Proof. The set of monomials $\text{in}_{\prec}(I)$ is a final segment with respect to P -divisibility; thus, it is finitely generated by Proposition 2.11. These generators are initial monomials of a finite subset G of elements of I . It follows that G is a P -Gröbner basis.

For the second statement, we verify that (4) holds in the characterization of Proposition 2.11. Let M be any final segment of Q with respect to $|_P$, and set $I = \langle M \rangle_{\mathbb{K}[Q]*P}$. By assumption, there is a finite set $G = \{g_1, \dots, g_k\} \subseteq I$ such that

$$M \subseteq \text{in}_{\prec}(I) = \text{in}_{\prec}(G) = \mathcal{F}(\{\text{in}_{\prec}(g_1), \dots, \text{in}_{\prec}(g_k)\}).$$

Now, each $g \in G$ has a representation of the form

$$g = \sum_{j=1}^d c_j q_j p_j m_j, \quad c_j \in \mathbb{K}, \quad q_j \in Q, \quad p_j \in P, \quad m_j \in M,$$

and since elements of P send monomials to scalar multiples of monomials, it follows that $\text{in}_{\prec}(g) = q \cdot \text{in}_{\prec}(pm)$ for some $q \in Q$, $p \in P$, and $m \in M$. In particular, we have $m |_P \text{in}_{\prec}(g)$. Thus, $\mathcal{F}(\{\text{in}_{\prec}(g_1), \dots, \text{in}_{\prec}(g_k)\}) \subseteq M$ and M is finitely generated. \square

Remark 2.13. Define a *monomial map* to be an element $p \in P$ that sends monomials to scalar multiples of monomials. Theorem 2.12 says that for a monoid P of monomial maps, P -divisibility is a well-partial-ordering if and only if every P -invariant ideal has a finite P -Gröbner basis. In our applications, the monoids P consist entirely of monomial maps. However, we do not know if the converse to Theorem 2.12 continues to hold when P is a more general set of maps, and it would be interesting to understand this situation better.

Using Proposition 2.10, the following finiteness result is immediate.

Corollary 2.14. *Let \preceq be a P -order. If P -divisibility $|_P$ is a well-partial-ordering, then every P -invariant ideal $I \subseteq \mathbb{K}[Q]$ is finitely generated over $\mathbb{K}[Q] * P$. In other words, $\mathbb{K}[Q]$ is a Noetherian $\mathbb{K}[Q] * P$ -module.*

We next define a general setup that allows us to go from global generation to local stabilization (Theorem 2.19). This can be seen as an analogue to [2, Theorem 4.7] which guaranteed stabilization of certain $\mathfrak{S}_{\mathbb{P}}$ -invariant chains over a polynomial ring in an infinite number of indeterminates. In fact, we shall show in the next section how the stabilization result of [2] follows from our theory.

Definition 2.15 (Filtrations). Let \preceq be a P -order, and suppose that $Q_n \subseteq Q$ and $P_{n,m} \subseteq P$ for nonnegative integers $m \geq n$. We say that Q_n and $P_{n,m}$ is a *filtration* of $\mathbb{K}[Q] * P$ if

- (1) Each Q_n is a submonoid of Q .
- (2) $Q_n \subseteq Q_{n+1}$ for all n .
- (3) $Q = \bigcup_n^\infty Q_n$ and $P = \bigcup_{n,m=1}^\infty P_{n,m}$.
- (4) $P_{n,m}Q_n \subseteq \mathbb{K}[Q_m]$ for all $m \geq n$.
- (5) Each $P_{n,m}$ contains the identity endomorphism.
- (6) If $q \in Q_n \setminus Q_{n-1}$ and $\text{in}_<(pq) \in Q_m$ for some $p \in P$, then there exists $p' \in P_{n,m}$ with $\text{in}_<(p'q) = \text{in}_<(pq)$.
- (7) Each Q_n is an *initial segment* with respect to \preceq (i.e., $u \preceq v$ and $v \in Q_n \Rightarrow u \in Q_n$).

Remark 2.16. From Lemma 2.3, we have $q_1 \preceq q_1q_2$ and $q_2 \preceq q_1q_2$ for any $q_1, q_2 \in Q$. In particular, if $q_1q_2 \in Q_n$, then (7) above implies that both $q_1, q_2 \in Q_n$.

Our most important example of a filtration arises from decomposing the monoid of increasing functions. It appears explicitly in the statements of Theorem 3.6 and Corollary 3.7, and will be used to prove the independent set conjecture of [14, Conj. 4.6] (Theorem 4.7).

Given a filtration of $\mathbb{K}[Q]*P$, we are interested in increasing chains I_\circ of ideals $I_n \subseteq \mathbb{K}[Q_n]$:

$$I_\circ := I_1 \subseteq I_2 \subseteq \cdots \subseteq I_n \subseteq \cdots,$$

simply called *chains* below. Of primary importance is when these ideals stabilize “up to the action” of the monoid P . For the purposes of this work, we will only consider a special class of chains; namely, a *P-invariant chain* is one for which $P_{n,m}I_n \subseteq I_m$ for all $m \geq n$. The stabilization definition alluded to above is as follows.

Definition 2.17. The *P-invariant chain* I_\circ *stabilizes* if there exists a positive integer n_0 such that

$$I_n = \bigcup_{k \leq n_0} \langle P_{k,n}I_k \rangle_{\mathbb{K}[Q_n]} \quad \text{for all } n \geq n_0.$$

In other words, a *P-invariant chain* I_\circ stabilizes when the ideals I_n can be generated by “lifting” the finite set of ideals $\{I_1, \dots, I_{n_0}\}$ in the chain.

Any *P-invariant chain* I_\circ naturally gives rise to an ideal $\mathcal{N}(I_\circ)$ over $\mathbb{K}[Q]*P$ by way of

$$\mathcal{N}(I_\circ) := \bigcup_{n \geq 1} I_n.$$

It is easy to see that if I_\circ stabilizes, then any set of $\mathbb{K}[Q_{n_0}]$ -generators for I_{n_0} will form a generating set of the $\mathbb{K}[Q]*P$ -module $I = \mathcal{N}(I_\circ)$. Our next result says that one can also move from global generation to chain stabilization; it will be a consequence of the following technical fact.

Lemma 2.18. *Let \preceq be a P-order and fix a filtration of $\mathbb{K}[Q]*P$. Suppose that $G \subseteq \mathbb{K}[Q_{n_0}]$ is a finite P-Gröbner basis for a P-invariant ideal $I \subseteq \mathbb{K}[Q]$. Then, if $0 \neq f \in \mathbb{K}[Q_n] \cap I$ with $n \geq n_0$, we have,*

$$\text{in}_<(f) = \text{in}_<(qpg) \quad \text{for some } q \in Q_n, p \in P_{k,n}, g \in G \cap \mathbb{K}[Q_k], k \leq n_0.$$

Proof. Let $0 \neq f \in \mathbb{K}[Q_n] \cap I$. Since G is a *P-Gröbner basis*, it follows that

$$\text{in}_<(f) = q \cdot \text{in}_<(pg)$$

for some $q \in Q$, $p \in P$, and $g \in G$. Since $\text{in}_<(f) \in Q_n$, Remark 2.16 implies that $q \in Q_n$ and $\text{in}_<(pg) \in Q_n$. Let $k \leq n_0$ be such that $\text{in}_<(g) \in Q_k \setminus Q_{k-1}$. From the exchange property (6) of Definition 2.15, there is a $p' \in P_{k,n}$ such that $\text{in}_<(p'g) = \text{in}_<(pg)$. Thus, $\text{in}_<(f) = q \cdot \text{in}_<(p'g) = \text{in}_<(qp'g)$. Finally, since Q_k is an initial segment, it follows that $g \in \mathbb{K}[Q_k]$. \square

Theorem 2.19 (Chain stabilization). *Let \preceq be a P -order. If P -divisibility $|_P$ is a well-partial-ordering, then every P -invariant chain stabilizes.*

Proof. Given an invariant chain I_\circ , construct the P -invariant ideal $I = \mathcal{N}(I_\circ)$ of $\mathbb{K}[Q]$, and let G be a finite P -Gröbner basis for I by Theorem 2.12. The result now follows from Lemma 2.18 using a descent argument as in Proposition 2.10. \square

3. EXAMPLES, COUNTEREXAMPLES, AND FIRST APPLICATIONS

In this section, we begin to apply the abstract theory from Section 2 to specific examples that make frequent appearances in applications. Although the finite Gröbner basis results we derive are for ideals invariant under the monoid of increasing functions, we can easily produce corollaries for the more familiar setting of ideals that are stable under a symmetric group action. In Section 4, we apply these ideas to Markov bases and other implicitization problems in algebraic statistics.

Our main monoid P of interest for constructing monoidal Gröbner bases will be the monoid of increasing functions (the *shift monoid*):

$$\Pi := \{\pi : \mathbb{P} \rightarrow \mathbb{P} : \pi(i) < \pi(i+1) \text{ for all } i \in \mathbb{P}\}.$$

Given a set R , let $X_R = \{x_r : r \in R\}$ denote the set of indeterminates indexed by R , and let $\mathbb{K}[X_R]$ be the (commutative) polynomial ring with coefficients in \mathbb{K} and indeterminates X_R . Of special interest is when R is a product of the form $R = R_1 \times \cdots \times R_m$. For $r \in \mathbb{P}$, let $[r] = \{1, 2, \dots, r\}$. Our first result concerns the case $R = [r] \times \mathbb{P}$ with the (linear) action of Π on $\mathbb{K}[X_{[r] \times \mathbb{P}}]$ being generated by its action on the second index of the indeterminates $X_{[r] \times \mathbb{P}}$:

$$\pi x_{i,j} := x_{i,\pi(j)}, \quad \pi \in \Pi.$$

Theorem 3.1. *The column-wise lexicographic term order $x_{i,j} \preceq x_{k,l}$ if $j < l$ or $(j = l$ and $i \leq k)$ is a Π -order on $\mathbb{K}[X_{[r] \times \mathbb{P}}]$ such that Π -invariant ideals of $\mathbb{K}[X_{[r] \times \mathbb{P}}]$ have finite Π -Gröbner bases. In particular, the ring $\mathbb{K}[X_{[r] \times \mathbb{P}}]$ is a Noetherian $\mathbb{K}[X_{[r] \times \mathbb{P}}] * \Pi$ -module.*

We call the Π -divisibility order induced by the column-wise lexicographic order in the statement of Theorem 3.1 the *shift order*. We shall prove Theorem 3.1 using Theorem 2.12 by showing that the Π -divisibility partial order on monomials in $\mathbb{K}[X_{[r] \times \mathbb{P}}]$ is a well-partial-order. Before verifying this fact, we recall the notion of a Higman partial order associated to a well-partial-order.

Definition 3.2 (The Higman Partial Order). Let (S, \preceq) be a partially-ordered set. Let (S_H, \preceq_H) be defined on the set $S_H = S^*$ of finite words of elements of S by:

$$u_1 u_2 \cdots u_n \preceq_H v_1 v_2 \cdots v_m$$

if and only if there is a $\pi \in \Pi$ such that $u_i \preceq v_{\pi(i)}$ for $i \in [n]$.

The main result about Higman partial orders is Higman's Lemma [13, 17].

Lemma 3.3 (Higman's Lemma). *If (S, \preceq) is a well-partial-order, then the Higman partial order (S_H, \preceq_H) is also a well-partial-order.*

Below, we shall apply Higman's Lemma to the set $S = \mathbb{N}^r$, partially ordered by inequality:

$$(s_1, \dots, s_r) \preceq (t_1, \dots, t_r) :\Leftrightarrow s_i \leq t_i \text{ for } i = 1, \dots, r.$$

This is a well-partial-order by Dickson's Lemma, and it can be interpreted as a well-partial-ordering on the monomials of $\mathbb{K}[X_{[r] \times \mathbb{P}}]$.

Example 3.4. In the Higman ordering on words $(\mathbb{N}^2)^*$ induced by the partial order above,

$$(1, 2)(3, 1)(2, 5) \preceq_H (1, 0)(1, 4)(5, 2)(1, 2)(2, 7),$$

witnessed by any shift monoid element $\pi \in \Pi$ that has $\pi(1) = 2$, $\pi(2) = 3$, $\pi(3) = 5$. Interpreted as a Π -divisibility relation between monomials in the polynomial ring $\mathbb{K}[X_{[2] \times \mathbb{P}}]$, this statement reads:

$$x_{1,1}x_{1,2}x_{2,2}^4x_{1,3}^5x_{2,3}^2x_{1,4}x_{2,4}^2x_{1,5}^2x_{2,5}^7 = x_{1,1}x_{2,2}^2x_{1,3}^2x_{2,3}x_{1,4}x_{2,4}^2x_{2,5}^2 \cdot \pi(x_{1,1}x_{2,1}^2x_{1,2}^3x_{2,2}^2x_{1,3}^5x_{2,3}^5).$$

□

Proof of Theorem 3.1. We first show that the column-wise lexicographic order is a Π -order on $\mathbb{K}[X_{[r] \times \mathbb{P}}]$. Each monomial in $\mathbb{K}[X_{[r] \times \mathbb{P}}]$ has the form $x^u = x_1^{u_1} \cdots x_n^{u_n}$ for some $n \in \mathbb{P}$, where $x_j^{u_j} = \prod_{i \in [r]} x_{i,j}^{u_{i,j}}$. Suppose that $x^u \prec x^v$. Then we can write

$$x^u = x_1^{u_1} \cdots x_k^{u_k} x_{k+1}^{v_{k+1}} \cdots x_m^{v_m}$$

for some $k \leq n$ in which $x_k^{u_k} \prec x_k^{v_k}$. For $\pi \in \Pi$, we have

$$\pi x^u = x_{\pi(1)}^{u_1} \cdots x_{\pi(k)}^{u_k} x_{\pi(k+1)}^{v_{k+1}} \cdots x_{\pi(n)}^{v_n},$$

$$\pi x^v = x_{\pi(1)}^{v_1} \cdots x_{\pi(k)}^{v_k} x_{\pi(k+1)}^{v_{k+1}} \cdots x_{\pi(n)}^{v_n}.$$

Since π is increasing, the right-most column index where πx^u and πx^v disagree is at $\pi(k)$, in which case $x_{\pi(k)}^{u_k} \prec x_{\pi(k)}^{v_k}$ so that $\pi x^u \prec \pi x^v$. Since multiplication by an ordinary monomial preserves \preceq for any term order, this proves that \preceq is a Π -order by Proposition 2.4.

Next, we must show that Π -divisibility on $\mathbb{K}[X_{[r] \times \mathbb{P}}]$ is a well-partial-order. In the Π -divisibility partial order, we have $x^u \mid_{\Pi} x^v$ if and only if there is a $\pi \in \Pi$ such that $\pi x^u \mid x^v$ (monomial division-wise). In turn, this happens if and only if there is a $\pi \in \Pi$ such that $x_{\pi(i)}^{u_i} \mid x_{\pi(i)}^{v_{\pi(i)}}$ for each $i \in [n]$. In other words, Π -divisibility is the Higman partial order of the standard divisibility partial order on the monomials of $\mathbb{K}[X_{[r] \times \mathbb{P}}]$ (viewed as elements of $(\mathbb{N}^r)^*$). Thus, Higman's Lemma implies that Π -divisibility is a well-partial-order. Theorem 2.12 now implies that $\mathbb{K}[X_{[r] \times \mathbb{P}}]$ has finite Gröbner bases; in particular, by Corollary 2.14 it is a Noetherian $\mathbb{K}[X_{[r] \times \mathbb{P}}] * \Pi$ -module. □

As a corollary to Theorem 3.1, we also deduce the Noetherian property for ideals that are stable under the action of the symmetric group $\mathfrak{S}_{\mathbb{P}}$. This was Theorem 1.1 from the introduction.

Corollary 3.5. *The polynomial ring $\mathbb{K}[X_{[r] \times \mathbb{P}}]$ is a Noetherian $\mathbb{K}[X_{[r] \times \mathbb{P}}] * \mathfrak{S}_{\mathbb{P}}$ -module.*

Proof. Each polynomial $f \in \mathbb{K}[X_{[r] \times \mathbb{P}}]$ depends on only finitely many column indices. Thus, if $\pi \in \Pi$, there exists $\sigma \in \mathfrak{S}_{\mathbb{P}}$ such that $\sigma \cdot f = \pi \cdot f$. Indeed, if the largest column index appearing in f is m , then σ can be chosen to be the identity on all $i > \pi(m)$. This implies that every $\mathfrak{S}_{\mathbb{P}}$ -stable ideal I is Π -stable and any $\mathbb{K}[X_{[r] \times \mathbb{P}}] * \Pi$ generating set of I is a $\mathbb{K}[X_{[r] \times \mathbb{P}}] * \mathfrak{S}_{\mathbb{P}}$ generating set. \square

Note that the $r = 1$ version of Corollary 3.5 is a main result of [2, 4]. Our proof and the material in Section 2 constitute a distillation and generalization of the proof in those papers. A second corollary concerns infinite chains of symmetric ideals, each contained in a finite polynomial ring.

Before stating this result, we must first introduce a filtration of $\mathbb{K}[X_{[r] \times \mathbb{P}}] * \Pi$. Let $Q_n \cong \mathbb{N}^{r \times n}$ be the set of monomials in the polynomial ring $\mathbb{K}[X_{[r] \times [n]}]$, and for $m \geq n$, let $\Pi_{n,m} \subset \Pi$ be the set of functions

$$\Pi_{n,m} = \{\pi \in \Pi : \pi([n]) \subseteq [m]\}.$$

Theorem 3.6. *The sets Q_n and $\Pi_{n,m}$ form a filtration of $\mathbb{K}[X_{[r] \times \mathbb{P}}] * \Pi$. In particular, every Π -invariant ascending chain I_{\circ} stabilizes.*

Proof. The seven conditions of Definition 2.15 are easy to check. The most difficult to parse is (6), which we describe in detail. In our setting, condition (6) says that if a monomial $x^u = x_1^{u_1} \cdots x_n^{u_n}$ has $u_n \neq 0$ and $\pi \in \Pi$ has $\pi(n) \leq m$, then there is a $\pi' \in \Pi_{n,m}$ such that $\pi'(x^u) = \pi(x^u)$. But if $\pi \in \Pi$ satisfies $\pi(n) \leq m$, then $\pi \in \Pi_{n,m}$ (since it is increasing). In particular, we can take $\pi' = \pi$. The second statement follows from Theorem 2.19 and the fact that Π -divisibility is a well-partial-order (from the proof of Theorem 3.1). \square

The most important and useful implication of Theorem 3.6 is the following corollary, which concerns chains of ideals stable under the action of the symmetric group. It is this fact, and its variations, that allow us to prove the theorems in algebraic statistics that appear in the next section. For simplicity of notation, we write \mathfrak{S}_n for $\mathfrak{S}_{[n]}$ below.

Corollary 3.7. *For each $n \in \mathbb{P}$, let $I_n \in \mathbb{K}[X_{[r] \times [n]}]$ be a \mathfrak{S}_n -invariant ideal. Suppose that the I_n form an invariant ascending chain:*

$$\mathfrak{S}_m I_n \subseteq I_m, \quad \text{for each } n \leq m.$$

Then there exists an $n_0 \in \mathbb{P}$ such that for all $n > n_0$, we have

$$\langle \mathfrak{S}_n I_{n_0} \rangle_{\mathbb{K}[X_{[r] \times [n]}]} = I_n.$$

In other words, ascending invariant chains are finitely generated up to symmetry.

Proof. An ascending invariant chain I_{\circ} with respect to the filtration of

$$\mathfrak{S}_{(\mathbb{P})} := \bigcup_{n \in \mathbb{P}} \mathfrak{S}_n$$

by the \mathfrak{S}_n is also an ascending invariant chain with respect to the filtration of Π by the $\Pi_{n,m}$. Hence, there exists an n'_0 with respect to which each I_n with $n \geq n'_0$ is generated by the generators of $I_{n'_0}$. Since $\Pi_{n,m} I_n \subseteq \mathfrak{S}_m I_n$, for all $m \geq n$, this implies that $n_0 = n'_0$ is sufficient in the corollary. \square

Beyond Theorem 3.1, we would like to have more general settings in which there is a priori knowledge that some family of ideals is finitely generated. In a certain sense, Theorem 3.1 is best possible for infinite dimensional polynomial rings, as the following example illustrates.

Example 3.8. The polynomial ring $\mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}]$ is naturally a $\mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}] * (\mathfrak{S}_{\mathbb{P}} \times \mathfrak{S}_{\mathbb{P}})$ -module, but this module is not Noetherian. For instance, the ideal

$$I = \langle x_{11}x_{12}x_{22}x_{21}, x_{11}x_{12}x_{22}x_{23}x_{33}x_{31}, \dots, x_{11}x_{12}x_{22} \cdots x_{mm}x_{m1}, \dots \rangle$$

is not finitely generated as a $\mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}] * (\mathfrak{S}_{\mathbb{P}} \times \mathfrak{S}_{\mathbb{P}})$ -module. Via the natural correspondence between square-free monomials in doubly indexed variables and bipartite graphs, the sequence of generators listed above are even length cycles. The fact that no even length cycle is a subgraph of any other even length cycle implies that this ideal is not finitely generated. \square

In spite of Example 3.8, it is possible to extend Theorem 3.1 via the theory from Section 2 to prove that certain ideals in rings such as $\mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}]$ are finitely generated up to the action of $\mathfrak{S}_{\mathbb{P}} \times \mathfrak{S}_{\mathbb{P}}$. This is done by combining the following elementary proposition with Corollary 3.11 below. The idea will be to focus on Π -stable ideals $J \subseteq \mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}]$ which contain a subideal $I \subseteq J$ such that $\mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}]/I$ is Noetherian (see Example 3.12).

Proposition 3.9. *Suppose that $L \subseteq M \subseteq N$ are R -modules, that L is finitely generated, and that N/L is a Noetherian R -module. Then M is finitely generated.*

Proof. Since N/L is Noetherian, M/L has a finite generating set, with representatives in M . These generators along with the generators of L form a finite generating set of M . \square

We next consider a natural class of rings that inherit Noetherianity from being contained in a Noetherian semigroup ring. The goal in applications will be to show that quotients as above are isomorphic to one of these special rings.

Definition 3.10. A subsemigroup ring $\mathbb{K}[Q'] \subseteq \mathbb{K}[Q]$ is called *divisible* if $q_1, q_2 \in Q'$ and $q_1 = q_3 q_2$ implies that $q_3 \in Q'$. The subsemigroup ring $\mathbb{K}[Q']$ is *P -invariant* if for all $q \in Q'$ and $p \in P$, we have $pq \in \mathbb{K}[Q']$.

Corollary 3.11. *Let $\mathbb{K}[Q']$ be a divisible P -invariant subring of $\mathbb{K}[Q]$. If \preceq is a P -order on Q , then*

- (1) \preceq is a P -ordering on Q' .

If, in addition, P -divisibility is a well-partial-ordering on Q , then

- (2) P -divisibility is a well-partial-ordering on Q' and
- (3) P -invariant ideals of $\mathbb{K}[Q']$ have finite P -Gröbner bases.

*If, in addition Q_n and $P_{n,m}$ are a filtration of $\mathbb{K}[Q] * P$, then*

- (4) $Q'_n = Q' \cap Q_n$ and $P_{n,m}$ are a filtration $\mathbb{K}[Q'] * P$, and
- (5) invariant chains I_\circ with $I_n \in \mathbb{K}[Q'_n]$ stabilize.

Proof. (1) If $q_1, q_2 \in Q'$, then $Q' \subseteq Q$ implies that the condition of Proposition 2.4 is satisfied. In particular, (1) holds regardless of whether Q' is divisible.

(2) Consider any infinite sequence of monomials in Q' . Since P -divisibility is a well-partial-ordering on Q , this sequence is good when considered as a subset of Q . Since $Q' \subseteq Q$ is

a divisible subsemigroup, the sequence is also good in Q' . Proposition 2.11 implies that P -divisibility is also a well-partial-order on Q' .

(3) This follows from (2) and Theorem 2.12.

(4) There are seven conditions to check in the definition of a filtration; all of them are straightforward.

(5) This follows from (4) and Theorem 2.19. \square

We close this section with an example illustrating how Proposition 3.9 and Corollary 3.11 will be used in Section 4.

Example 3.12. Consider $\mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}]$ as a module over $\mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}] * \Pi$ with Π acting on both indices simultaneously (i.e. $\pi x_{i,j} = x_{\pi(i),\pi(j)}$). By Example 3.8, $\mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}]$ is not a Noetherian $\mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}] * \Pi$ -module. However, consider a Π -stable ideal $J \subseteq \mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}]$ such that $I \subseteq J$, where

$$I = \left\langle \det \begin{pmatrix} x_{i_1, j_1} & x_{i_1, j_2} \\ x_{i_2, j_1} & x_{i_2, j_2} \end{pmatrix} : i_1, i_2, j_1, j_2 \in \mathbb{P} \right\rangle_{\mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}]}$$

is the ideal of two-by-two minors of the matrix $X_{\mathbb{P} \times \mathbb{P}}$. We have the following isomorphism of $\mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}] * \Pi$ -modules:

$$\mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}]/I \cong \mathbb{K}[y_{1,i}y_{2,j} : i, j \in \mathbb{P}],$$

the map being induced by $x_{i,j} \mapsto y_{1,i}y_{2,j}$. Thus, $\mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}]/I$ has the same module structure as that of $\mathbb{K}[y_{1,i}y_{2,j} : i, j \in \mathbb{P}]$ as a $\mathbb{K}[y_{1,i}y_{2,j} : i, j \in \mathbb{P}] * \Pi$ -module. Since $\mathbb{K}[y_{1,i}y_{2,j} : i, j \in \mathbb{P}]$ is a Π -stable divisible semigroup ring that is a subring of $\mathbb{K}[Y_{[2] \times \mathbb{P}}]$, we see that $\mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}]/I$ is a Noetherian $\mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}] * \Pi$ -module by Corollary 3.11. Since I is also finitely generated as a $\mathbb{K}[X_{\mathbb{P} \times \mathbb{P}}] * \Pi$ -module, it follows that J is finitely generated by Proposition 3.9. \square

4. APPLICATIONS IN ALGEBRAIC STATISTICS

In this section, we apply the theory developed in the previous two sections to give new proofs of some classical finiteness results about Markov bases of hierarchical models [14, 19]. These finiteness theorems follow from Corollary 3.7 for finite generation of chains of increasing symmetric ideals. We also extend these results using Proposition 3.9 and Corollary 3.11 to give an affirmative solution to the independent set conjecture [14, Conj. 4.6]. Finally, we explain how these finiteness results extend beyond hierarchical models to more general statistical models.

Let $r_1, \dots, r_m \in \mathbb{P}$ and for a subset $F \subseteq [m]$, set $R_F = \prod_{i \in F} [r_i]$ to be the Cartesian product of the index sets $[r_i]$. If $F = [m]$, we use the shorthand $R = R_F$. For an algebraic object \mathbb{A} (e.g. a field, semiring, monoid) and a finite set M , let \mathbb{A}^M be the Cartesian product of \mathbb{A} with itself $\#M$ times, with coordinates indexed by M .

Let $\mathbf{i} \in R$ denote an index vector. For each $F \subset [m]$, let $\mathbf{i}_F := (i_f)_{f \in F}$ be the substring $\mathbf{i}_F \in R_F$ obtained from \mathbf{i} by the natural projection. For $u \in \mathbb{R}^R$ and $F \subseteq [m]$, also let $u|_F \in \mathbb{R}^{R_F}$ be the F -marginal of u , defined by linearly extending

$$e_{\mathbf{i}}|_F := e_{\mathbf{i}_F}.$$

Here, $e_{\mathbf{i}}$ is the standard unit table in \mathbb{R}^R , having a 1 in the \mathbf{i} position and zero elsewhere, and similarly $e_{\mathbf{i}_F}$ is the standard unit table in \mathbb{R}^{R_F} .

Given a collection $\Gamma = \{F_1, F_2, \dots\}$ of subsets of $[m]$, we define the Γ -marginal map by

$$\pi_{\Gamma,r} : \mathbb{R}^R \rightarrow \bigoplus_{F \in \Gamma} \mathbb{R}^{R_F}$$

$$u \mapsto (u|_{F_1}, u|_{F_2}, \dots).$$

From the linear transformation $\pi_{\Gamma,r}$, we can extract the matrix $A_{\Gamma,r}$ representing it. This matrix $A_{\Gamma,r}$ is called the *design matrix* of the hierarchical model associated to Γ in algebraic statistics [12].

Associated to any linear transformation $A : \mathbb{Z}^r \rightarrow \mathbb{Z}^d$ is the lattice $\ker_{\mathbb{Z}} A$. Among the many important spanning sets for a lattice are the Markov bases, which are special sets that allow one to take random walks over the fibers $(u + \ker_{\mathbb{Z}} A) \cap \mathbb{N}^r$.

Definition 4.1. A *Markov basis* for the matrix A (or lattice $\ker_{\mathbb{Z}} A$) is a finite subset $\mathcal{B} \subset \ker_{\mathbb{Z}}(A)$ such that for all $u, v \in \mathbb{N}^r$ with $Au = Av$, there exists a sequence $b_1, \dots, b_L \in \mathcal{B}$ such that

$$u = v + \sum_{i=1}^L b_i \quad \text{and} \quad v = u + \sum_{i=1}^l b_i \in \mathbb{N}^r, \quad l = 1, 2, \dots, L.$$

Elements of a Markov basis are called *moves*. A Markov basis for A is *minimal* if no proper subset is a Markov basis of A .

Markov bases of the matrices $A_{\Gamma,r}$ are useful for performing statistical hypothesis tests by running random walks over contingency tables (see [7] or Chapter 1 in [12]). Note that Markov bases are not in general unique, even if we assume the Markov basis is minimal.

One of the main mathematical questions about Markov bases of hierarchical models is the following: How does the structure of a Markov basis depend on Γ and r_1, \dots, r_m ? A specific problem of this type is to determine what finiteness properties of the Markov bases should be expected when we fix Γ and send one or more values of $r_i \rightarrow \infty$. Questions about finiteness for these Markov bases are natural in our setting because the lattice $\ker_{\mathbb{Z}} A_{\Gamma,r}$ is stable under the action of the product of symmetric groups $\mathfrak{S}_{r_1} \times \dots \times \mathfrak{S}_{r_m}$, where \mathfrak{S}_{r_i} acts by permuting the i th index. Furthermore, given any Γ , $r \in \mathbb{P}^m$, and $t \in \mathbb{N}^m$, a vector $b \in \ker_{\mathbb{Z}} A_{\Gamma,r}$ can be naturally lifted into $\ker_{\mathbb{Z}} A_{\Gamma,r+t}$ by padding with zeroes. Denote the resulting vector by $\text{pad}_{r+t}(b)$.

We now make precise the notion of sending some $r_i \rightarrow \infty$. First, fix a collection of indices $T \subseteq [m]$ which will “go to infinity”. For each fixed set of values r_i with $i \in [m] \setminus T$, we consider the Markov bases of the matrices $A_{\Gamma,r}$ (here, r_i is allowed to be arbitrary when $i \in T$). We have *finite Markov bases up to symmetry* in this situation if for every fixed set of values r_i with $i \in [m] \setminus T$, there exist r_i with $i \in T$ and a finite set of moves $\mathcal{B} \subseteq \ker_{\mathbb{Z}} A_{\Gamma,r}$, such that for all $t \in \mathbb{N}^m$ with $t_i = 0$ for $i \in [m] \setminus T$, the set

$$\mathfrak{S}_{r_1+t_1} \times \dots \times \mathfrak{S}_{r_m+t_m} \cdot \{\text{pad}_{r+t}(b) : b \in \mathcal{B}\}$$

is a Markov basis for $A_{\Gamma,r+t}$. Otherwise, there is no finite Markov basis up to symmetry.

We represent two examples illustrating that in some situations the Markov basis is finite up to symmetry and in other cases it is not.

Example 4.2. Let $\Gamma = \{\{1\}, \{2\}\}$. Then $\pi_\Gamma : \mathbb{Z}^{[r_1] \times [r_2]} \rightarrow \mathbb{Z}^{[r_1]} \oplus \mathbb{Z}^{[r_2]}$ is the map that computes the row and column sums of an $r_1 \times r_2$ table. Thus $\ker_{\mathbb{Z}} A_\Gamma$ consists of all integral tables whose row and column sums are equal to zero.

If both $r_1, r_2 \geq 2$, the Markov basis for this model consists of the $2\binom{r_1}{2}\binom{r_2}{2}$ moves:

$$\mathcal{B} = \{e_{i_1 j_1} + e_{i_2 j_2} - e_{i_1 j_2} - e_{i_2 j_1} : i_1, i_2 \in [r_1], j_1, j_2 \in [r_2]\}.$$

For example, for $r_1 = 3, r_2 = 4$, a typical element in \mathcal{B} is the 3×4 table

$$\begin{pmatrix} 1 & 0 & -1 & 0 \\ -1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

Up to the natural action of $\mathfrak{S}_{r_1} \times \mathfrak{S}_{r_2}$, permuting rows and columns of the matrices, there is only one move in the Markov basis [7]. \square

On the other hand, these types of finite Markov basis descriptions are known not to hold for general Γ when we let many of the numbers $r_i \rightarrow \infty$.

Example 4.3. Let $\Gamma = \{\{1, 2\}, \{1, 3\}, \{2, 3\}\}$ be the three cycle hierarchical model (also called the model of no 3-way interaction). Then $\pi_\Gamma : \mathbb{Z}^{[r_1] \times [r_2] \times [r_3]} \rightarrow \mathbb{Z}^{[r_1] \times [r_2]} \oplus \mathbb{Z}^{[r_1] \times [r_3]} \oplus \mathbb{Z}^{[r_2] \times [r_3]}$ is the map that computes all 2-way marginals of the three way table u . For all m , the move

$$\sum_{i=1}^m (e_{i,i,1} - e_{i,i,2}) + e_{m,1,2} - e_{m,1,1} + \sum_{i=1}^{m-1} (e_{i,i+1,2} - e_{i,i+1,1})$$

belongs to every minimal Markov basis for Γ for all $r_1, r_2 \geq m$ and $r_3 \geq 2$ [7]. When $r_3 = 2$, these Markov basis elements can be represented as two $r_1 \times r_2$ matrices obtained from extracting slices where $i_3 = 1$ and $i_3 = 2$ respectively. When $r_1 = r_2 = 5$, the Markov basis element is:

$$\begin{pmatrix} 1 & -1 & 0 & 0 & 0 \\ 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 1 & -1 \\ -1 & 0 & 0 & 0 & 1 \end{pmatrix} \quad \begin{pmatrix} -1 & 1 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & -1 & 1 & 0 \\ 0 & 0 & 0 & -1 & 1 \\ 1 & 0 & 0 & 0 & -1 \end{pmatrix}.$$

In particular, Markov bases for Γ are not finite up to $\mathfrak{S}_{r_1} \times \mathfrak{S}_{r_2} \times \mathfrak{S}_{r_3}$ symmetry on $r_1 \times r_2 \times r_3$ arrays for $r_3 \geq 2$ as r_1 and r_2 both tend to infinity. \square

These two examples illustrate a dichotomy between cases where we send more than one of the $r_i \rightarrow \infty$. In some situations the Markov basis is finite up to symmetry, and in other cases it is not. If we only send one of the r_i to infinity, however, there is always a finite Markov basis up to symmetry [14, 19]:

Theorem 4.4. *For any Γ and fixed r_1, \dots, r_{m-1} , there exists an $N = N(\Gamma; r_1, \dots, r_{m-1})$ such that the Markov basis for A_Γ for r_1, \dots, r_m with $r_m > N$ is determined up to symmetry by the Markov basis for $r_1, r_2, \dots, r_{m-1}, N$.*

We provide a new proof of Theorem 4.4 below. An important ingredient will be the fundamental theorem of Markov bases, which translates questions about Markov bases into questions about generating sets of toric ideals.

Given any matrix $A = (a_{ij}) \subseteq \mathbb{Z}^{k \times r}$, consider the ring homomorphism:

$$\phi : \mathbb{K}[x_1, \dots, x_r] \rightarrow \mathbb{K}[y_1^{\pm 1}, \dots, y_k^{\pm 1}], \quad x_j \mapsto \prod_{i=1}^k y_i^{a_{ij}}.$$

The kernel of ϕ is the *toric ideal* $I_A := \ker \phi$. The ideal I_A is a prime ideal that gives an algebraic encoding of the integer kernel of the matrix A since

$$I_A = \langle x^u - x^v : u, v \in \mathbb{N}^r, Au = Av \rangle.$$

Note that $\mathbb{K}[X]/I_A$ is a semigroup ring, the ring generated by the monomials $\phi(x_1), \dots, \phi(x_r)$.

The following theorem establishes the connection between Markov bases of the lattice $\ker_{\mathbb{Z}} A$ and the toric ideal I_A . (Below, the vectors $b^+ \in \mathbb{N}^r$ and $b^- \in \mathbb{N}^r$ are the nonnegative and nonpositive part, respectively, of $b = b^+ - b^- \in \mathcal{B}$).

Theorem 4.5 (Fundamental theorem of Markov bases). [7] *A finite subset $\mathcal{B} \subseteq \ker_{\mathbb{Z}} A$ is a Markov basis for A if and only if the set of binomials*

$$\left\{ x^{b^+} - x^{b^-} : b \in \mathcal{B} \right\}$$

is a generating set of the toric ideal I_A .

Proof of Theorem 4.4. Applying the fundamental theorem of Markov bases, it suffices to show that the associated toric ideals are finitely generated up to symmetry. For each value of $r_m \in \mathbb{P}$, let A_{r_m} be the matrix representing the linear transformation π_{Γ} for a table of size r_1, \dots, r_m . That is $A_{r_m} = A_{\Gamma, r}$, but where we are paying special attention to the changing value of r_m . Each of the ideals $I_{A_{r_m}}$ is contained in $\mathbb{K}[X_{R_{[m-1] \times [r_m]}}]$. Taking $k = \prod_{i=1}^{m-1} r_i$ and identifying $\prod_{i=1}^{m-1} [r_i]$ with $[k]$, we see that each ideal naturally lies in $\mathbb{K}[X_{[k] \times [r_m]}]$. Furthermore, each ideal $I_{A_{r_m}}$ is stable under the action of \mathfrak{S}_{r_m} . On tables, \mathfrak{S}_{r_m} acts by permuting “slices” of the table. The ideals are also nested:

$$\langle \mathfrak{S}_{r_{m+1}} I_{A_{r_m}} \rangle_{\mathbb{K}[X_{[k] \times [r_{m+1}]}]} \subset I_{A_{r_{m+1}}},$$

which on the level of tables corresponds to the fact that we can always add a slice of all zeroes to an element $b \in \ker_{\mathbb{Z}} A_r$ and obtain an element $b' \in \ker_{\mathbb{Z}} A_{r+1}$. Thus, the sequence of ideals I_{A_1}, I_{A_2}, \dots forms an ascending invariant chain. Therefore, by Corollary 3.7 they have a finite generating set up to the filtration of $\mathfrak{S}(\mathbb{P})$ by the \mathfrak{S}_{r_m} . \square

Our new proof of Theorem 4.4 has the advantage over the proofs from [14, 19] that it puts these finiteness properties into a very general framework. On the other hand, the older proofs produce bounds on the number $N(\Gamma; r_1, \dots, r_{m-1})$. Part of the reason for introducing our more general framework is that it can produce finiteness results in situations where the ideas from [14, 19] do not generalize. The technique in [14, 19] is to show that the universal Gröbner basis is finite up to $\mathfrak{S}_{\mathbb{P}}$ symmetry, which implies finite generation up to symmetry (a *universal Gröbner basis* is a set of polynomials that is a Gröbner basis with respect to every term order). That idea does not work in the more general settings considered below because

the universal Gröbner basis is not, in general, finite up to symmetry (e.g. the universal Gröbner basis of the ideal of 2×2 minors in $\mathbb{K}[X_{[k] \times [k]}]$ requires polynomials of degree k).

Example 4.3 shows that there cannot be a general finiteness result when two or more of the r_i are sent to infinity. However, we can still produce finiteness theorems when some of the $r_i \rightarrow \infty$ and Γ satisfies some extra properties.

Definition 4.6. A subset $T \subseteq [m]$ is called an *independent subset* of Γ if $\#(T \cap F) \leq 1$ for all $F \in \Gamma$.

Equivalently, the independent subsets of Γ are precisely the independent sets of the 1-skeleton of Γ (that is, of the underlying graph).

The main theorem of this section is a finiteness property for Markov bases in models Γ that have independent vertex sets. This provides a proof of the independent set conjecture of Hoşten and the second author [14, Conj. 4.6].

Theorem 4.7. *Let $\Gamma \subseteq 2^{[m]}$, and suppose that $T \subseteq [m]$ is an independent set of Γ . Fix the table dimensions r_s such that $s \in [m] \setminus T$. Then A_Γ has a finite Markov basis up to the natural action of $\mathfrak{S}_{r_1} \times \cdots \times \mathfrak{S}_{r_m}$ as $r_t \rightarrow \infty$ for all $t \in T$.*

Proving Theorem 4.7 requires two intermediate results. First of all, we shall need to understand the relationships between toric ideals I_{A_Γ} for varying Γ . Secondly, we will need to understand an important family of Γ that are called decomposable.

One simplification we can make about Γ is to assume it is a simplicial complex; that is, if $S \in \Gamma$ and $T \subseteq S$ then $T \in \Gamma$ as well. We may make this assumption without loss of generality since “the marginal of a marginal is a marginal”. In other words, adding T to Γ when $S \in \Gamma$ and $T \subseteq S$ does not change $\ker A_\Gamma$.

Lemma 4.8. *Suppose that $\Gamma_1 \subseteq \Gamma_2$, in the sense that for each $S \in \Gamma_1$, there is a $T \in \Gamma_2$ such that $S \subseteq T$. Then $\ker A_{\Gamma_2} \subseteq \ker A_{\Gamma_1}$ and the toric ideals satisfy $I_{A_{\Gamma_2}} \subseteq I_{A_{\Gamma_1}}$.*

Proof. If $S \subseteq T$, then $u|_S = (u|_T)|_S$. Thus, if $\Gamma_1 \subseteq \Gamma_2$, the marginal map π_{Γ_1} factors through π_{Γ_2} . \square

A simplicial complex Δ has a *reducible decomposition* (Δ_1, S, Δ_2) if $\Delta = \Delta_1 \cup \Delta_2$, $\Delta_1 \cap \Delta_2 = 2^S$ (where 2^S is the power set of S), and neither Δ_1 nor $\Delta_2 = 2^S$. A simplicial complex with a reducible decomposition is called *reducible*. A simplicial complex is *decomposable* if it is either a simplex (of the form 2^K) or it is reducible and both Δ_1 and Δ_2 are decomposable. The following theorem characterizes the generating sets of the toric ideals I_{A_Γ} whenever Γ is a decomposable simplicial complex.

Theorem 4.9. [8, 21] *If Γ is a decomposable simplicial complex, then I_{A_Γ} is generated by quadratic binomials. As $r_1, \dots, r_m \rightarrow \infty$, there is a finite set of quadratic binomials that generate I_{A_Γ} up to the action of $\mathfrak{S}_{r_1} \times \cdots \times \mathfrak{S}_{r_m}$. Furthermore, let $T \subseteq [m]$ and fix the table dimensions r_s where $s \in [m] \setminus T$, and let $r_t = r$ for all $t \in T$. Let \mathfrak{S}_r act diagonally on $[r]^{\#T}$. Then, the generators of I_{A_Γ} stabilize up to the action of \mathfrak{S}_r after $r \geq 2\#T$.*

Proof of Theorem 4.7. By the fundamental theorem of Markov bases, it suffices to show that the corresponding toric ideals I_{A_Γ} are finitely generated up to symmetry. It also suffices to show the finiteness result when considering the action of a much smaller group contained

inside of $\mathfrak{S}_{r_1} \times \cdots \times \mathfrak{S}_{r_m}$. Namely, we will send $r_t \rightarrow \infty$ for $t \in T$ simultaneously and consider the diagonal action of \mathfrak{S}_r acting on the indices i_t with $t \in T$. This is sufficient because every Markov basis move for a small table embeds as a Markov basis element for a table of larger dimensions, by the padding operation.

For each $r \in \mathbb{P}$, let I_{A_r} be the corresponding toric ideal, which belongs to the ring

$$\mathbb{K}[Q_r] := \mathbb{K}[X_{R_{[m] \setminus T} \times [r]^{\#T}}].$$

Let $\mathbb{K}[Q]$ denote the limiting ring

$$\mathbb{K}[Q] := \mathbb{K}[X_{R_{[m] \setminus T} \times \mathbb{P}^{\#T}}].$$

Let Π act on $\mathbb{K}[Q]$ by acting diagonally on $\mathbb{P}^{\#T}$. Then the Q_r and $\Pi_{n,r}$ form a filtration of $\mathbb{K}[Q] * \Pi$, and the sequence of ideals $I_\circ = I_{A_1} \subseteq I_{A_2} \subseteq \cdots$ is an invariant chain. Let $J_\Gamma = \mathcal{N}(I_\circ) = \cup_{n \geq 1} I_{A_n}$. Our goal is to show that the chain I_\circ stabilizes.

Consider the following decomposable simplicial complex:

$$\Gamma' = \{([m] \setminus T) \cup \{t\} : t \in T\} \cup 2^{[m] \setminus T}.$$

For each $r \in \mathbb{P}$, let I_{B_r} be the toric ideal $I_{\Gamma'}$ which is in the ring $\mathbb{K}[Q_r]$. The I_{B_r} form a chain with respect to the filtration of $\mathbb{K}[Q] * \Pi$. Since Γ' is decomposable, this chain stabilizes by Theorem 4.9. Let $J_{\Gamma'} \subseteq \mathbb{K}[Q]$ denote the union of this chain. Since T is an independent set of Γ , we have $\Gamma \subseteq \Gamma'$, which implies $I_{B_r} \subseteq I_{A_r}$ by Lemma 4.8. We now want to apply Corollary 3.11 and Proposition 3.9 to deduce that the chain I_\circ stabilizes.

For each $r \in P$, the ideal I_{B_r} is a toric ideal, and hence $\mathbb{K}[Q_r]/I_{B_r}$ is a semigroup ring. The limiting ring $\mathbb{K}[Q]/J_{\Gamma'}$ is also a semigroup ring, and it is generated by all monomials appearing in the ring homomorphisms ϕ associated to the matrices B_r . This can be explicitly obtained by looking at the effect of the linear transformation $\pi_{\Gamma'}$ on standard unit vectors. Let $S = [m] \setminus T$. Then,

$$\pi_{\Gamma'}(e_i) = \bigoplus_{t \in T} e_{i_{S \cup \{t\}}} \in \bigoplus_{t \in T} \mathbb{R}^{R_S \times [r]}.$$

For each $F \in \Gamma'$ and $\mathbf{j} \in R_S \times [r]$ we have a variable $y_{\mathbf{j}}^F$. The formula for $\pi_{\Gamma'}$ implies that for each $\mathbf{i} \in R$,

$$\phi(x_{\mathbf{i}}) = \prod_{t \in T} y_{\mathbf{i}, i_t}^{S \cup \{t\}}.$$

This implies that

$$\mathbb{K}[Q]/J_{\Gamma'} =: \mathbb{K}[Q'] \cong \mathbb{K} \left[\prod_{t \in T} y_{\mathbf{i}, t, j_t} : \mathbf{i} \in R_S, j_t \in \mathbb{P} \right].$$

In particular, $\mathbb{K}[Q']$ is a subsemigroup ring of $\mathbb{K}[Y_{R_{[m] \setminus T} \times T \times \mathbb{P}}]$. (This is obtained by replacing the cumbersome superscript $S \cup \{t\}$ with a simple t subscript.)

We now show that $\mathbb{K}[Q']$ is a divisible subsemigroup ring of $\mathbb{K}[Y_{R_{[m] \setminus T} \times T \times \mathbb{P}}]$. Consider the \mathbb{K} -algebra homomorphism ψ from $\mathbb{K}[Y_{R_{[m] \setminus T} \times T \times \mathbb{P}}]$ to $\mathbb{K}[Z_{R_{[m] \setminus T}}]$ that maps $y_{\mathbf{i}, t, j_t}$ to $z_{\mathbf{i}}$. A monomial $y^\alpha \in \mathbb{K}[Y_{R_{[m] \setminus T} \times T \times \mathbb{P}}]$ belongs to $\mathbb{K}[Q']$ if and only if $\psi(y^\alpha)$ is of the form $(z^\beta)^{\#T}$ for some β . Now, if $\psi(y^{\alpha^1}) = (z^{\beta^1})^{\#T}$, $\psi(y^{\alpha^2}) = (z^{\beta^2})^{\#T}$, and $y^{\alpha^1} | y^{\alpha^2}$ then $\psi(y^{\alpha^2 - \alpha^1}) = (z^{\beta^2 - \beta^1})^{\#T}$. This implies that $\mathbb{K}[Q']$ is a divisible subsemigroup of $\mathbb{K}[Y_{R_{[m] \setminus T} \times T \times \mathbb{P}}]$.

Letting Π act on \mathbb{P} and since $R_{[m]\setminus T} \times T$ is a finite set, we have that Π -divisibility on $\mathbb{K}[Y_{R_{[m]\setminus T} \times T \times \mathbb{P}}]$ is a well-partial-ordering. Then, Corollary 3.11 implies that Π -divisibility is also a well-partial-ordering on $\mathbb{K}[Q']$.

Consider the filtration on $\mathbb{K}[Y_{R_{[m]\setminus T} \times T \times \mathbb{P}}] * \Pi$ using $\mathbb{K}[\tilde{Q}_r] = \mathbb{K}[Y_{R_{[m]\setminus T} \times T \times [r]}]$ with $\Pi_{n,m}$. Let $Q'_r = \tilde{Q}_r \cap Q'$ be the induced filtration in $\mathbb{K}[Q'] * \Pi$. Corollary 3.11 also implies that chains with respect to this filtration stabilize. In particular, the chain I_{A_r}/I_{B_r} stabilizes. That is, there is an $r_0 \geq 2$ and a finite generating set \mathcal{F} of $I_{A_{r_0}}/I_{B_{r_0}}$ such that $\Pi_{r,r_0}\mathcal{F}$ generates I_{A_r}/I_{B_r} for all $r > r_0$. Since $r_0 \geq 2$, Theorem 4.9 implies $I_{B_{r_0}}$ is generated by quadrics which also generate all I_{B_r} up to the action of Π . Finally, Proposition 3.9 implies that $\langle \Pi_{r_0,r} I_{A_{r_0}} \rangle_{\mathbb{K}[Q_r]} = I_{A_r}$, which proves the theorem. \square

Example 4.10 (6-cycle). The six cycle $\Gamma = \{\{1, 2\}, \{2, 3\}, \{3, 4\}, \{4, 5\}, \{5, 6\}, \{1, 6\}\}$ has the independent set $T = \{2, 4, 6\}$. If we fix r_1, r_3, r_5 and send $r_2, r_4, r_6 \rightarrow \infty$, then there will be a finite Markov basis for A_Γ up to the natural action of the symmetric group. \square

Theorems 4.4 and 4.7 are finiteness results for Markov bases, but it is also natural to extend these ideas to other statistical situations. Indeed, the Markov bases under consideration are useful tools for studying hierarchical models. As sets, these models are families of probability distributions inside the probability simplex

$$\Delta_R = \left\{ p \in \mathbb{R}^R : \sum_{i \in R} p_i = 1, p_i \geq 0, i \in R \right\}.$$

Each point $p \in \Delta_R$ is a probability distribution for an m -dimensional discrete random vector $Y = (Y_1, \dots, Y_m)$ with state space equal to R . The i th coordinate is the probability of the event $Y = i$, and $p_i = P(Y = i)$.

The hierarchical model \mathcal{M}_Γ is defined as the set $\mathcal{M}_\Gamma = V(I_{A_\Gamma}) \cap \Delta_R$ of solutions to the toric ideal I_{A_Γ} inside the probability simplex. Turning this around, the homogeneous vanishing ideal $\mathcal{I}^h(\mathcal{M}_\Gamma) = I_{A_\Gamma}$ encodes an implicit description of the model that is finite up to symmetry as the number of states of some of the random variables go to infinity.

Using reasoning similar to that found in the preceding proofs, one can deduce finiteness for the implicit representations of families of statistical models as the number of states of some of the variables tend to infinity. We give brief proofs, which follow the same outlines as those of Theorems 4.4 and 4.7.

Theorem 4.11. *For each $r \in \mathbb{P}$, let $\mathcal{M}_r \subseteq \Delta_{R \times [r]}$, where $R = \prod_{i=1}^m [r_i]$, be a statistical model for $(m+1)$ -dimensional discrete random vectors. Suppose that each homogeneous vanishing ideal $I_r = \mathcal{I}^h(\mathcal{M}_r) \subseteq \mathbb{K}[X_{R \times [r]}]$ is stable under the action of \mathfrak{S}_r , and that for each r , we have $I_r \subseteq I_{r+1}$. Then, up to symmetry there is a finite set of polynomials that generates the ideals I_r for all r .*

Proof. The sequence of ideals I_1, I_2, \dots forms an ascending invariant chain. Therefore, by Corollary 3.7 they have a finite generating set up to the filtration of $\mathfrak{S}(\mathbb{P})$ by the \mathfrak{S}_{r_m} . \square

For each $r \in \mathbb{P}^n$, let $\mathcal{M}_r \subseteq \Delta_{\mathcal{S} \times R}$, where $\mathcal{S} = [s_1] \times \dots \times [s_m]$ and $R = [r_1] \times \dots \times [r_n]$, be a statistical model for an $(m+n)$ -dimensional discrete random vector $(Y, Z) = (Y_1, \dots, Y_m, Z_1, \dots, Z_n)$. Suppose that each homogeneous vanishing ideal $I_r = \mathcal{I}^h(\mathcal{M}_r) \subseteq$

$\mathbb{R}[X_{\mathcal{S} \times \mathcal{R}}]$ is stable under the action of $\mathfrak{S}_{r_1} \times \cdots \times \mathfrak{S}_{r_n}$ and assume that for each $r \in \mathbb{P}^n$ and $t \in \mathbb{N}^n$, we have $I_r \subseteq I_{r+t}$. To generalize Theorem 4.7 to arbitrary statistical models, we need to explain what should be meant by the condition that a collection of vertices forms an independent set. The simplest (algebraic) way to guarantee such a generalization is to require that for each r , we have $I_{B_r} \subseteq I_r$, where I_{B_r} is the toric ideal of the hierarchical model whose simplicial complex Γ has facets $\{[m] \cup \{i'\} : i' \in \{1', 2', \dots, n'\}\}$. Note that this is the same ideal appearing in the proof of Theorem 4.7.

In more statistical language, the condition $I_{B_r} \subseteq I_r$ for all r is equivalent to the random vector (Y, Z) satisfying the conditional independence statement $Z_1 \perp\!\!\!\perp Z_2 \perp\!\!\!\perp \cdots \perp\!\!\!\perp Z_n | Y$ (see Chapter 3 in [12] for connections between conditional independence and hierarchical/graphical models). We state our result using the language of conditional independence.

Theorem 4.12. *For each $r \in \mathbb{P}^n$, let $\mathcal{M}_r \subseteq \Delta_{\mathcal{S} \times \mathcal{R}}$, where $\mathcal{S} = [s_1] \times \cdots \times [s_m]$ and $\mathcal{R} = [r_1] \times \cdots \times [r_n]$, be a statistical model for an $(m+n)$ -dimensional discrete random vector $(Y, Z) = (Y_1, \dots, Y_m, Z_1, \dots, Z_n)$. Suppose that each homogeneous vanishing ideal $I_r = \mathcal{I}^h(\mathcal{M}_r) \subseteq \mathbb{R}[X_{\mathcal{S} \times \mathcal{R}}]$ is stable under the action of $\mathfrak{S}_{r_1} \times \cdots \times \mathfrak{S}_{r_n}$ and assume that for each $r \in \mathbb{P}^n$ and $t \in \mathbb{N}^n$, we have $I_r \subseteq I_{r+t}$. If, in addition, the \mathcal{M}_r all satisfy the conditional independence constraint $Z_1 \perp\!\!\!\perp Z_2 \perp\!\!\!\perp \cdots \perp\!\!\!\perp Z_n | Y$, then up to symmetry there is a finite set of polynomials that generates the ideals I_r for all r .*

Proof. The key feature of this theorem is the conditional independence constraint

$$Z_1 \perp\!\!\!\perp Z_2 \perp\!\!\!\perp \cdots \perp\!\!\!\perp Z_n | Y.$$

Let Γ' be the simplicial complex with facets $[m] \cup \{i'\}$ such that $i' \in \{1', 2', \dots, n'\}$; this is the decomposable simplicial complex that appeared in the proof of Theorem 4.7. The conditional independence statement holding for the model \mathcal{M}_r is equivalent to $I_{B_r} \subseteq I_r$ (see Chapter 3 in [12]). The remainder of the proof now follows closely that of Theorem 4.7. \square

5. FURTHER DIRECTIONS

From the standpoint of computational algebra, we have proved theorems asserting the existence of finite generating sets of ideals up to symmetry. Many open problems remain about how to transition from these existence theorems to effective versions and, in particular, how to develop specific algorithms for computing with symmetric ideals. We outline some of these challenges here.

Many chains of ideals in algebraic statistics arise as kernels of ring homomorphisms. Besides knowing that these chains eventually stabilize and have finite generating sets, one desires upper bounds on when stabilization occurs in terms of the input data. To be more precise, for each $r \in \mathbb{P}$, let $\phi_r : \mathbb{K}[X_{[k] \times [r]}] \rightarrow R$ be a ring homomorphism and let $I_r = \ker \phi_r$. Suppose that each I_r is invariant under the action of \mathfrak{S}_r and that this sequence of kernels is nested: $I_r \subseteq I_{r+1}$. We call such a chain a *chain of kernels*.

Question 5.1. Given a chain of kernels I_\circ , find upper bounds on n_0 such that

$$\langle \mathfrak{S}_n I_{n_0} \rangle_{\mathbb{K}[X_{[k] \times [n]}]} = I_n \quad \text{for all } n > n_0$$

in terms of the ring homomorphisms ϕ_r . Of special interest is when each I_r is a toric ideal, in which case $\phi_r = \phi_A$ for a integral matrix A .

In Section 3, we showed that Π -invariant divisible semigroup rings $\mathbb{K}[Q]$ that are subrings of $\mathbb{K}[X_{[k] \times \mathbb{P}}]$ are Noetherian $\mathbb{K}[Q] * \Pi$ modules. A natural question is to what extent this property generalizes.

Question 5.2. Let $\mathbb{K}[Q] \subseteq \mathbb{K}[X_{[k] \times \mathbb{P}}]$ be a Π -invariant semigroup ring which is finitely generated under the action of Π . Is it true that $\mathbb{K}[Q]$ is a Noetherian $\mathbb{K}[Q] * \Pi$ -module?

Alexei Krasilnikov constructed a remarkably simple example which shows that the answer to Question 5.2 is “no”.

Example 5.3 (Krasilnikov). Let $k = 2$ and let $\mathbb{K}[Q] \subseteq \mathbb{K}[X_{[2] \times \mathbb{P}}]$ be the semigroup generated by the monomials $x_{1i}x_{2j}$ where $i < j$. Note that this semigroup ring is finitely generated up to the action of Π by the single monomial $x_{11}x_{22}$.

For $n \geq 3$ define the element $w_n \in Q$ by

$$w_n = x_{11}x_{2n} \prod_{i=1}^{n-1} x_{1i}x_{2i+1}.$$

Consider the the multigrading on the ring $\mathbb{K}[Q]$ defined by $\deg x_{ij} = e_j \in \mathbb{N}^{\mathbb{N}}$. In particular, the $\deg w_n = (2, 2, \dots, 2, 0, 0, \dots)$. Suppose that some $w_m |_{\Pi} w_n$. Then there is a $p \in \Pi$ such that $pw_m = hw_n$. Now, $\deg pw_m \in \{0, 2\}^{\mathbb{N}}$ so $\deg h \in \{0, 2\}^{\mathbb{N}}$ as well. Examining at the right-most nonzero entry in $\deg h$, we see that $x_{2k_1}^2 | h$ for some k_1 . Also, the right-most nonzero entry in $\deg pw_m$ implies that $x_{2k_2}^2 | pw_m$ for some k_2 . This implies that $x_{2k_1}^2 x_{2k_2}^2$ divides w_n which is impossible. Hence, the sequence w_3, w_4, \dots is a bad sequence and by Proposition 2.11 and Theorem 2.12, $\mathbb{K}[Q]$ is not a Noetherian $\mathbb{K}[Q] * \Pi$ -module. \square

While we have been mainly interested in ideals that are invariant under the action of the symmetric group, we needed to restrict to actions of the monoid of increasing functions Π in order to prove our finiteness theorems. We are lead to wonder if this strategy could always be used to prove Noetherianity under symmetric group actions or if there might be some pathological counterexamples or obstructions.

In particular, let R be a ring equipped with an $\mathfrak{S}_{\mathbb{P}}$ action. We say that this action is $\mathfrak{S}_{\mathbb{P}}$ -finite if for every $f \in R$ there is an $m \in \mathbb{P}$ such that $\sigma \cdot f = \sigma' \cdot f$ for all $\sigma, \sigma' \in \mathfrak{S}_{\mathbb{P}}$ such that $\sigma(i) = \sigma'(i)$ for all $i \leq m$. If R has a $\mathfrak{S}_{\mathbb{P}}$ -finite action, it also has a natural action by the monoid of increasing functions Π .

Question 5.4. Is there a ring R with a $\mathfrak{S}_{\mathbb{P}}$ -finite action such that R is a Noetherian $R * \mathfrak{S}_{\mathbb{P}}$ -module but not a Noetherian $R * \Pi$ -module?

One of the lessons we have learned about proving Noetherianity of $\mathbb{K}[X_{[k] \times \mathbb{P}}]$ as a $\mathbb{K}[X_{[k] \times \mathbb{P}}] * \mathfrak{S}_{\mathbb{P}}$ -module is that it is not possible to define Gröbner bases in this setting. This suggests that an approach for computing with ideals that have a natural symmetry group using Gröbner bases might not work well if the entire symmetry group is used. However, working with a semigroup that has a P -order might be natural and useful, and not require the bookkeeping of a full symmetry group. This suggests the following implementation challenge.

Problem 5.5. Develop and implement algorithms for computing with symmetric ideals by using monoids of transformations and P -orders.

For some recent work along these lines, including an algorithm for computing with certain classes of invariant ideals, we refer the reader to [3].

ACKNOWLEDGMENTS

We thank Alexei Krasilnikov for providing us with references to the work of Cohen and for the use of Example 5.3.

REFERENCES

- [1] S. Aoki and A. Takemura. Minimal basis for connected Markov chain over $3 \times 3 \times K$ contingency tables with fixed two dimensional marginals. *Australian and New Zealand Journal of Statistics* **45** (2003) 229–249.
- [2] M. Aschenbrenner and C. Hillar. Finite generation of symmetric ideals, *Trans. Amer. Math. Soc.*, **359** (2007), 5171–5192; erratum, *ibid.* **361** (2009), 5627–5627.
- [3] A. E. Brouwer and J. Draisma, Equivariant Gröbner bases and the Gaussian two-factor model, *Math. Comp.* **80** (2011), 1123–1133.
- [4] D.E. Cohen, *On the laws of a metabelian variety*, *J. Algebra* **5** (1967), 267–273.
- [5] D.E. Cohen, *Closure relations, Buchberger’s algorithm, and polynomials in infinitely many variables*, *Computation theory and logic*, (1987), 78–87.
- [6] D. Cox, J. Little, and D. O’Shea, *Ideals, Varieties, and Algorithms*, Springer-Verlag, New York, 1997.
- [7] P. Diaconis and B. Sturmfels. Algebraic algorithms for sampling from conditional distributions. *Ann. Statist.* **26** (1998), no.1, 363–397.
- [8] A. Dobra. Markov bases for decomposable graphical models. *Bernoulli* **9** No. 6 (2003), 1093–1108.
- [9] J. Draisma. Finiteness for the k-factor model and chirality varieties, *Adv. Math.* **223** (2010), 243–256.
- [10] V. Drensky and R. La Scala, *Gröbner bases of ideals invariant under endomorphisms*, *J. Symbolic Comput.* **41** (2006), 835–846.
- [11] M. Drton, B. Sturmfels, and S. Sullivant, Algebraic factor analysis: Tetrads, pentads and beyond, *Probability Theory and Related Fields* **138** (2007) 463–493.
- [12] M. Drton, B. Sturmfels, and S. Sullivant. *Lectures on Algebraic Statistics*, Oberwolfach Seminar Series **39**, Birkhauser, 2009.
- [13] G. Higman, Ordering by divisibility in abstract algebras, *Proc. London Math. Soc.* (3) **2** (1952), 326–336.
- [14] S. Hoşten and S. Sullivant. A finiteness theorem for Markov bases of hierarchical models. *J. Combin. Theory Ser. A* **114** (2007), no. 2, 311–321.
- [15] A. Kemer, *Analog of Hilbert basis theorem for infinitely generated commutative algebras*, *Asian-European J. Math.* **1** (2008), 555–564.
- [16] J. B. Kruskal. The theory of well-quasi-ordering: A frequently discovered concept, *J. Combinatorial Theory Ser. A* **13** (1972), 297–305.
- [17] C. St. J. A. Nash-Williams. On well-quasi-ordering finite trees, *Proc. Cambridge Philos. Soc.* **59** (1963), 833–835.
- [18] E. Ruch, A. Schönhofer, and I. Ugi. Die Vandermondesche Determinante als Näherungsansatz für eine Chiralitätsbeobachtung, ihre Verwendung in der Stereochemie und zur Berechnung der optischen Aktivität, *Theor. Chim. Acta* **7** (1967), 420–432.
- [19] F. Santos and B. Sturmfels. Higher Lawrence Configurations. *J. Combin. Theory Ser. A*, **103** (2003) 151–164.
- [20] B. Sturmfels. Gröbner bases and Stanley decompositions of determinantal rings, *Mathematische Zeitschrift* **205** (1990) 137–144.
- [21] A. Takken. Monte Carlo goodness of fit tests for discrete data. Ph.D. Thesis, Stanford University, 2000.

MATHEMATICAL SCIENCES RESEARCH INSTITUTE, 17 GAUSS WAY, BERKELEY, CA 94720
E-mail address: `chillar@msri.org`

DEPARTMENT OF MATHEMATICS, NORTH CAROLINA STATE UNIVERSITY, RALEIGH, NC
E-mail address: `smsulli2@ncsu.edu`