This is the first author's version of a work that was accepted for publication in the Journal of Algebra. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. A definitive version was subsequently published in
J. Algebra **306** (2006) 175–90.
http://dx.doi.org/10.1016/j.jalgebra.2005.11.040

p-restriction of partitions and homomorphisms between Specht modules

Matthew Fayers
Queen Mary, University of London, Mile End Road, London E1 4NS, U.K.

Sinéad Lyle
School of Mathematics and Statistics F07, University of Sydney, NSW 2006, Australia

Stuart Martin
DPMMS, Wilberforce Road, Cambridge CB3 0WB, U.K.

2000 Mathematics subject classification: 20C30

Abstract

Let λ be a partition of n. We show that the space of $\mathbb{F}\mathfrak{S}_n$-homomorphisms between the Specht modules S^λ and $S^{\lambda^{\mathrm{r}}}$ is one-dimensional, where $\mathbb{F}$ is a field of characteristic p and λ^{r} is the 'p-restriction' of λ. Equivalently, our result proves the corresponding theorem for the homomorphism space $\mathrm{Hom}_{\mathbb{F}\mathfrak{S}_n}(S^{\lambda^{\mathrm{reg}}}, S^\lambda)$, where λ^{reg} is the 'p-regularisation' of λ, as defined by James.

1 Introduction

Let n be a non-negative integer, let $\mathfrak{S}_n$ denote the symmetric group on n letters and let $\mathbb{F}$ be a field of characteristic $p > 0$. For each partition λ of n, one defines a *Specht module* S^λ for the group algebra $\mathbb{F}\mathfrak{S}_n$. When λ is p-restricted, S^λ has a simple socle D_λ, and $\{D_\lambda \mid \lambda \text{ is a } p\text{-restricted partition of } n\}$ is a complete set of non-isomorphic, irreducible $\mathbb{F}\mathfrak{S}_n$-modules.

The main problem in the modular representation theory of $\mathfrak{S}_n$ is the calculation of the decomposition numbers, i.e. the composition multiplicities of the simple modules D_μ in the Specht modules S^λ, for λ and μ partitions of n with μ p-restricted. Many results concerning this problem have been proved, but it remains very difficult in general. One of the earliest such results was proved by James in [3]; he found, for each partition λ, the most dominant p-restricted partition μ such that $[S^\lambda : D_\mu] > 0$, and showed moreover that this decomposition number equals 1. The partition μ is constructed in a combinatorial way from λ, via a process we call 'p-restriction'.

A problem of similar interest and difficulty to the decomposition number problem is the determination of the homomorphism space $\mathrm{Hom}_{\mathbb{F}\mathfrak{S}_n}(S^\lambda, S^\mu)$ for partitions λ and μ. In this paper, we solve this problem in the case where μ is the p-restriction of λ, showing that the homomorphism space is one-dimensional. Our methods are elementary, involving manipulation of Young tableaux.

As with many results in the representation theory of the symmetric groups, James's decomposition number result generalises to the representation theory of the Iwahori–Hecke algebra $\mathcal{H}_{\mathbb{F},q}(\mathfrak{S}_n)$; the notion of p-restriction must be replaced with e-restriction, where e is the multiplicative order of q. It is tempting to speculate that our result also carries over to this setting, but this seems difficult to prove using our methods.

1.1 Background and notation

We recall the basic representation theory of the symmetric group from James's book [4]. In particular, we use the notions of composition, partition, tableau and row and column equivalence discussed there, as well as the permutation module M^λ and the Specht module S^λ. We may abuse notation by identifying a partition with its Young diagram. Throughout, we work over a field $\mathbb{F}$ of prime characteristic p.

1.1.1 Homomorphisms from Specht modules to permutation modules

We now recall the results we shall need on homomorphisms; most of this is taken from [4], but we shall also need some results from [2].

If λ is a composition of n, then we write t^λ for the λ-tableau (of type (1^n)) formed by writing the integers $1, \dots, n$ along successive rows. If μ is another composition of n, then we define a bijection from the set of μ-tabloids to the set $\mathcal{T}(\lambda, \mu)$ of λ-tableaux of type μ: given a μ-tabloid $\{\mathfrak{s}\}$, we define the corresponding λ-tableau S by

$$S(x, y) = \text{the number of the row in which } \mathrm{t}^\lambda(x, y) \text{ appears in } \{\mathfrak{s}\}.$$

Using this bijection, we regard $\mathcal{T}(\lambda, \mu)$ as a basis for M^μ.

Now suppose we have a λ-tableau T of type μ. We define a homomorphism $\Theta_T : M^\lambda \to M^\mu$ by specifying

$$\Theta_T : \{\mathrm{t}^\lambda\} \longmapsto \sum_{S \sim_{\mathrm{row}} T} S,$$

and extending homomorphically. The restriction of Θ_T to the Specht module S^λ is written as $\hat{\Theta}_T$. The maps $\hat{\Theta}_T$ turn out to be very useful; our main theorem will be proved by explicitly constructing a tableau T and using the corresponding homomorphism.

For $d \geqslant 1$ and $0 \leqslant t < \lambda_{d+1}$, there is also a homomorphism $\psi_{d,t}$ from M^λ to a permutation module M^ν, whose importance is illustrated by the following theorem.

Theorem 1.1. The Kernel Intersection Theorem [4, Corollary 17.18] *If λ is a partition of n, then*

$$S^\lambda = \bigcap_{d \geqslant 1} \bigcap_{t=0}^{\lambda_{d+1}-1} \ker \psi_{d,t}.$$

The Kernel Intersection Theorem is very useful in finding homomorphisms between Specht modules. If λ and μ are partitions of n, and if θ is a homomorphism from S^λ to M^μ, then the image of θ lies inside the Specht module S^μ if and only if $\psi_{d,t} \circ \theta = 0$ for all d, t. We shall make use of this observation in Section 5.

Now we cite some results from [2] concerning basic manipulation of homomorphisms. Given any tableau T, we write T_i^j for the number of entries equal to i in row j.

Lemma 1.2. [2, Lemma 7] *Suppose λ is a partition of n and μ a composition, and that $T \in \mathcal{T}(\lambda, \mu)$ is row standard. Suppose $r < s$, and that i is an integer appearing a times in row s of T. Let $\mathcal{V}(T)$ be the set of row standard tableaux which may be obtained from T by interchanging the entries equal to i in row s with some a entries not equal to i in row r, and re-ordering the entries in each row. Then*

$$\hat{\Theta}_T = (-1)^a \sum_{V \in \mathcal{V}(T)} \prod_{j \geqslant 1} \binom{V^s_j}{T^s_j} \hat{\Theta}_V.$$

Lemma 1.3. *Suppose T is a λ-tableau of type μ, and let U be the tableau formed from rows $r+1, r+2, \dots$ of T, for some r. If $\hat{\Theta}_U = 0$, then $\hat{\Theta}_T = 0$.*

Proof. This is a special case of [2, Lemma 4]. □

1.1.2 The process of p-restriction

In this section, we describe the process of p-restriction of partitions which motivates our main theorem. We use the prime $p = \mathrm{char}(\mathbb{F})$, although the notion of p-restriction works for any integer $p > 1$. For an integer $i \geqslant 1$, the ith *ramp* in $\mathbb{N} \times \mathbb{N}$ is defined to be

$$\{(x, y) \mid (p-1)x + y = i + p - 1\}.$$

If λ is a partition, then the ith ramp of λ is the intersection of this ramp with the Young diagram of λ. We say that a ramp is *full* if every node of that ramp is a node of λ. The *p-restriction* of λ, denoted λ^{r}, is defined to be the partition whose Young diagram is obtained by moving all the nodes of λ as far down their ramps as they will go. It is a fairly easy exercise to show that we actually obtain the Young diagram of a p-restricted partition by this procedure. For example, if $p = 3$ and $\lambda = (9, 7, 3, 3)$, then $\lambda^{\mathrm{r}} = (7, 6, 5, 3, 1)$. We may see this by comparing the Young diagrams of these partitions; for each node we write the number of the ramp in which it appears:

1	2	3	4	5	6	7	8	9
3	4	5	6	7	8	9		
5	6	7						
7	8	9						

1	2	3	4	5	6	7
3	4	5	6	7	8	
5	6	7	8	9		
7	8	9				
9						

.

James [3, Theorem A] showed that every row of the decomposition matrix for the symmetric group $\mathfrak{S}_n$ in characteristic p contains a 1; his result may be stated as follows.

Theorem 1.4. *Suppose λ and μ are partitions of n, with μ p-restricted. Then $[S^\lambda : D_{\lambda^{\mathrm{r}}}] = 1$, while $[S^\lambda : D_\mu] = 0$ unless $\lambda^{\mathrm{r}} \trianglerighteq \mu$.*

A quicker proof of James's result appears in [5]. Our main theorem may be regarded as an analogue of this theorem for homomorphisms between Specht modules.

Theorem 1.5. *If λ is a partition of n, then*

$$\dim_{\mathbb{F}}(\mathrm{Hom}_{\mathbb{F}\mathfrak{S}_n}(S^\lambda, S^{\lambda^{\mathrm{r}}})) = 1.$$

To prove that this homomorphism space has dimension at most 1 is easy.

Proposition 1.6.

$$\dim_{\mathbb{F}}(\mathrm{Hom}_{\mathbb{F}\mathfrak{S}_n}(S^\lambda, S^{\lambda^{\mathrm{r}}})) \leqslant 1.$$

Proof. Since the socle of $S^{\lambda^{\mathrm{r}}}$ is a simple module $D_{\lambda^{\mathrm{r}}}$ which appears exactly once as a composition factor of S^λ, the coimage of any non-zero homomorphism from S^λ to $S^{\lambda^{\mathrm{r}}}$ must be the unique quotient M of S^λ having $D_{\lambda^{\mathrm{r}}}$ as its socle. So

$$\begin{aligned}\dim_{\mathbb{F}}(\mathrm{Hom}_{\mathbb{F}\mathfrak{S}_n}(S^\lambda, S^{\lambda^{\mathrm{r}}})) &= \dim_{\mathbb{F}}(\mathrm{Hom}_{\mathbb{F}\mathfrak{S}_n}(M, S^{\lambda^{\mathrm{r}}}))\\ &\leqslant \dim_{\mathbb{F}}(\mathrm{Hom}_{\mathbb{F}\mathfrak{S}_n}(\mathrm{soc}(M), \mathrm{soc}(S^{\lambda^{\mathrm{r}}})))\\ &= \dim_{\mathbb{F}}(\mathrm{Hom}_{\mathbb{F}\mathfrak{S}_n}(D^{\lambda^{\mathrm{r}}}, D^{\lambda^{\mathrm{r}}}))\\ &= 1,\end{aligned}$$

since every field is a splitting field for $\mathfrak{S}_n$. □

In order to prove Theorem 1.5, therefore, it suffices to find a non-zero homomorphism from S^λ to $S^{\lambda^{\mathrm{r}}}$. This is done in Theorem 2.2.

In some instances, our homomorphism occurs as the composition of 'known' homomorphisms between Specht modules. For example, suppose $p = 3$ and $\lambda = (5,3)$, so that $\lambda^{\mathrm{r}} = (4,3,1)$. Then there are non-zero homomorphisms

$$\hat{\Theta}_1 : S^{(5,3)} \to S^{(5,2,1)}, \qquad \hat{\Theta}_2 : S^{(5,2,1)} \to S^{(4,3,1)}$$

(these are 'one-node Carter–Payne homomorphisms' [1]), and it is easy to check (using the more explicit construction in [2]) that the composition $\hat{\Theta}_2 \circ \hat{\Theta}_1$ is non-zero. Similarly, Koppinen [6] (working in an algebraic groups setting) described certain pairs of partitions for which the homomorphism space between the corresponding Weyl modules is non-zero; this gives the existence of homomorphisms between the corresponding Specht modules, which can be used to construct homomorphisms $S^\lambda \to S^{\lambda^{\mathrm{r}}}$ in certain cases.

In general, however, our homomorphism cannot be constructed in this way. Consider the case where $p = 3$ and $\lambda = (6)$, so that $\lambda^{\mathrm{r}} = (2^3)$. Then λ and λ^{r} are not close (in the sense of Koppinen), nor is there a Carter–Payne homomorphism between them. Nor can our homomorphism be written as a composition of homomorphisms between Specht modules: the only partition μ such that $(6) \rhd \mu \rhd (2^3)$ and $\mathrm{Hom}_{\mathbb{F}\mathfrak{S}_6}(S^{(6)}, S^\mu) \neq 0$ is $\mu = (5,1)$, and it is easy to check (by writing out all possible maps in terms of semistandard homomorphisms, or otherwise) that $\mathrm{Hom}_{\mathbb{F}\mathfrak{S}_6}(S^{(5,1)}, S^{(2^3)}) = 0$.

1.1.3 *p*-regularisation

Before proceeding with the proof of Theorem 1.5, we briefly discuss p-regularisation of partitions. Write λ' for the partition conjugate to λ, and define the *p-regularisation* λ^{reg} of λ to be $(\lambda')^{\mathrm{r}\prime}$. p-regularisation is perhaps a more familiar concept than p-restriction, and is more appropriate when using James's parameterisation $\{D^\lambda \mid \lambda \text{ a } p\text{-regular partition}\}$ of the irreducible $\mathbb{F}\mathfrak{S}_n$-modules; indeed, Theorem 1.4 was originally stated in these terms. An equivalent version of our main theorem, stated in terms of p-regularisation, is as follows.

Theorem 1.7. *If λ is a partition of n, then*

$$\dim_{\mathbb{F}}(\mathrm{Hom}_{\mathbb{F}\mathfrak{S}_n}(S^{\lambda^{\mathrm{reg}}}, S^{\lambda})) = 1.$$

That this is equivalent to Theorem 1.5 is easy to see using the fact [4, Theorem 8.15] that $S^{\lambda'} \cong (S^{\lambda})^* \otimes \mathrm{sgn}$, where sgn denotes the signature representation of $\mathfrak{S}_n$. From this it follows that for any λ and μ we have

$$\dim_{\mathbb{F}}(\mathrm{Hom}_{\mathbb{F}\mathfrak{S}_n}(S^{\lambda}, S^{\mu})) = \dim_{\mathbb{F}}(\mathrm{Hom}_{\mathbb{F}\mathfrak{S}_n}(S^{\mu'}, S^{\lambda'})),$$

which immediately gives the equivalence of Theorems 1.5 and 1.7. We have chosen to work with p-restriction in this paper simply because it is easier to construct homomorphisms from S^{λ} to $S^{\lambda^{\mathrm{r}}}$.

1.1.4 Miscellaneous notation

- We frequently use row and column removal operations on partitions, and we use $\bar{\ }$ and $^{|}$ to denote these: so if $\lambda = (\lambda_1, \lambda_2, \dots)$, then

  $$\bar{\lambda} = (\lambda_2, \lambda_3, \dots)$$

  and

  $$^{|}\lambda = (\max(\lambda_1 - 1, 0), \max(\lambda_2 - 1, 0), \dots).$$

- We write $\mathbb{1}(S)$ for the indicator function of the truth of a statement S.

- We use a circumflex accent to denote the omission of an item from a list.

2 Magic tableaux

The advantage of working with the partitions λ and λ^{r} is that we shall be able to express our homomorphism $S^{\lambda} \to S^{\lambda^{\mathrm{r}}}$ in terms of a single λ-tableau (of type $S^{\lambda^{\mathrm{r}}}$), which we shall call a *magic tableau*. Our construction begins with the following lemma.

Lemma 2.1. *Suppose $\lambda = (\lambda_1, \lambda_2, \dots)$ is a partition, and j is a positive integer with $\lambda^{\mathrm{r}}_j > 0$. The following are equivalent.*

1.
$$(\bar{\lambda})^{\mathrm{r}}_i = \begin{cases} \lambda^{\mathrm{r}}_i - p + 1 & (i < j) \\ \lambda^{\mathrm{r}}_{i+1} & (i \geqslant j). \end{cases}$$

2.
$$\lambda^{\mathrm{r}}_j + (p-1)(j-1) = \lambda_1.$$

Furthermore, there exists at least one j for which these conditions hold.

We call a j for which the conditions of Lemma 2.1 hold a *magic value*.

Proof. (1)⇒(2) is easy, since $|\lambda^{\mathrm{r}}| = |\bar{\lambda}^{\mathrm{r}}| + \lambda_1$. So we suppose that (2) holds.

We write $\mathrm{mis}^{\lambda}(l)$ for the number of nodes in ramp l which are not nodes of λ. Then we have

$$\lambda^{\mathrm{r}}_i = |\{l > (i-1)(p-1) \mid \mathrm{mis}^{\lambda}(l) < i\}|$$

while

$$\bar{\lambda}^{\mathrm{r}}_i = |\{l > i(p-1) \mid \mathrm{mis}^{\lambda}(l) < i + \mathbb{1}(l > \lambda_1)\}|.$$

In other words,

$$\bar{\lambda}^{\mathrm{r}}_i = |\{i(p-1) < l \leqslant \lambda_1 \mid \mathrm{mis}^{\lambda}(l) < i\}| + |\{l > \max(\lambda_1, i(p-1)) \mid \mathrm{mis}^{\lambda}(l) < i+1\}|.$$

Suppose first that $i < j$. By (2), the last node in row j of λ^{r} lies in ramp λ_1, which means that $\mathrm{mis}^{\lambda}(l) \geqslant j$ for $l > \lambda_1$. Hence $\mathrm{mis}^{\lambda}(l) \geqslant i+1$ for $l > \lambda_1$, so the second part of the above sum is zero. So we have

$$\bar{\lambda}^{\mathrm{r}}_i = |\{i(p-1) < l \leqslant \lambda_1 \mid \mathrm{mis}^{\lambda}(l) < i\}|.$$

Now if $(i-1)(p-1) < l \leqslant i(p-1)$, then ramp l contains exactly i nodes. If $l \leqslant \lambda_1$ then $\mathrm{mis}^{\lambda}(l) < i$ for $(i-1)(p-1) < l \leqslant i(p-1)$, and we deduce that

$$\begin{aligned}\bar{\lambda}^{\mathrm{r}}_i &= |\{(i-1)(p-1) < l \leqslant \lambda_1 \mid \mathrm{mis}^{\lambda}(l) < i\}| - p + 1\\ &= |\{(i-1)(p-1) < l \mid \mathrm{mis}^{\lambda}(l) < i\}| - p + 1\\ &= \lambda^{\mathrm{r}}_i - p + 1.\end{aligned}$$

Now suppose $i \geqslant j$. Every l with $i(p-1) < l \leqslant \lambda_1$ satisfies $\mathrm{mis}^{\lambda}(l) < j$, and so certainly $\mathrm{mis}^{\lambda}(l) < i$. So

$$\begin{aligned}\bar{\lambda}^{\mathrm{r}}_i &= (\lambda_1 - i(p-1))\mathbb{1}(i(p-1) \leqslant \lambda_1) + |\{l > \max(\lambda_1, i(p-1)) \mid \mathrm{mis}^{\lambda}(l) < i+1\}|\\ &= |\{l > i(p-1) \mid \mathrm{mis}^{\lambda}(l) < i+1\}|\\ &= \lambda^{\mathrm{r}}_{i+1}.\end{aligned}$$

To show that a magic value exists, we must find j such that the last node in row j of λ^{r} lies in ramp λ_1. Consider $j = \mathrm{mis}^{\lambda}(\lambda_1) + 1$. Certainly row j contains a node of ramp λ_1; we must show that it does not contain a node of ramp $\lambda_1 + 1$. Letting $\mathrm{rp}(l)$ denote the number of nodes in ramp l of λ, we must show that $\mathrm{rp}(\lambda_1 + 1) < \mathrm{rp}(\lambda_1) + \mathbb{1}((p-1) \mid \lambda_1)$. But every node in ramp $\lambda_1 + 1$ of λ (except the node in column 1, if $(p-1) \mid \lambda_1$) has a node of ramp λ_1 immediately to the left of it; furthermore the node $(1, \lambda_1)$ in ramp λ_1 does not lie immediately to the left of a node of λ, and the inequality follows. □

Note that the magic value j constructed at the end of the above proof is the smallest magic value, since it corresponds to the row containing the highest node of ramp λ_1 of λ^{r}.

Now suppose we have a sequence $i_1 < i_2 < \dots$ of integers. We define a *magic λ-tableau on* $i_1, i_2, \dots$ to be any λ-tableau obtained using the following recursive procedure:

1. choose some a magic value j for λ;
2. fill in the first row of λ (in increasing order) with $p-1$ entries equal to i_1, $p-1$ entries equal to i_2, and so on up to i_{j-1}, and then λ^{r}_j entries equal to i_j;
3. fill in the remaining rows with a magic $\bar{\lambda}$-tableau on $i_1, i_2, \dots, \widehat{i_j}, \dots$.

Remarks.

1. There are often several choices of j in Lemma 2.1. For example, if $p = 2$ and λ is the 2-core $(r, r-1, \dots, 2, 1)$, then any $j \in \{1, 2, \dots, r\}$ will do. Hence there are usually several different magic tableaux; in fact, if we let $\lambda^{(s)}$ denote the partition $(\lambda_{s+1}, \lambda_{s+2}, \dots)$, then the number of magic λ-tableaux is the product over all s of the number of magic values for $\lambda^{(s)}$. We also note that set of magic values for λ is a set of *consecutive* integers – this follows easily using characterisation (1) of magic values in Lemma 2.1.

2. It is easy to find the content of a magic tableau using Lemma 2.1. A magic λ-tableau on $i_1, i_2, \dots$ contains λ^{r}_k entries equal to i_k, for each k.

We define the *first magic λ-tableau* $T_\lambda = T_\lambda(i_1, i_2, \dots)$ on $i_1, i_2, \dots$ to be the magic tableau in which we choose the smallest possible magic value at each stage. Note that this tableau is straightforward to write down: the smallest possible magic value for λ will be i_j, where j is the highest row of λ^{r} which contains a node in ramp λ_1. Subsequent magic values are chosen similarly.

Example. Take $p = 3$ and $\lambda = (10, 10, 9, 3, 3, 3)$. Then $\lambda^{\mathrm{r}} = (9, 8, 7, 6, 4, 3, 1)$ and the first magic λ-tableau on $1, 2, \dots$ is

1	1	2	2	2	2	2	2	2	2
1	1	3	3	4	4	4	4	4	4
1	1	3	3	5	5	6	6	6	
1	1	1							
3	3	3							
5	5	7							

.

In view of Proposition 1.6, Theorem 1.5 will follow from the following result.

Theorem 2.2. *Suppose λ is a partition, and T is a magic λ-tableau on* $1, 2, \dots$. *Then $\hat{\Theta}_T : S^\lambda \to M^{\lambda^{\mathrm{r}}}$ is a non-zero homomorphism whose image lies inside the Specht module $S^{\lambda^{\mathrm{r}}}$.*

Remark. Given Theorem 2.2, it is therefore easy to write down a non-zero homomorphism $\Theta : S^\lambda \to S^{\lambda^{\mathrm{r}}}$: we take $\Theta = \hat{\Theta}_T$ where T is the first magic tableau on $1, 2, \dots$. As noted above, it is a simple matter to construct T.

2.1 Alternative characterisations of magic tableaux

It will be useful in the proofs in later sections to have two more descriptions of magic tableaux.

Lemma 2.3. *Suppose T is a λ-tableau on $i_1, i_2, \dots$. Define $f : \{1, \dots, \lambda'_1\} \to \{i_1, i_2, \dots\}$ by $k \mapsto T(k, \lambda_k)$. Then T is a magic tableau on $i_1, i_2, \dots$ if and only if the following all hold.*

1. *For each k, the number of entries of T equal to i_k is λ^{r}_k.*

2. *The entries in each row of T are weakly increasing.*

3. *f is injective.*

4. *If $f(r) = i_k$ and $s > r$, then i_k does not appear in row s.*

5. *If $f(r) = i_k$, $s < r$ and $f(s) > i_k$, then i_k appears exactly $p - 1$ times in row s.*
6. *If i_k does not lie in the image of f, then i_k appears exactly $p - 1$ times in row s if $f(s) > i_k$, and does not appear in row s otherwise.*

Proof. Suppose first that T is magic, and that j is the magic value chosen for the first row. Write $\bar{T}$ for the $\bar{\lambda}$-tableau formed by removing the first row of T. Then $\bar{T}$ is magic on $i_1, i_2, \dots, \widehat{i_j}, \dots$, while row 1 has entries $i_1^{p-1}, i_2^{p-1}, \dots, i_{j-1}^{p-1}, i_j^{\lambda^{\mathrm{r}}_j}$; in particular, all the i_js in T occur in the first row. The content of T is correct by our earlier remarks, and Conditions (2–6) follow by induction.

Conversely, suppose that (1–6) hold, and that $f(1) = i_j$. The conditions imply that the first row of T has entries $i_1^{p-1}, i_2^{p-1}, \dots, i_{j-1}^{p-1}, i_j^{\lambda^{\mathrm{r}}_j}$, and that i_j does not appear anywhere else in T. Since the number of i_js in T is λ^{r}_j, this means that $\lambda^{\mathrm{r}}_j + (j-1)(p-1) = \lambda_1$, so that j is a magic value for λ. So T is magic on $i_1, i_2, \dots$ if and only if $\bar{T}$ is magic on $i_1, i_2, \dots, \widehat{i_j}, \dots$, which is true by induction: the content of $\bar{T}$ is correct because j is a magic value for λ, and conditions (2–6) are true for $\bar{T}$ because they are true for T. □

Our next characterisation of magic tableaux requires some additional notation. Suppose T is a λ-tableau on $i_1, i_2, \dots$. Say that T is *pre-magic* if and only if there is an integer m such that $T(x, y) = i_1$ if and only if either $(x < m,\ y \leqslant p - 1)$ or $x = m$. We allow the possibility $m > \lambda'_1$, which means that the entries in the first $p - 1$ columns of T are the entries equal to i_1. If T is pre-magic, we define the composition λ° by

$$\lambda^\circ_k = \begin{cases} \lambda_k - p + 1 & (k < m) \\ \lambda_{k+1} & (k \geqslant m), \end{cases}$$

and define the λ°-tableau T° by

$$T^\circ(x, y) = \begin{cases} T(x, y + p - 1) & (x < m) \\ T(x + 1, y) & (x \geqslant m). \end{cases}$$

Then T° is a λ°-tableau on $i_2, i_3, \dots$. We call the procedure by which T° is obtained from T *L-removal*.

Example. Let T be the magic $(10, 10, 9, 3, 3, 3)$-tableau of the last example. Then we have

$T^\circ =$

2	2	2	2	2	2	2	2
3	3	4	4	4	4	4	4
3	3	5	5	6	6	6	
3	3	3					
5	5	7					

.

Lemma 2.4. *Suppose T is a λ-tableau with values from the set $\{i_1, i_2, \dots\}$. Then T is magic on $i_1, i_2, \dots$ if and only if the following hold.*

1. *For each k, the number of entries equal to i_k in T is λ^{r}_k.*
2. *T is pre-magic.*

3. *λ° is a partition.*

4. *T° is magic on $i_2, i_3, \dots$.*

Proof. Suppose first that T is magic. Then (1) is certainly true. Define f as above.

If $f(m) = i_1$ for some m, then the conditions in Lemma 2.3 imply that T is pre-magic (with this value of m). The number of i_1s in T equals λ_1^{r}, which is the number of full ramps in λ. So the node (m, λ_m) of λ lies in a full ramp. This means that either $m = 1$ or $\lambda_{m-1} - \lambda_m \geqslant p - 1$, and either way λ° is a partition. Furthermore, by counting nodes in ramps, we find that $(\lambda^\circ)^{\mathrm{r}} = \bar{\lambda}^{\mathrm{r}}$. This implies that the conditions of Lemma 2.3 hold for T°, and so T° is magic.

If i_1 does not lie in the image of f, then i_1 appears $p - 1$ times in each row of T. We put $m = \lambda_1' + 1$, and essentially repeat the above argument (without the proof that λ° is a partition, which is trivial in this case).

Conversely, suppose that (1–4) hold. Then conditions (1–6) of Lemma 2.3 hold for T°, and so they hold for T. So T is magic. □

Recall that $T_\lambda = T_\lambda(i_1, i_2, \dots)$ denotes the first magic λ-tableau. We can characterise this as follows.

Lemma 2.5. *Suppose T is a magic λ-tableau on $i_1, i_2, \dots$. Then $T = T_\lambda$ if and only if there do not exist j, k, r such that:*

- *i_k appears at the end of row r of T;*
- *the last entry in row r of T not equal to i_k is equal to i_j;*
- *the number of entries in row r equal to i_k equals the number of entries strictly below row r which equal i_j.*

Proof. Suppose such j, k, r exist. Form the tableau $\hat{T}$ by replacing each i_k in row r of T with i_j, and replacing each i_j below row r with i_k. We claim that $\hat{T}$ is magic. By assumption, $\hat{T}$ has the same content as T, so condition (1) of Lemma 2.3 is satisfied. Conditions (2–6) are easy to verify from the corresponding conditions for T, and so the claim is true. $\hat{T}$ has an earlier magic value in row r than T, so T is not the first magic λ-tableau.

Conversely, suppose $T \neq T_\lambda$. If the first row of T agrees with the first row of T_λ, then we may remove this row from both and use induction on the number of rows. So we suppose that T and T_λ differ in the first row. Let k be the magic value for λ chosen in the construction of T, namely, k is such that the last entry of the first row of T equals i_k. The magic value chosen for T_λ must be strictly less than k, since T_λ is the first magic λ-tableau. Since the set of magic values of λ is a set of consecutive integers, $k - 1$ is a magic value for λ. If we construct a magic λ-tableau $\check{T}$ by choosing the magic value $k - 1$ in the first row, and in subsequent rows making the same choice of magic value as in T, then $\check{T}$ may be obtained from T by replacing all the entries equal to i_k in row 1 with i_{k-1}s, and replacing all the entries equal to i_{k-1} below row 1 with i_ks. The fact that T and $\check{T}$ have the same content then means that T satisfies the conditions of the lemma, with $j - k = 1$ and $r = 1$. □

Now we show that the L-removal procedure described above preserves the property of being the first magic tableau.

Lemma 2.6. *Let $T = T_\lambda(i_1, i_2, \dots)$, and define T° as above. Then $T^\circ = T_{\lambda^\circ}(i_2, i_3, \dots)$.*

Proof. T° is magic on $i_2, i_3, \dots$ by Lemma 2.4. The fact that T° is the first magic λ°-tableau follows easily from Lemma 2.5. □

3 An important lemma

We begin this section with the following well-known result about binomial coefficients.

Lemma 3.1. [4, Corollary 22.5] *Suppose that $a \equiv -1 \mod p$. Then*

$$\binom{a+1}{1}, \binom{a+2}{2}, \dots, \binom{a+p-1}{p-1} \equiv 0 \mod p.$$

This will be used to prove the following important lemma, which will help us to show both that the image of a 'magic homomorphism' lies in the Specht module $S^{\lambda^{\mathrm{r}}}$, and that if one magic homomorphism is non-zero then they all are.

Lemma 3.2. *Suppose R is a magic λ-tableau. Suppose $t \geqslant 1$, and that $(x_1, y_1), \dots, (x_t, y_t)$ are nodes of λ and $m_1, \dots, m_t, m$ are integers such that for each $1 \leqslant i \leqslant t$:*

- *$R(x_i, y_i) = m$;*
- *$m_i < m$;*
- *m_i appears at the end of row r_i of R, for some $r_i < x_i$.*

Define the λ-tableau S by

$$S(x, y) = \begin{cases} m_i & ((x, y) = (x_i, y_i) \text{ for some } i) \\ R(x, y) & (\text{otherwise}). \end{cases}$$

Then $\hat{\Theta}_S = 0$.

Proof. By re-ordering, we may assume that $r_1 \leqslant \cdots \leqslant r_t$, and we let g be maximal such that $r_1 = r_g$. Then $m_1 = \cdots = m_g \neq m_i$ for $i > g$. We want to use Lemma 1.2 to move the entries equal to m_1 in positions $(x_1, y_1), \dots, (x_g, y_g)$ of S up to row r_1. Lemma 1.2 tells us that $\hat{\Theta}_S$ is a linear combination of maps $\hat{\Theta}_U$, where each U is obtained by interchanging these entries with entries not equal to m_1 in row r_1. We will show that for each such U either the map $\hat{\Theta}_U$ is zero or the coefficient of $\hat{\Theta}_U$ given by Lemma 1.2 is divisible by p. Write $a_i = U(x_i, y_i)$ for $i = 1, \dots, g$.

Suppose first that for some j, a_j does not appear the end of any of rows $r_1 + 1, \dots, x_j - 1$. Then there are $p - 1$ entries equal to a_j in row x_j of R, and by Lemma 1.2 the coefficient of $\hat{\Theta}_U$ in $\hat{\Theta}_T$ includes a factor $\binom{p-1+a}{p-1}$, where a is the number of $1 \leqslant k \leqslant g$ such that $x_k = x_j$ and $a_k = a_j$. Since there are $p - 1$ entries equal to a_j in row r_1 of R, we have $a \leqslant p - 1$, and so the binomial coefficient $\binom{p-1+a}{p-1}$ is divisible by p, by Lemma 3.1.

So we may assume that for each j, the integer a_j appears at the end of some row between r_1 and x_j. Now if we let $\bar{R}$ be the tableau formed by rows $r_1 + 1, r_1 + 2, \dots$ of U, then by induction on λ'_1 (replacing $m_1, \dots, m_g$ with $a_1, \dots, a_g$) we find that $\hat{\Theta}_{\bar{R}} = 0$. Hence $\hat{\Theta}_U = 0$ by Lemma 1.3. □

4 $\hat{\Theta}_T$ is non-zero when T is magic

To prove that $\hat{\Theta}_T$ is non-zero when T is a magic tableau, we consider the image of a chosen polytabloid under $\hat{\Theta}_T$, and try to isolate a particular tabloid such that we can find the coefficient of this tabloid in the image. Suppose $S \sim_{\mathrm{row}} T$ and that the entries in each column of S are distinct. Let $\mathcal{U}(S)$ be the set of λ-tableaux U such that $U \sim_{\mathrm{row}} T$ and $U \sim_{\mathrm{col}} S$. For $U \in \mathcal{U}(S)$, define ϵ_{US} to be the sign of the column permutation taking U to S (this is well-defined, since the entries in each column of S are distinct). Say that S *is special for* T if $\sum_{U\in\mathcal{U}(S)} \epsilon_{US}$ is not divisible by p. Then we have the following.

Lemma 4.1. *$\hat{\Theta}_T$ is non-zero if and only if there is a tableau S which is special for T.*

Proof. Let C_{t^λ} denote the column stabiliser of the tableau t^λ and set $\kappa_{\mathrm{t}^\lambda} = \sum_{w\in C_{\mathrm{t}^\lambda}} \mathrm{sgn}(w)w$. Consider the image of the polytabloid $e_{\mathrm{t}^\lambda} = \{\mathrm{t}^\lambda\}\kappa_{\mathrm{t}^\lambda}$ (which generates S^λ) under the map $\hat{\Theta}_T$. Recall that

$$\hat{\Theta}_T : \{\mathrm{t}^\lambda\}\kappa_{\mathrm{t}^\lambda} \longmapsto \left(\sum_{U\sim_{\mathrm{row}}T} U\right)\kappa_{\mathrm{t}^\lambda}$$

where we regard $\mathcal{T}(\lambda, \lambda^{\mathrm{r}})$ as a basis for $M^{\lambda^{\mathrm{r}}}$. Each $w \in C_{\mathrm{t}^\lambda}$ acts on $U \sim_{\mathrm{row}} T$ by permuting the entries in its columns. If $S \in \mathcal{T}(\lambda, \lambda^{\mathrm{r}})$ has repeated entries in some column then the coefficient of S in $\hat{\Theta}_T(e_{\mathrm{t}^\lambda})$ is zero ([3, Lemma 13.12]); otherwise S has coefficient in $\hat{\Theta}_T(e_{\mathrm{t}^\lambda})$ equal to

$$\sum_{\substack{U\sim_{\mathrm{row}}T\\ U\sim_{\mathrm{col}}S}} \epsilon_{US} 1_{\mathbb{F}}.$$

Suppose S is special for T. Then the coefficient of S in $\hat{\Theta}_T(e_{\mathrm{t}^\lambda})$ is $\sum_{U\in\mathcal{U}(S)} \epsilon_{US}1_{\mathbb{F}}$, which is non-zero.

Conversely, suppose R is a λ-tableau of type λ^{r} such that the coefficient of R in $\hat{\Theta}_T(e_{\mathrm{t}^\lambda})$ is non-zero. This implies that the entries in each column of R are distinct, and that there is a tableau S such that $T \sim_{\mathrm{row}} S \sim_{\mathrm{col}} R$, with the coefficient of S in $\hat{\Theta}_T(e_{\mathrm{t}^\lambda})$ being ϵ_{RS} times the coefficient of R, and hence non-zero. Since this coefficient equals $\sum_{U\in\mathcal{U}(S)} \epsilon_{US}1_{\mathbb{F}}$, we find that S is special for T. □

Our proof that $\hat{\Theta}_T \neq 0$ will be by induction, and we shall often need to switch between different magic tableaux for the same λ. So we need the following result.

Lemma 4.2. *If T_1 and T_2 are magic λ-tableaux on $i_1, i_2, \dots$, then $\hat{\Theta}_{T_1} = \pm\hat{\Theta}_{T_2}$. Hence $\hat{\Theta}_{T_1}$ is non-zero if and only if $\hat{\Theta}_{T_2}$ is non-zero.*

Proof. The magic tableau T_1 is specified by choosing a magic value j_1 for λ, choosing a magic value j_2 for $\bar{\lambda}$, and so on. Similarly, T_2 is specified by a magic value j'_1 for λ, a magic value j'_2 for $\bar{\lambda}$, and so on. We say that T_1 and T_2 are *adjacent* if for some r we have

$$j'_i = \begin{cases} j_i \pm 1 & (i = r) \\ j_i & (i \neq r). \end{cases}$$

Since the set of magic values for a partition is a set of consecutive integers, and since the relation $\hat{\Theta}_{T_1} = \pm\hat{\Theta}_{T_2}$ on the set of magic λ-tableaux is an equivalence relation, it suffices to consider the case where T_1 and T_2 are adjacent. So we suppose that the integers j_i and j_i' are as above, with $j_r' = j_r + 1$.

Let l be the last entry in row r of T_1, and let m be the last entry in row r of T_2. Then for any node $\mathfrak{n}$ of λ below row r, we have $T_1(\mathfrak{n}) = m$ if and only if $T_2(\mathfrak{n}) = l$, and T_2 may be obtained from T_1 by replacing all the entries equal to m below row r with ls, and replacing all but $p-1$ of the entries equal to l in row r with ms. We shall apply Lemma 1.2 to T_1, to move the entries equal to m below row r up to row r. Lemma 1.2 tells us that $\hat{\Theta}_{T_1}$ is equal to a linear combination of homomorphisms $\hat{\Theta}_U$, where U is a row standard tableau obtained from T_1 by interchanging the entries equal to m below row r with entries not equal to m in row r. Of course, T_2 is such a tableau, and we wish to show that the coefficient of $\hat{\Theta}_{T_2}$ in this linear combination is ± 1, while for any $U \neq T_2$ either the coefficient of $\hat{\Theta}_U$ is zero, or the map $\hat{\Theta}_U$ equals zero.

If all the entries that we bring down from row r are equal to l then, since T_1 does not contain any entries equal to l below row r, the binomial coefficients occurring in Lemma 1.2 are all equal to 1, and so we find that the coefficient of $\hat{\Theta}_{T_2}$ is ± 1. Now we suppose that U is a row standard tableau obtained as above, different from T_2. Suppose the entries equal to m below row r in T_1 appear in positions $(x_1, y_1), \dots, (x_u, y_u)$, and set $a_i = U(x_i, y_i)$ for each i. By re-ordering, we may find t such that $a_i = l$ if and only if $i > t$. The fact that $U \neq T_2$ means that $t \geqslant 1$.

Suppose that for some $j \leqslant t$ the integer a_j does not appear at the end of any of rows $r+1, \dots, x_j - 1$ of T_1. Then a_j appears exactly $p-1$ times in row x_j of T_1, and so the coefficient of $\hat{\Theta}_U$ in $\hat{\Theta}_{T_1}$ includes the binomial coefficient $\binom{p-1+a}{p-1}$, where a is the number of k such that $a_k = a_j$ and $x_k = x_j$. Since row r of T_1 contains $p-1$ entries equal to a_j, we have $a \leqslant p-1$, so the coefficient $\binom{p-1+a}{p-1}$ is divisible by p.

So we may assume that for each $j \leqslant t$ the integer a_j appears at the end of some row of T_1 above row x_j. Now we may apply Lemma 3.2, letting R be the tableau formed from rows $r+1, r+2, \dots$ of T_2, and S the tableau formed from rows $r+1, r+2, \dots$ of U, with $m_i = a_i$ for $i = 1, \dots, t$. We find that $\hat{\Theta}_S = 0$, which means that $\hat{\Theta}_U = 0$ by Lemma 1.3. □

Now we describe the inductive step in our proof. Given a λ-tableau T, write ${}^{|}T$ for the ${}^{|}\lambda$-tableau with ${}^{|}T(x, y) = T(x, y)$ for all x, y. (So ${}^{|}T$ is T with the last entry deleted from each row.)

Lemma 4.3. *Let $T = T_\lambda(i_1, i_2, \dots)$. Then ${}^{|}T$ is a magic ${}^{|}\lambda$-tableau on $i_1, i_2, \dots$.*

This relies on the following comparison between λ^{r} and $({}^{|}\lambda)^{\mathrm{r}}$.

Lemma 4.4.

$$({}^{|}\lambda)^{\mathrm{r}}_k = \begin{cases} \lambda^{\mathrm{r}}_k - 1 & (\textit{if } k \textit{ appears at the end of some row of } T_\lambda(i_1, i_2, \dots)) \\ \lambda^{\mathrm{r}}_k & (\textit{otherwise}). \end{cases}$$

Proof. We work by induction on λ_1', the case where λ has no non-zero parts being trivial. We assume the lemma holds for $\bar{\lambda}$, that is

$$({}^{|}\bar{\lambda})^{\mathrm{r}}_k = \begin{cases} (\bar{\lambda})^{\mathrm{r}}_k - 1 & (\text{if } k \text{ appears at the end of some row of } T_{\bar{\lambda}}(i_1, i_2, \dots)) \\ (\bar{\lambda})^{\mathrm{r}}_k & (\text{otherwise}). \end{cases}$$

Let j be the first magic value for λ and j^* the first magic value for ${}^{|}\lambda$ so that $j = \mathrm{mis}^{\lambda}(\lambda_1) + 1$ and $j^* = \mathrm{mis}^{{}^{|}\lambda}({}^{|}\lambda_1) + 1$. Therefore either $j = j^*$ or $j = j^* + 1$.

Let $\bar{M}$ be the set of numbers which appear at the ends of the rows of $T_{\bar{\lambda}}(i_1, i_2, \ldots)$ and M be the set of numbers appearing at the ends of the rows of $T_{\lambda}(i_1, i_2, \ldots)$. By construction,

$$M = \{i_k \mid i_k \in \bar{M} \text{ and } k < j\} \cup \{i_j\} \cup \{i_k \mid i_{k-1} \in \bar{M} \text{ and } k > j\}.$$

Now, from Lemma 2.1

$$(\bar{\lambda})^{\mathrm{r}}_k = \begin{cases} \lambda^{\mathrm{r}}_k - p + 1 & (k < j) \\ \lambda^{\mathrm{r}}_{k+1} & (k \geqslant j) \end{cases}$$

and

$$({}^{|}\bar{\lambda})^{\mathrm{r}}_k = \begin{cases} ({}^{|}\lambda)^{\mathrm{r}}_k - p + 1 & (k < j^*) \\ ({}^{|}\lambda)^{\mathrm{r}}_{k+1} & (k \geqslant j^*). \end{cases}$$

Therefore (omitting the case $k = j^*$ for the moment)

$$\begin{aligned} ({}^{|}\lambda)^{\mathrm{r}}_k &= \begin{cases} ({}^{|}\bar{\lambda})^{\mathrm{r}}_k + p - 1 & (k < j^*) \\ ({}^{|}\bar{\lambda})^{\mathrm{r}}_{k-1} & (k > j^*) \end{cases} \\ &= \begin{cases} (\bar{\lambda})^{\mathrm{r}}_k + p - 2 & (k < j^*,\ i_k \in \bar{M}) \\ (\bar{\lambda})^{\mathrm{r}}_{k-1} - 1 & (k > j^*,\ i_{k-1} \in \bar{M}) \\ (\bar{\lambda})^{\mathrm{r}}_k + p - 1 & (k < j^*,\ i_k \notin \bar{M}) \\ (\bar{\lambda})^{\mathrm{r}}_{k-1} & (k > j^*,\ i_{k-1} \notin \bar{M}). \end{cases} \end{aligned}$$

Now note that since either $j = j^*$ or $j = j^* + 1$, a situation where $k \geqslant j$ and $k < j^*$ cannot arise; and a situation where $k - 1 < j$ and $k > j^*$ can only arise if $j = j^* + 1$ and $k = j$. Therefore for $k \neq j, j^*$,

$$({}^{|}\lambda)^{\mathrm{r}}_k = \begin{cases} \lambda^{\mathrm{r}}_k - 1 & (k < j,\ k < j^*,\ i_k \in \bar{M}) \\ \lambda^{\mathrm{r}}_k & (k < j,\ k < j^*,\ i_k \notin \bar{M}) \\ \lambda^{\mathrm{r}}_k - 1 & (k - 1 \geqslant j,\ k > j^*,\ i_{k-1} \in \bar{M}) \\ \lambda^{\mathrm{r}}_k & (k - 1 \geqslant j,\ k > j^*,\ i_{k-1} \notin \bar{M}). \end{cases}$$

Comparing this with the expression for M above, we find that for $k \neq j, j^*$,

$$({}^{|}\lambda)^{\mathrm{r}}_k = \begin{cases} \lambda^{\mathrm{r}}_k - 1 & (i_k \in M) \\ \lambda^{\mathrm{r}}_k & (\text{otherwise}). \end{cases}$$

It is also clear that $({}^{|}\lambda)^{\mathrm{r}}_j = \lambda^{\mathrm{r}}_j - 1$ since the node at the end of row j of λ^{r} is the highest node on ramp λ_1. If λ is a partition of n, then ${}^{|}\lambda$ is a partition of $n - |M|$; hence the result must also hold for $k = j^*$. □

Proof of Lemma 4.3. Since T is magic, we may construct the tableau T° as defined before Lemma 2.4. We define the tableau ${}^{|}T^\circ$ by removing the last entry from each row of T°, or equivalently by using the L-removal procedure on ${}^{|}T$ (which is certainly pre-magic). Since $T^\circ = T_{\lambda^\circ}(i_2, i_3, \dots)$ by Lemma 2.6, we find by induction that ${}^{|}T^\circ$ is magic. So we find that conditions (2–4) of Lemma 2.4 hold for ${}^{|}T$; condition (1) follows from Lemma 4.4, and so ${}^{|}T$ is magic. □

Proposition 4.5. *Suppose T is a magic λ-tableau on $1, 2, \dots$. Then $\hat{\Theta}_T \neq 0$.*

Proof. By Lemma 4.2, we may assume that $T = T_\lambda(1, 2, \dots)$. If we construct ${}^{|}T$ as above, then ${}^{|}T$ is magic, so by induction $\hat{\Theta}_{{}^{|}T}$ is non-zero. So there exists a ${}^{|}\lambda$-tableau S which is special for ${}^{|}T$. Let $f(r)$ denote the entry at the end of row r of T, and then for any ${}^{|}\lambda$-tableau R define the λ-tableau R^+ by

$$R^+(x, y) = \begin{cases} R(x, y-1) & (y > 1) \\ f(x) & (y = 1). \end{cases}$$

We claim that S^+ is special for T. First we need to show that

$$\mathcal{U}(S^+) = \{U^+ \mid U \in \mathcal{U}(S)\},$$

i.e. that if V is a λ-tableau such that $V \sim_{\mathrm{row}} S^+$ and $V \sim_{\mathrm{col}} S^+$, then $V(x, 1) = f(x)$ for all x. We prove this by induction on x. There is exactly one entry equal to $f(x)$ in the first column of V. Now $f(x)$ appears at the end of row x of T, and so by the construction of magic tableaux, all the entries of T equal to $f(x)$ occur in or above row x. We have $V \sim_{\mathrm{row}} S^+ \sim_{\mathrm{row}} T$, so the entry $f(x)$ in the first column of V appears in one of the positions $(1, 1), (2, 1), \dots (x, 1)$; but by induction we know that $V(z, 1) = f(z) \neq f(x)$ for $1 \leqslant z < x$, and so we have $V(x, 1) = f(x)$. So the above description of $\mathcal{U}(S^+)$ follows. It is clear that $\epsilon_{U^+S^+} = \epsilon_{US}$ for $U \in \mathcal{U}(S)$, and hence we find that S^+ is special for T. □

5 The image of a magic homomorphism lies in the Specht module $S^{\lambda^{\mathrm{r}}}$

In this section, we show that if T is a magic λ-tableau, then the image of $\hat{\Theta}_T$ lies in the Specht module $S^{\lambda^{\mathrm{r}}}$, which will complete the proof of Theorem 2.2. Following the discussion in Section 1.1.1, we wish to show that the compositions $\psi_{d,t} \circ \hat{\Theta}_T$ are all equal to zero.

Fortunately, it is easy to calculate $\psi_{d,t} \circ \hat{\Theta}_T$. Let $\Psi_{d,t}(T)$ be the set of row standard λ-tableaux which may be obtained by replacing $\lambda^{\mathrm{r}}_{d+1} - t$ of the entries equal to $d+1$ in T with ds. Recall the notation U_i^d from Section 1.1.1, and for each $U \in \Psi_{d,t}(T)$ define

$$c_U = \prod_{i \geqslant 1} \binom{U_i^d}{T_i^d},$$

Then we have the following.

Lemma 5.1. [2, Lemma 5]

$$\psi_{d,t} \circ \hat{\Theta}_T = \sum_{U \in \Psi_{d,t}(T)} c_U 1_{\mathbb{F}} \hat{\Theta}_U.$$

Our task is to show that if T is a magic λ-tableau, then for each $U \in \Psi_{d,t}(T)$ either c_U is divisible by p or $\hat{\Theta}_U = 0$.

Lemma 5.2. *Suppose that $U \in \Psi_{d,t}(T)$ is such that, for some i, $T_i^d = p-1$ and $U_i^d = b+p-1$, where $1 \leqslant b \leqslant p-1$. Then $p \mid c_U$.*

Proof. c_U contains a factor $\binom{b+p-1}{b} \equiv 0 \mod p$. □

Lemma 5.3. *If $U \in \Psi_{d,t}(T)$ then either c_U is divisible by p or $\hat{\Theta}_U = 0$.*

Proof. Let M be the set of numbers that appear at the ends of the rows of T. We consider four separate cases.

1. $d, d+1 \notin M$.

   Let $U \in \Psi_{d,t}(T)$. Then U satisfies the conditions of Lemma 5.2 and hence c_U is divisible by p.

2. $d \in M, d+1 \notin M$.

   Suppose that d appears at the end of row m_d of T. If for some $i < m_d$ we have that $U_i^d > T_i^d$ then U satisfies the conditions of Lemma 5.2 and $p \mid c_U$. Otherwise, all nodes on which T and U differ must be in rows below m_d, and by Lemma 3.2 we find that that $\hat{\Theta}_U = 0$.

3. $d \notin M, d+1 \in M$.

   Suppose that $d+1$ appears at the end of row m_{d+1} of T. If for some $i < m_{d+1}$ row i of U contains more entries equal to d then T does then U satisfies the conditions of Lemma 5.2 and $p \mid c_U$. Similarly if row m_{d+1} of U contains b more entries equal to d in row m_{d+1} than T does, with $1 \leqslant b \leqslant p-1$, Lemma 5.2 shows that $p \mid c_U$. We are left with the situation where U is formed from T by changing $b \geqslant p$ entries equal to $d+1$ into d in row m_{d+1}, and we claim that in this situation we have $\hat{\Theta}_U = 0$. Note that

   $$\begin{aligned} U_{d+1}^{m_{d+1}} &< T_{d+1}^{m_{d+1}} - p + 1 \\ &\leqslant \sum_{i>m_d} T_d^i \end{aligned}$$

   (since λ^{r} is a partition)

   $$= \sum_{i>m_d} U_d^i.$$

   We now apply Lemma 1.2 to U to bring the entries equal to d below row m_{d+1} into row m_{d+1}. We find that $\hat{\Theta}_U$ is a linear combination of maps $\hat{\Theta}_V$, where V is a row standard tableau obtained from U by interchanging all the entries equal to d below row m_{d+1} with some entries not equal to d in row m_{d+1}. Suppose $(x_1, y_1), \ldots, (x_u, y_u)$ are the nodes below row m_{d+1} such that $U(x_i, y_i) = d$. Given a tableau V as described above, write $a_i = V(x_i, y_i)$ for $i = 1, \ldots, u$. By re-ordering, we may find t such that $a_i = d+1$ if and only if $i > t$. By the above inequality, we have $t \geqslant 1$.

Suppose first that for some $j \leqslant t$ the integer a_i does not appear at the end of any of rows $m_{d+1}+1, \dots, x_i - 1$ of U. As in the proof of Lemma 4.2, we find that the coefficient of $\hat{\Theta}_V$ in the expression for $\hat{\Theta}_U$ given by Lemma 1.2 is zero. So we suppose that for each $j \leqslant t$ the integer a_j does appear at the end of one of rows $m_{d+1}, \dots, x_j - 1$. Let R' be the tableau formed by rows $m_{d+1}+1, m_{d+1}+2, \dots$ of U (or equivalently, of T), and let S be the tableau formed from the corresponding rows of V. Then R' is a magic tableau on some set I of integers which includes d but not $d+1$. If we form the tableau R by changing all the ds in R' into $d+1$s, then R is magic on $(I \setminus \{d\}) \cup \{d+1\}$. Now we may apply Lemma 3.2, with $m_i = a_i$ for $i = 1, \dots, t$, and we get $\hat{\Theta}_S = 0$, and hence $\hat{\Theta}_V = 0$, by Lemma 1.3. Hence $\hat{\Theta}_U = 0$.

4. $d, d+1 \in M$.

   Suppose that d appears at the end of row m_d and $d+1$ appears at the end of row m_{d+1}. If $m_d < m_{d+1}$ then we may repeat the argument of Case 2 above. If $m_d > m_{d+1}$, we may repeat the argument of Case 3.

□

This completes the proof of Theorem 2.2, and hence of Theorem 1.5.

Acknowledgement. The second author would like to thank Cambridge University for the use of its facilities.

References

[1] R. Carter & M. Payne, 'On homomorphisms between Weyl modules and Specht modules', *Math. Proc. Cambridge Philos. Soc.* **87** (1980), 419–425.

[2] M. Fayers & S. Martin, 'Homomorphisms between Specht modules', *Math. Z.* **248** (2004), 395–421.

[3] G. James, 'On the decomposition matrices of the symmetric groups II', *J. Algebra* **43** (1976), 45–54.

[4] G. James, *The representation theory of the symmetric groups*, Lecture notes in mathematics **682**, Springer-Verlag, New York/Berlin, 1978.

[5] G. James & A. Kerber, *The representation theory of the symmetric group*, Encyclopædia of Mathematics and its Applications **16**, Addison–Wesley, 1981.

[6] M. Koppinen, 'Homomorphisms between neighboring Weyl modules', *J. Algebra* **103** (1986), 302–319.