Discrete Mathematics 68 (1988) 281–292
North-Holland

INDUCTIVE AND INJECTIVE PROOFS OF LOG CONCAVITY RESULTS

Bruce E. SAGAN
Department of Mathematics, University of Pennsylvania, Philadelphia, PA 19104–6395, U.S.A.

Received 14 July 1986

Given a triangular array of non-negative integers we give a necessary condition to insure that every row of the array is log concave. This result is then used to inductively construct injections showing the log concavity of the binomial coefficients and Stirling numbers of both kinds. Finally these proofs are related to the graphical interpretation of these numbers given by Wilf.

To Herbert Wilf, whose tantalizing questions inspired this paper.

1. Introduction

A sequence of non-negative integers $(a_k)_{0\leq k\leq n}$ is called *log concave* if

$$a_{k-1}a_{k+1}\leq a_k^2,\quad \text{for all } k,\ 0<k<n. \tag{1}$$

It is easy to show that this is equivalent to the seemingly stronger condition

$$a_k a_l \leq a_{k+i}a_{l-i}, \tag{2}$$

for all k, l and i satisfying $0\leq k\leq l\leq n$ and $0\leq i\leq k-l$. Proving that various sequences are log concave has been a topic of increasing interest in recent years [2, 3, 5, 6, 7]. If the sequence (a_k) is enumerative, then it may be possible to give a combinatorial proof of (1) by exhibiting an injection from pairs counted by $a_{k-1}a_{k+1}$ to those corresponding to a_k^2. This approach will be taken up in Section 3.

Given a set of non-negative integers t_{nk} defined for $n\geq 0$ and $0\leq k\leq n$, we say that t_{nk} is *log concave in* k if for any fixed n the sequence $(t_{nk})_{0\leq k\leq n}$ is log concave. If the t_{nk} are arranged in a triangular array with n and k being the row and column indices respectively, then log concavity in k corresponds to the log concavity of each row. Three famous triangles of this type are Pascal's triangle and Stirling's triangles of the first and second kinds.

In addition, the binomial coefficients and Stirling numbers all satisfy two-term recurrence relations. In Section 2 we give conditions on the coefficients of such recurrences that guarantees log concavity in k. As a consequence we obtain "one-line" proofs that the three examples above satisfy this condition. Also unwinding the recursion leads us to the injective log concavity proofs mentioned in the first paragraph.

Finally, we show that these combinatorial constructions have a simple interpretation in terms of paths through graphical structures developed by Wilf [8, 9]. In this setting it is reminiscent of a technique of Gessel and Viennot [4] for counting lattice paths.

2. Inductive proofs

Now suppose that for $n \geq 1$ and $0 < k < n$ the t_{nk} satisfy the recurrence

$$t_{nk} = c_{nk} t_{n-1\,k-1} + d_{nk} t_{n-1\,k}, \tag{3}$$

where the coefficients c_{nk}, d_{nk} are non-negative integers. To eliminate boundary conditions, it is convenient to extend the definitions of t_{nk} by setting $t_{nk} = 0$ for $k < 0$ or $k > n$. We assume that in this extended range c_{nk} and d_{nk} can be chosen so that (3) continues to hold. Also note that if the t_{nk} were log concave in k for $0 \leq k \leq n$, then they continue to be for all integral k since the new inequalities are satisfied trivially.

With these preliminaries we can introduce our primary tool.

Theorem 1. *Let t_{nk} be an extended triangular array satisfying $t_{nk} = c_{nk} t_{n-1\,k-1} + d_{nk} t_{n-1\,k}$, for all $n \geq 1$ where t_{nk}, c_{nk} and d_{nk} are all non-negative integers. Suppose that*

(i) *c_{nk} and d_{nk} are log concave in k,*
(ii) *$c_{nk-1} d_{nk+1} + c_{nk+1} d_{nk-1} \leq 2 c_{nk} d_{nk}$, for all $n \geq 1$,*

then t_{nk} is log concave in k.

Proof. Induct on n. The $n = 0$ row of the array is automatically log concave since it contains at most one non-zero entry. Now suppose that the $(n-1)$st row is log concave.

To prove $t_{nk-1} t_{nk+1} \leq t_{nk}^2$ we expand both sides by the recurrence and compare corresponding terms. Thus it is enough to show that the following three equations hold

$$c_{nk-1} c_{nk+1} t_{n-1\,k-2} t_{n-1\,k} \leq c_{nk}^2 t_{n-1\,k-1}^2, \tag{4}$$

$$d_{nk-1} d_{nk+1} t_{n-1\,k-1} t_{n-1\,k+1} \leq d_{nk}^2 t_{n-1\,k}^2, \tag{5}$$

and

$$c_{nk-1} d_{nk+1} t_{n-1\,k-2} t_{n-1\,k+1} + c_{nk+1} d_{nk-1} t_{n-1\,k-1} t_{n-1\,k} \leq 2 c_{nk} d_{nk} t_{n-1\,k-1} t_{n-1\,k}. \tag{6}$$

Log concavity of row $n-1$ (in the form given in Section 1, eq. (2)) together with assumptions (i) and (ii) are precisely what is needed to demonstrate these three formulae. $\square$

The reader may feel that Theorem 1 is not much of a labor-saving device as it endeavours to demonstrate the log concavity of a single array by proving that *two* others are log concave with an extra condition besides! However, this is offset by the fact that the coefficient arrays are usually much simpler than the original one. In particular, this will be true of the three examples mentioned in the introduction.

If n is a non-negative integer then let $[n]$ denote the set $\{1, 2, \ldots, n\}$. Recall that the binomial coefficients $\binom{n}{k}$, (signless) Stirling numbers of the first kind $c(n, k)$, and Stirling numbers of the second kind $S(n, k)$ count subsets of $[n]$ with k elements, permutations of $[n]$ with k cycles, and partitions of $[n]$ with k blocks respectively. The recurrence relations for these functions are well known and follow directly from their directions:

$$\binom{n}{k} = \binom{n-1}{k-1} + \binom{n-1}{k},$$
$$c(n, k) = c(n-1, k-1) + (n-1)c(n-1, k),$$
$$S(n, k) = S(n-1, k-1) + kS(n-1, k).$$

Corollary 2. $\binom{n}{k}$, $c(n, k)$ and $S(n, k)$ *are all log concave in* k.

Proof. In each case we verify conditions (i) and (ii). For $\binom{n}{k}$, $c_{nk} = d_{nk} = 1$ when $n \geq 1$ and constant sequences are trivially log concave. Furthermore,

$$c_{nk-1}d_{nk+1} + c_{nk+1}d_{nk-1} = 2 = 2c_{nk}d_{nk}.$$

The Stirling numbers of the first kind have $c_{nk} = 1$, $d_{nk} = n - 1$ for all $n \geq 1$ which are also constant (with respect to k) hence satisfying (i). In this case (ii) reduces to $2(n-1) \leq 2(n-1)$. Finally for $S(n, k)$ we let

$$c_{nk} = \begin{cases} 1, & \text{for } k \geq 0, \\ 0, & \text{for } k < 0, \end{cases} \qquad d_{nk} = \begin{cases} k, & \text{for } k \geq 0, \\ 0, & \text{for } k < 0. \end{cases}$$

so that $d_{nk} \geq 0$, always. Log concavity for d_{nk} follows from the fact that $(k-1)(k+1) = k^2 - 1 \leq k^2$. Also

$$\begin{aligned} c_{nk-1}d_{nk+1} + c_{nk+1}d_{nk-1} &= \begin{cases} (k+1) + (k-1) = 2k, & \text{if } k > 0, \\ 0, & \text{if } k \leq 0, \end{cases} \\ &= 2c_{nk}d_{nk} \quad \text{in both cases.} \qquad \square \end{aligned}$$

3. Injective proofs

We can now construct an injective demonstration of log concavity in k for any triangular array t_{nk} satisfying the hypotheses of Theorem 1 as follows.

(a) Define the injection for small n in some reasonable way. (Often this definition will be forced on us because t_{nk} will also be small and so there will not be many choices.)

(b) Use eqs. (4), (5) and (6) to extend the injection to larger n until a pattern evolves.

(c) Prove that the pattern holds for all n.

In the interests of brevity, only the results of step (c) will be given below for our three canonical examples.

Let B_{nkl} be the set of all pairs $(S, T) \subseteq [n] \times [n]$ with $|S| = k$ and $|T| = l$, $0 \leq k \leq l \leq n$. Also let $S_i = S \cap [i]$ and $S_i^c = S - S_i$. Now given $(S, T) \in B_{nk-1k+1}$ consider the sequence of pairs

$$(S, T) = (S_n, T_n), (S_{n-1}, T_{n-1}), \ldots, (S_0, T_0) = (\emptyset, \emptyset).$$

Since $|T_n| = |S_n| + 2$ and $|T_0| = |S_0|$ there must be a largest index i such that $|T_i| = |S_i| + 1$. Now define a map f_n by

$$f_n(S, T) = (T_i \cup S_i^c, S_i \cup T_i^c). \tag{7}$$

Theorem 3. *f_n as defined by eq. (7) is an injection from $B_{nk-1k+1}$ to B_{nkk}.*

Proof. We first show that f_n is well defined, i.e., that $f_n(S, T) \in B_{nkk}$. But if $|T_i| = |S_i| + 1 = m + 1$, say, then

$$|T_i \cup S_i^c| = m + 1 + [(k-1) - m] = k$$

and

$$|S_i \cup T_i^c| = m + [(k+1) - (m+1)] = k.$$

We will give two proofs that f_n is one-to-one. The first will make f_n's heritage from Theorem 1 apparent, but the second will be simpler.

Demonstration 1. Extend f_n to a function $f_n : B_{nkl} \to B_{nk+1l-1}$, where $l - k \geq 2$ by again finding the first index i such that $|T_i| = |S_i| + 2$ and then defining f_n by (7). It is easy to show that f_n is still well defined and injectivity will be proved by induction on n. In fact we will actually be proving log concavity in the form given by eq. (2).

Clearly f_n is an injection if $n \leq 2$ so assume the result for $n - 1$. Given $S, T \subseteq [n]$ construct $S', T' \subseteq [n-1]$ by

$$(S', T') = \begin{cases} f_{n-1}(S_{n-1}, T_{n-1}), & \text{if } |T_{n-1}| \geq |S_{n-1}| + 2, \\ (T_{n-1}, S_{n-1}), & \text{if } |T_{n-1}| = |S_{n-1}| + 1. \end{cases}$$

From the definition of f we have

$$f_n(S, T) = \begin{cases} (S' \cup \{n\}, T' \cup \{n\}), & \text{if } n \in S, n \in T, & (4') \\ (S', T'), & \text{if } n \notin S, n \notin T, & (5') \\ (S' \cup \{n\}, T'), & \text{if } n \in S, n \notin T, & (6a') \\ (S', T' \cup \{n\}), & \text{if } n \notin S, n \notin T. & (6b') \end{cases}$$

Note that eqs. (4′) and (5′) correspond to inequalities (4) and (5) in Theorem 1 while (6a′) and (6b′) mirror the first and second terms in (6). From the inductive hypothesis and the placement of n's above we see that f_n is one-to-one.

Demonstration 2. Injectivity will follow if we can construct an inverse map from the image of f_n back to its domain. Let (S', T') be a pair mapped onto by f_n. Thus there must be some index, and so a largest one, such that $|T_i'| = |S_i'| - 1$. Letting $f_n^{-1}(S', T') = (T_i' \cup S_i'^{c}, S_i' \cup T_i'^{c})$ we see that f_n^{-1} is such an inverse. □

We should note that f_n was independently discovered by Butler [2]. See Section 5 for details.

Next consider permutations π of $[n]$. Given a cycle c we define c_i to be the cycle whose elements are those of $c \cap [i]$ in the same relative order as they are in c. For example, if $c = (1, 6, 5, 3, 9, 4, 7)$, then $c_5 = (1, 5, 3, 4)$. But any permutation π of $[n]$ is uniquely decomposable as a product of disjoint cycles. So let π_i be the permutation of $[i]$ whose cycles are the c_i for all cycles c in the decomposition of π (empty cycles are deleted).

From now on we will always write π in *standard cycle notation*, where each cycle is led by its smallest element and the cycles are ordered lexicographically. If

$$\pi = (a_1, a_2, \ldots, a_p)(a_{p+1}, \ldots, a_q)\ldots(a_s, \ldots, a_{n-1})$$

is a permutation of $[n-1]$, then there are n positions where we can insert an n so as to obtain a permutation π' of $[n]$ in standard form. Position s, $s < n$, occurs in the same cycle with and directly after a_s. Position n is a space for forming a new cycle (n) at the end of π. Figure 1 shows the result of inserting 4 into a permutation of [3] in all possible ways.

Let S_{nkl} be the set of all pairs of permutations (π, σ) where π and σ are products of k and l disjoint cycles respectively. Given $(\pi, \sigma) \in S_{nk-1k+1}$, then, as before, there is a largest index i such that σ_i has one more cycle than π_i. Now define $(\pi', \sigma') = g_n(\pi, \sigma)$ by constructing a sequence of pairs of permutations

$$(\pi_i', \sigma_i') = (\sigma_i, \pi_i),\ (\pi_{i+1}', \sigma_{i+1}'), \ldots, (\pi_n', \sigma_n') = (\pi, \sigma),$$

where π_{j+1}' (respectively σ_{j+1}') is obtained from π_j' (respectively σ_j') by inserting

	Position of 4	Resulting π'
	1	(1, 4, 3)(2)
	2	(1, 3, 4)(2)
$\pi = (1, 3)(2)$	3	(1, 3)(2, 4)
	4	(1, 3)(2)(4)

Fig. 1. Insertion of 4 into a permutation of [3].

$j+1$ in the same position as in passing from π_j to π_{j+1} (respectively σ_j to σ_{j+1}). For example, if

$$(\pi, \sigma) = ((1,5,3)(2,4), (1)(2)(3,5)(4)) = (\pi_5, \sigma_5),$$

then

$$\begin{aligned}
(\pi_4, \sigma_4) &= ((1,3)(2,4), (1)(2)(3)(4)),\\
(\pi_3, \sigma_3) &= ((1,3)(2), (1)(2)(3)),\\
(\pi_3', \sigma_3') &= ((1)(2)(3), (1,3)(2)),\\
(\pi_4', \sigma_4') &= ((1)(2)(3,4), (1,3)(2)(4)),\\
(\pi_5', \sigma_5') &= ((1,5)(2)(3,4), (1,3)(2,5)(4)) = (\pi', \sigma').
\end{aligned}$$

It is easy to verify that $(\pi', \sigma') \in S_{nkk}$ and that g_n has an inverse which is a step-by-step reversal of the definition of g_n. Hence

Theorem 4. *$g_n : S_{nk-1k+1} \to S_{nkk}$ defined above is an injection.*

The construction for partitions λ of $[n]$ is similar to the one for permutations. We will separate the blocks of λ using slashes to avoid a profusion of parentheses, e.g.,

$$\lambda = 1,3,5/2,6/4 \quad \text{has blocks} \quad \{1,3,5\}, \quad \{2,6\} \text{ and } \{4\}.$$

The *standard notation* for partitions is to write the elements of each block in lexicographic order and then order the blocks lexiographically as was done in the previous example. As expected, λ_i is the partition obtained from λ by intersecting each block with $[i]$ and discarding empty blocks.

Given λ a partition of $[n-1]$ with k blocks, there are $k+1$ positions in which to place an n so as to obtain a partition λ' of $[n]$. We can put n in position s by placing it at the end of the sth block for $1 \le s \le k$. The other possibility is to make $\{n\}$ a singleton block at the end of λ. For convenience in what follows, this is called position 0. Figure 2 illustrates all ways of placing a 4 in a partition of [3].

Let the set of all pairs (λ, μ) of partitions of $[n]$ with k and l blocks respectively be denoted P_{nkl}. If $(\lambda, \mu) \in P_{nk-1k+1}$, then we find the largest i such that μ_i has

	Position of 4	Resulting π'
	0	$1, 3/2/4$
$\lambda = 1, 3/2$	1	$1, 3, 4/2$
	2	$1, 3/2, 4$

Fig. 2. Placing a 4 in a partition of [3].

one more block than λ_i and build $(\lambda', \mu') = h_n(\lambda, \mu)$ from $(\lambda_i', \mu_i') = (\mu_i, \lambda_i)$ inductively.

It is important to note that whenever $(\lambda_j, \mu_j) \in B_{jlm}$, then λ_j' and μ_j' will be constructed so that $(\lambda_j', \mu_j') \in B_{jl+1m-1}$. This is certainly the case for (λ_i', μ_i'), so we can assume that $(\lambda_{j-1}', \mu_{j-1}')$ has this property and build λ_j' from λ_{j-1}' (respectively μ_j' from μ_{j-1}') by placing j in the same position as in passing from λ_{j-1} to λ_j (respectively μ_{j-1} to μ_j). The only time this is not possible is when j is put in position m to form μ_j since there is no position m in μ_{j-1}'.

This brings us to two exceptional cases.

(i) If j was put in positions $s \geq 1$ and m to form λ_j and μ_j respectively, then place j in positions $l+1$ and s to construct λ_j' and μ_j'. Notice that by the inductive hypothesis on $(\lambda_{j-1}', \mu_{j-1}')$, positions s and $l+1$ are guaranteed to exist in this case.

(ii) If j was put in positions $s = 0$ and m to form λ_j and μ_j respectively, then we proceed as follows. Under these assumptions $(\lambda_j, \mu_j) \in B_{jlm}$ implies $(\lambda_{j-1}', \mu_{j-1}') \in B_{j-1lm-1}$ with $(m-1) - l \geq 1$ since $m - l \geq 2$. Hence we can find a largest index p, $i < p < j$ such that μ_p' has exactly one more block than λ_p'. Let $(\lambda_p'', \mu_p'') = (\mu_p', \lambda_p')$ and form the sequence

$$(\lambda_p'', \mu_p''), (\lambda_{p+1}'', \mu_{p+1}''), \ldots, (\lambda_{j-1}'', \mu_{j-1}''),$$

using the same rules as above with all unprimed symbols replaced by primed ones and all primes replaced by double-primes. Finally we let (λ_j', μ_j') be $(\lambda_{j-1}'', \mu_{j-1}'')$ with j placed in positions $l+1$ and 0 respectively. The astute reader will have noticed that the double-primed sequence comes from a composition of maps

$$B_{qlm} \xrightarrow{h_q} B_{ql+1m-1} \xrightarrow{h_q} B_{ql+2m-2},$$

i.e.,

$$(\lambda_q, \mu_q) \xrightarrow{h_q} (\lambda_q', \mu_q') \xrightarrow{h_q} (\lambda_q'', \mu_q'').$$

The following example will illustrate the general method, including both exceptional cases.

$$
\begin{aligned}
(\lambda, \mu) &= (123456/7/8,\ 1/28/34/5/67) = (\lambda_8, \mu_8),\\
(\lambda_7, \mu_7) &= (123456/7,\ 1/2/34/5/67),\\
(\lambda_6, \mu_6) &= (123456,\ 1/2/34/5/6),\\
(\lambda_5, \mu_5) &= (12345,\ 1/2/34/5),\\
(\lambda_4, \mu_4) &= (1234,\ 1/2/34),\\
(\lambda_3, \mu_3) &= (123,\ 1/2/3),\\
(\lambda_2, \mu_2) &= (12,\ 1/2),\\
(\lambda_2', \mu_2') &= (1/2,\ 12),
\end{aligned}
$$

$$
\begin{aligned}
(\lambda_3', \mu_3') &= (13/2, 12/3),\\
(\lambda_4', \mu_4') &= (13/24, 124/3),\\
(\lambda_5', \lambda_5') &= (135/24, 124/3/5),\\
(\lambda_6', \mu_6') &= (1356/24, 124/3/5/6),\\
(\lambda_5'', \mu_5'') &= (124/3/5, 135/24),\\
(\lambda_6'', \mu_6'') &= (1246/3/5, 135/24/6),\\
(\lambda_7', \mu_7') &= (1246/3/57, 135/24/6/7),\\
(\lambda_8', \mu_8') &= (1246/3/57/8, 135/248/6/7) = (\lambda', \mu').
\end{aligned}
$$

Theorem 5. *The map $h_n(\lambda, \mu) = (\lambda', \mu')$ defined above is an injection from $B_{nk-1k+1}$ to B_{nkk}.*

Proof. The fact that h_n is well-defined is an easy induction. To construct the inverse, first note that if neither of the exceptional cases have come into play in the construction of (λ', μ'), then one can apply the same steps used in inverting the map g_n of Theorem 4. To recognize when an exception has been used merely note whether, for some j, the last block of λ_j' contains a j with at least one other element. Distinguishing between exceptions (i) and (ii) is done by finding the position of j in μ_j' (non-zero or zero respectively). Now it is a simple matter to construct a step-by-step inverse for the exception. □

4. Graphical interpretations

Wilf [8, 9] has shown that various recursive structures can be modeled using paths through a labeled digraph (directed graph). He used this interpretation to provide a unified method for sequencing, ranking and selecting combinatorial objects. These ideas will also shed light on the injective proofs of the preceding section.

Let our digraph have as vertices the set of all integral lattice points (x, y) of the Cartesian plane satisfying $x \geqslant 0$ and $0 \leqslant y \leqslant x$. Given a triangular array t_{nk} satisfying the recurrence (3), we direct c_{nk} labeled arcs from vertex (n, k) to vertex $(n-1, k-1)$ and d_{nk} labeled arcs from (n, k) to $(n-1, k)$. These arcs are called *diagonal* and *horizontal* arcs respectively. The boundary vertices $(n, 0)$ (respectively (n, n)) are only assigned horizontal (respectively diagonal) outgoing arcs. If $t_{00} = 1$, then it is clear that t_{nk} counts the number of paths from (n, k) to $(0, 0)$. We should note that the labeling of arcs in our examples will differ slightly from that found in [9] but this will have no effect on the overall content.

For the binomial coefficients, all horizontal arcs are labeled 0 and the diagonal arc leaving (n, k) has label n. The set formed by the non-zero labels on a path

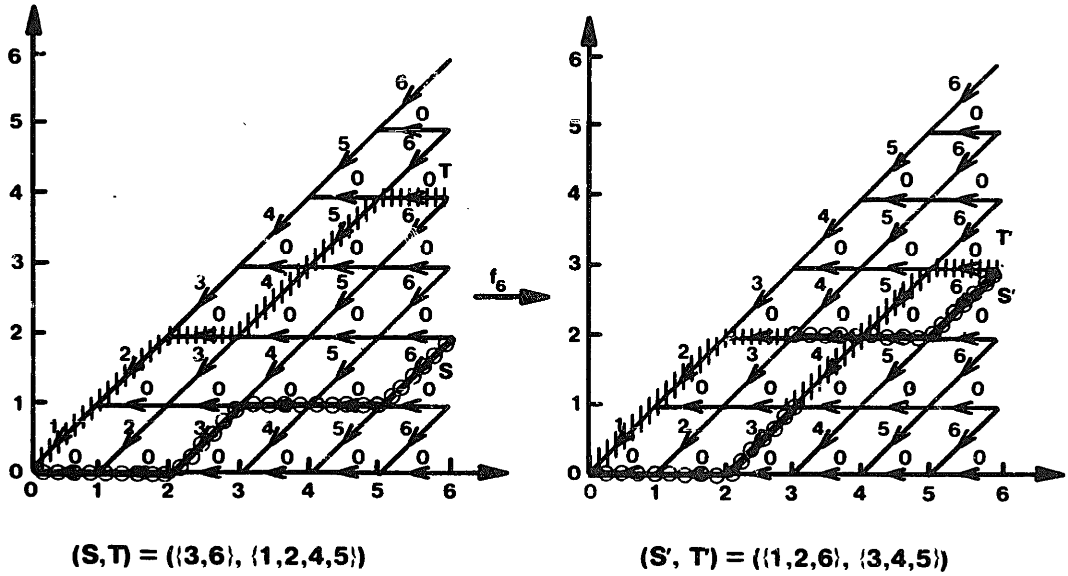

Fig. 3. Graphical interpretation of $f_6: B_{624} \to B_{633}$.

from (n, k) to $(0, 0)$ is thus a k-element subset of $[n]$. The map f_n of Section 3 has a straight-forward interpretation in terms of such paths.

A pair $(S, T) \in B_{nk-1k+1}$ is a pair of paths from $(n, k-1)$ and $(n, k+1)$ back to $(0, 0)$. The index i represents the first time that the T-path is exactly one unit above the S-path. Finally $(S', T') = f_n(S, T)$ represents the result of moving the portion of the S-path between $x = i$ and $x = n$ one unit up, and moving the corresponding piece of the T-path down in the obvious way. Notice that S' refers to the path which agrees with T for $x \leqslant i$ and is a translate of S for $x \geqslant i$, vice-versa for T'. An example is given in Fig. 3.

The digraph for permutations has $n-1$ horizontal arcs and one diagonal arc emanating from (n, k). (We can disregard any arcs from points $(n, 0)$, $n \geqslant 1$, since it is impossible to reach $(0, 0)$ from such vertices.) Label these arcs 1 through n, with the sole arc to $(n-1, k-1)$ receiving the label n as in Fig. 4. Thus traveling along arc s represents passing from $\pi = \pi_n$ to π_{n-1} by eliminating n from the sth position. Index i plays the same role as before and the paths are exchanged via a pair of label-preserving translations. Figure 5 illustrates the example of $g_5: S_{524} \to S_{533}$ worked out in the previous section.

In the partition digraph the outdegree of (n, k) is $k+1$ with a unique diagonal arc labeled 0 and the remaining horizontal arcs labeled 1 to k, see Fig. 6. The reason why the injection is more complicated in this setting is now apparent: vertical label-preserving translations are not always possible. However, the reader should have no trouble formulating the exceptional cases (i) and (ii) in graphical terms, so this is left as an exercise.

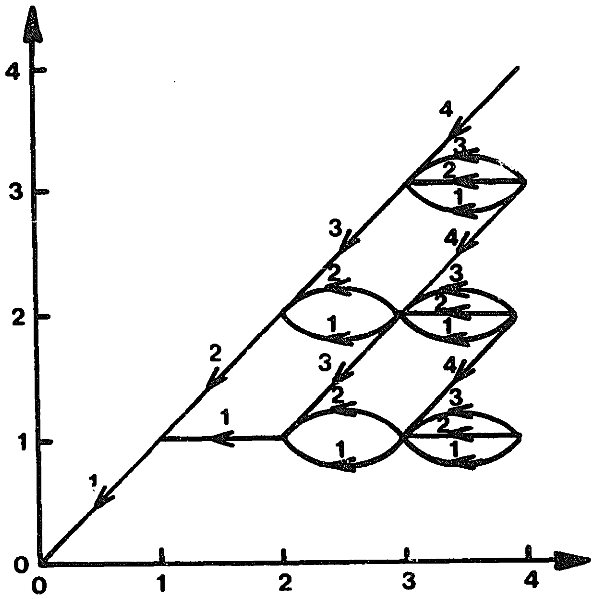

Fig. 4. Digraph for permutations.

5. Concluding remarks

Injective proofs of log concavity results have appeared elsewhere in the literature. Daykin, Daykin and Paterson [3] used explicit injections to prove that various sequences counting order preserving maps of posets were log concave. Building on an idea of Bhatt and Leiserson [1], Lynne Butler [2] has investigated

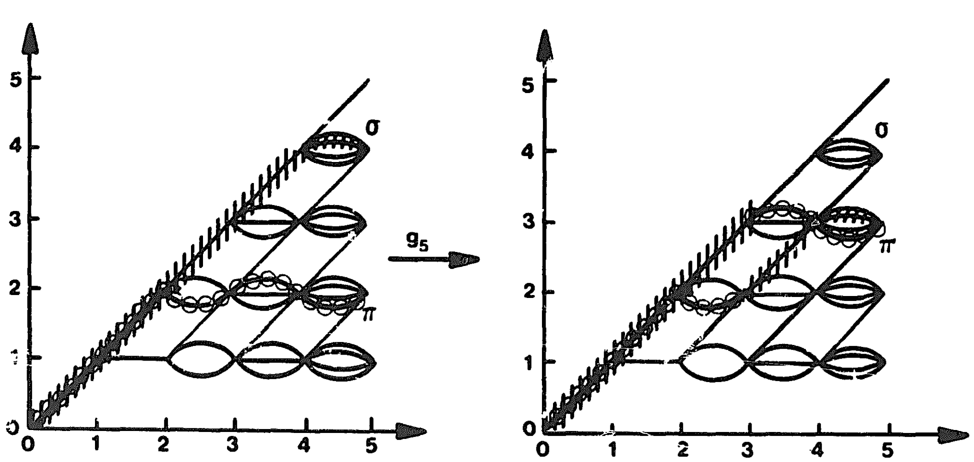

$(\pi,\sigma) = ((1,5,3)(2,4), (1)(2)(3,5)(4))$ $(\pi', \sigma') = ((1,5)(2)(3,4), (1,3)(2,5)(4))$

Fig. 5. Graphical interpretation of $g_5 : S_{524} \to S_{533}$. Note: labels and arrows have been suppressed for readability.

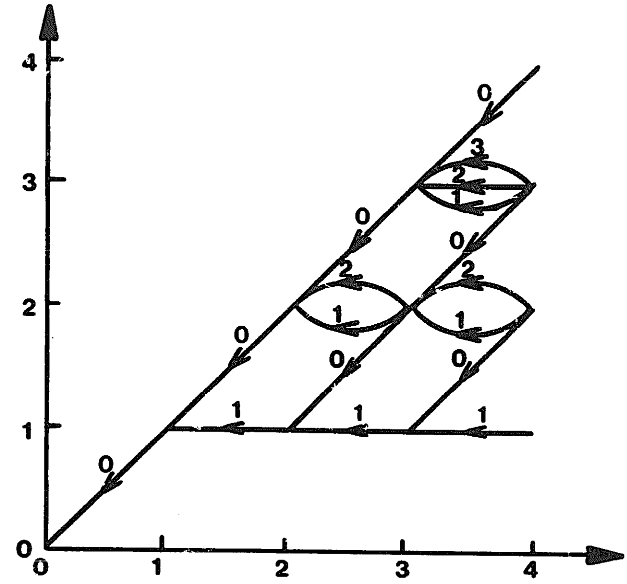

Fig. 6. Digraph for partitions.

log concavity of the Gaussian polynomials

$$\begin{bmatrix} n \\ k \end{bmatrix}_q = \frac{(q^n-1)(q^{n-1}-1)\cdots(q-1)}{(q^k-1)(q^{k-1}-1)\cdots(q-1)(q^{n-k}-1)(q^{n-k-1}-1)\cdots(q-1)}.$$

In particular, she showed combinatorially that

$$\begin{bmatrix} n \\ k \end{bmatrix}_q^2 - q\begin{bmatrix} n \\ k-1 \end{bmatrix}_q \begin{bmatrix} n \\ k+1 \end{bmatrix}_q$$

has nonnegative coefficients and when $q = 1$ her construction specializes to the map f_n of Theorem 3.

The graphical interpretation in Section 4 bears a resemblence to work of Gessel and Viennot [4] on combinatorial proofs of determinantal identities. By assigning signs to an n-tuple of paths, the exchange of a path pair can be interpreted as cancellation in a determinant and many beautiful results follow. This leads one to wonder whether there are other settings where path interchanges can be brought into play.

References

[1] S.N. Bhatt and C.E. Leiserson, How to assemble tree machines, *Proc. ACM Symposium on Theory of Computing* 14 (1982) 77–84.

[2] L.M. Butler, A unimodality result in the enumeration of subgroups of a finite abelian group, preprint.

[3] D.E. Daykin, J.W. Daykin and M.S. Paterson, On log concavity for order-preserving maps of partial orders, Discrete Math. 50 (1984) 221–226.
[4] I. Gessel and G. Viennot, Binomial determinants, paths and hook length formulae, Adv. in Math. 58 (1985) 300–321.
[5] R.L. Graham, Linear extensions of partial orders and the FKG inequality, in: I. Rival, ed., Ordered Sets (Reidel, Dordrecht, 1982) 213–236.
[6] P.D. Seymour and D.J.A. Welsh, Combinatorial applications of an inequality from statistical mechanics, Math. Proc. Comb. Phil. Soc. 77 (1975) 485–495.
[7] R.P. Stanley, Two combinatorial applications of the Aleksandrov–Fenchel inequalities, J. Combin. Theory Ser. A 31 (1981) 56–65.
[8] H.S. Wilf, A unified setting for sequencing, ranking and selection algorithms for combinatorial objects, Adv. in Math. 24 (1977) 281–291.
[9] H.S. Wilf, A unified setting for selection algorithms (II), Ann. Discrete Math. 2 (1978) 135–148.