Combinatorics, Probability and Computing (2007) 16, 29–41. © 2006 Cambridge University Press
doi:10.1017/S0963548306007772 First published online 14 August 2006 Printed in the United Kingdom

Positive Influence and Negative Dependence

DEVDATT DUBHASHI,[1†] JOHAN JONASSON[2‡]
and DESH RANJAN[3§¶]

[1]Department of Computing Science, Chalmers University of Technology, SE 412 96 Göteborg, Sweden
(e-mail: dubhasi@cs.chalmers.se)

[2]Department of Mathematics, Chalmers University of Technoloy, SE 412 96 Göteborg, Sweden
(e-mail: jonasson@math.chalmers.se)

[3]Department of Computer Science, New Mexico State University, Las Cruces, New Mexico 88003, USA
(e-mail: dranjan@cs.nmsu.edu)

Received 30 June 2004; revised 6 July 2005

We study negative dependence properties of a sampling process due to Srinivasan to produce distributions on level sets with given marginals. We give a simple proof that the distribution satisfies negative association. We also show that under a linear match schedule it satisfies the stronger condition of conditional negative association via a non-trivial application of the Feder–Mihail theorem. This method involves the notion of a variable of positive influence. We give some results and related counter-examples which might shed some light on its role in a theory of negative dependence.

1. Introduction

Pemantle [14] points out that there is a need for a theory of negatively dependent events analogous to the existing natural and useful theory of positively dependent events. He gives a survey of various notions of negative dependence, three of which are reviewed below, and points out several applications of consequences of one of these concepts: negative association. These examples include natural stochastic processes and structures such as the uniform random spanning tree [12], the simple exclusion process [10], the random cluster model and occupancy numbers of competing bins [3]. Negative dependence may also

† Work done partly while the author was visiting BRICS.
‡ Work supported by the Swedish Research Council.
§ Work done partly while the author was visiting the Max–Planck–Institut für Informatik and BRICS.
¶ Work partially supported by the ESPRIT Basic Research Actions Program of the EC under contract No. 7141 (project ALCOM II).

https://doi.org/10.1017/S0963548306007772 Published online by Cambridge University Press

be used to obtain information on the distribution of functionals such as the sum of the variables involved (or more generally, a nondecreasing function). Newman [13] shows that under negative dependence, one obtains a Central Limit Theorem (CLT) for stationary sequences of random variables. Dubhashi and Ranjan [3] show that under negative dependence conditions, one can employ classical tools of concentration of measure such as the Chernoff–Hoeffding (CH) bounds and related martingale inequalities.

In this paper, we contribute to the incipient theory of negative dependence by adding to the stock of examples of negative dependence. This class of examples involves distributions on level sets with given marginals that is introduced in the next subsection. We also add to the body of available techniques, by giving a non-trivial application of the Feder–Mihail theorem [5] to this class of examples. We now describe briefly an application of level set distributions to the design and analysis of randomized rounding algorithms for hard combinatorial optimization problems.

1.1. Randomized rounding and distributions on level sets

Recent results in approximation algorithms have suggested new applications of negative dependence via the very successful technique of randomized rounding [16]. Various hard combinatorial optimization problems can often be solved approximately via the following steps.

(1) Formulate the problem as an integer linear program (ILP) involving decision variables taking binary values.
(2) Relax the ILP to the corresponding linear program (LP) and solve it optimally. (There are several polynomial time algorithms to do this.)
(3) Round the fractional values of the optimal LP solution by a randomized rounding rule that uses the optimal LP solution as probabilities: $X_i = 1$ with probability x_i^* and $X_i = 0$ with probability $1 - x_i^*$.

Often the ILP involves *hard cardinality constraints* corresponding to sharing of resources, of the form

$$\sum_{i \in I} x_i = k.$$

If the randomized rounding step is performed independently for all variables, it is very likely that such constraints are violated. Srinivasan [17] uses this as motivation to introduce the problem of generating distributions on level sets with given marginals. A *level set* distribution is a distribution on $0, 1$ variables which is concentrated on the event that their sum is a fixed value k. The randomized rounding technique motivates the need to generate distributions on n binary random variables such that:

P1 it is a level set distribution, *i.e.*, concentrated on the event $\sum_i X_i = k$ for some integer $k > 0$,

P2 the marginals are given: $P[X_i = 1] = p_i, i \in [n]$.

In addition, for applying standard tools of concentration of measure such as the CH bounds, we need that:

P3 the distribution is negatively dependent.

https://doi.org/10.1017/S0963548306007772 Published online by Cambridge University Press

Srinivasan [17] gave a simple process that generated such distributions and gave a proof of a negative dependence property that is needed in his applications. This yields improved approximation algorithms for problems such as (a) low congestion multi-path routing, (b) maximum coverage set cover, (c) group Steiner tree problems. Gandhi, Khuller, Parthasarathy and Srinivasan [7] extended the method, combining it with a pipage rounding technique due to Ageev and Sviridenko, to develop new randomized rounding approaches to obtain fractional vectors on edge sets of bipartite graphs. This yielded richer random graph models for graphs with a given degree sequence, improved approximation algorithms for (a) throughput maximization in broadcast scheduling, (b) delay minimization in broadcast scheduling, (c) capacitated vertex cover, and (d) fair scheduling of jobs on unrelated parallel machines.

In this paper, we give the following results. First we give a very simple proof of a stronger negative dependence property of Srinivasan's distribution on level sets: namely, we show that it satisfies negative association. Second, we give a proof of a (potentially) even stronger negative dependence condition, namely, conditional negative association via a non-trivial application of the Feder–Mihail theorem [5]. This theorem can be used to establish negative dependence, somewhat paradoxically, via a notion of a variable of *positive influence*. This concept could possibly play an important role in a theory of negative dependence, but at the moment it is not sufficiently well understood. A very basic and natural question in the general theory is: *For which combinations of a function and a distribution does a variable of positive influence exist?* Here we provide some first simple results related to this question and we give some general results and counter-examples that might shed more light on its role in a theory of negative dependence.

In Section 2 we give the definitions of various concepts of negative dependence and positive influence. In Section 3 we give general results and counter-examples on the existence of a variable of positive influence. In Section 4 we establish negative and conditional negative association for various level set distributions using the Feder–Mihail theorem. We give a simple proof of negative association for Srinivasans's level set distribution and then a non-trivial application of the Feder–Mihail theorem to establish conditional negative association.

2. Negative dependence, positive influence and the Feder–Mihail theorem

Intuitively, a set of variables is negatively dependent if, conditioned on a subset of the variables taking 'high' values, it is more likely that the complementary subset of variables takes on 'low' values. This notion can be formalized in a number of ways [3], out of which, we focus on the following two.

Definition 1. (Negative and conditional negative association) A set of random variables $X_1, \ldots, X_n$ is *negatively associated* (NA) if, for very two disjoint subsets $I, J \subseteq [n]$ of variables, and all nondecreasing functions f, g,

$$E[f(X_i, i \in I)g(X_j, j \in J)] \leqslant E[f(X_i, i \in I)]E[g(X_j, j \in J)].$$

https://doi.org/10.1017/S0963548306007772 Published online by Cambridge University Press

The variables are *conditionally negatively associated* (CNA) if, conditioned on setting any subset of the variables to fixed values, the resulting conditional distribution on the remaining set of variables is NA.

These definitions are much stronger than the simple notion of *pairwise negative correlation*: $E[X_iX_j] \leqslant E[X_i]E[X_j]$ for distinct $i, j \in [n]$. The Feder–Mihail theorem yields the surprising conclusion that for binary random variables this weak notion of negative dependence yields the very strong notion CNA whenever there is a variable of *positive influence*. Intuitively, a variable has a positive influence on a function if one expects the value of the function to increase if the value of the variable is increased, assuming that the other variables are chosen at random. To be more precise, positive influence is defined as follows [3].

Definition 2. (Positive influence) Let $X_1, \ldots, X_n$ be real-valued random variables and let μ be any probability distribution on $X_1, \ldots, X_n$. Let f be any real-valued function of $X_1, \ldots, X_n$. We say that X_j is a *positive influence variable* for (f, μ) if $E_\mu[f(X_1, \ldots, X_n)|X_j = t]$ is a nondecreasing function of t.

Other notions of influence of variables on functions have been defined and used before. In particular, it is instructive to compare our definition of positive influence with that of influence in Bourgain, Kahn, Kalai, Katznelson and Linial [1]. The influence of a variable X_j on a function $f(X_1, \ldots, X_n)$ there is defined to be the probability that the univariate function of X_j that results by picking all the other X_is at random is a non-constant.

For simplification of notation, when it is clear what distribution we are talking about, or we are talking about an arbitrary fixed distribution, we will often omit μ from the subscript in the expectation. Also, for a family of objects $\{x_i : i \in I\}$ and a set $A \subseteq I$, we will write x_A for the sub-family $\{x_i : i \in A\}$.

Theorem 2.1. (Feder–Mihail) *Let $X_1, \ldots, X_n$ be binary random variables such that for any $A \subseteq [n]$ and any $x_A \in \{0,1\}^A$, the conditional distribution of $X_{[n]\setminus A}$ given that $X_A = x_A$, satisfies the following two conditions.*

Positive influence. *For any nondecreasing function f, there is a variable of positive influence.*
Pairwise negative correlation. *Any two variables are negatively correlated.*

Then the variables $X_1, \ldots, X_n$ satisfy CNA (conditional negative association).

3. Positive influence: General results and counter-examples

Proposition 3.1. *Let $X_1, X_2, \ldots, X_n$ be binary random variables. A sufficient condition for the existence of a positive influence variable is*

$$\operatorname{cov}\left(f, \sum_i X_i\right) \geqslant 0. \tag{3.1}$$

https://doi.org/10.1017/S0963548306007772 Published online by Cambridge University Press

Proof. By linearity of expectation,

$$0 \leqslant \operatorname{cov}\left(f, \sum_i X_i\right) = \sum_i \operatorname{cov}(f, X_i).$$

Hence there exists an $i \in [n]$ such that $\operatorname{cov}(f, X_i) \geqslant 0$. Since X_i is binary, i is a variable of positive influence. □

One of the most useful cases where condition (3.1) holds is for distributions on level sets.

Corollary 3.2. *Suppose that the distribution μ of the binary X_is is concentrated on a fixed level $k \geqslant 0$, i.e., $\sum_i X_i = k$. Then there is a variable of positive influence.*

Another set of circumstances in which positive influence holds for a pair (μ, f) is if either μ or f is symmetric.

Proposition 3.3. *Let f be symmetric, i.e., $f(a_1, \ldots, a_n) = f(a_{\sigma(1)}, \ldots, a_{\sigma(n)})$ for any permutation σ, and nondecreasing and let the X_is be binary. Then there is a variable of positive influence for any μ.*

Proof. Since f is symmetric and the X_is are binary we have that f only depends on $\sum_i X_i$. Since f is increasing it follows that f is an increasing function of $\sum_i X_i$. Thus Chebyshev's inequality entails that $\operatorname{cov}(f, \sum_i X_i) \geqslant 0$. Now apply the above proposition. □

Proposition 3.4. *Let f be nondecreasing, let the X_is be binary and let μ be symmetric (exchangeable), i.e., $\mu(X_1 = a_1, \ldots, X_n = a_n) = \mu(X_1 = a_{\sigma(1)}, \ldots, X_n = a_{\sigma(n)})$ for any permutation σ. Then there is a variable of positive influence for any μ.*

Proof. This will follow from Proposition 3.1 if it is established that, given $\sum_i X_i = k$, the distribution of $(X_1, X_2, \ldots, X_n)$ is stochastically increasing in k.

However, the symmetry of μ implies that this conditional distribution is uniform over all k-subsets of $[n]$. This trivially implies the desired property. □

3.1. Counter-example: What one cannot hope for

One might begin thinking naively that if a function $f(X_1, \ldots, X_n)$ is nondecreasing then at least one of $X_1, \ldots, X_n$ must have a positive influence on the function f no matter what the probability distribution on $X_1, \ldots, X_n$ is. This is quickly seen to be false. We provide a counterexample here that shows that this is not true even for two 3-valued symmetric random variables and even if the functions are additionally restricted to be binary and symmetric.

https://doi.org/10.1017/S0963548306007772 Published online by Cambridge University Press

Consider the following probability distribution on X, Y that takes values in $\{0,1,2\}$ with

$$\Pr[X=0, Y=2] = p_2,$$
$$\Pr[X=1, Y=1] = p_1,$$
$$\Pr[X=2, Y=0] = p_2,$$

where $p_1, p_2 > 0$ and $2p_2 + p_1 = 1$. Consider the nondecreasing binary symmetric function $f(X,Y)$ defined as follows:

$$f(X,Y) = \begin{cases} 1 & \text{if } X^2 + Y^2 > 3, \\ 0 & \text{otherwise.} \end{cases}$$

Then neither X nor Y has a positive influence on f as

$$E[f(X,Y)|X=0] = E[f(X,Y)|Y=0] = 1$$

whereas

$$E[f(X,Y)|X=1] = E[f(X,Y)|Y=1] = 0.$$

3.2. Counter-example: Binary variables

Even when the variables are restricted to be binary, there exist probability distributions for which there are nondecreasing binary functions that do not have any positive influence variable. This is evidenced by the example below. Consider the distribution μ on six random variables:

a_1	a_2	a_3	b_1	b_2	b_3	$\mu(X_1=a_1, X_2=a_2, X_3=a_3, Y_1=b_1, Y_2=b_2, Y_3=b_3)$
0	0	0	1	1	1	1/5
0	1	1	0	1	1	1/5
1	0	1	1	0	1	1/5
1	1	0	1	1	0	1/5
1	1	1	0	0	0	1/5

Consider the binary function $f(X_1, X_2, X_3, Y_1, Y_2, Y_3)$ defined as follows:

$$f(X_1, X_2, X_3, Y_1, Y_2, Y_3) = \begin{cases} 1 & \text{if } (X_1+X_2+X_3)^2 + (Y_1+Y_2+Y_3)^2 > 8, \\ 0 & \text{otherwise.} \end{cases}$$

Then, for $i = 1,2,3$, $E[f|X_i=0] = E[f|Y_i=0] = \frac{1}{2}$ whereas $E[f|X_i=1] = E[f|Y_i=1] = \frac{1}{3}$. Hence, none of the variables is a positive influence variable for f under μ. Notice that neither μ nor f is symmetric.

4. Negative dependence via the Feder–Mihail theorem

We give four applications of the Feder–Mihail theorem to distributions on binary random variables that are concentrated on level sets, *i.e.*, where the number of ones is fixed. Here

https://doi.org/10.1017/S0963548306007772 Published online by Cambridge University Press

the positive influence property holds automatically, so it remains to check the pairwise negative correlation property to establish CNA.

4.1. Example 1: Fermi–Dirac statistics

An (n,k)-Fermi–Dirac statistic is a distribution on n binary variables obtained by picking, from a set of n objects, a k-subset uniformly at random. Conditioned on setting some of the variables, the resulting distribution is just another Fermi–Dirac distribution, and negative covariance holds for any Fermi–Dirac distribution because of the simple log-concavity property of the binomial-coefficients: $\binom{n}{k-2}\binom{n}{k} \leqslant \left(\binom{n}{k-1}\right)^2$. Thus we get CNA for Fermi–Dirac statistics. A generalization of Fermi–Dirac statistics is as follows.

4.2. Example 2: Conditional Poisson sampling

Suppose we want to sample from a population of n individuals in such a way that the probability that individual $i \in [n]$ is included in the sample equals a prescribed value p_i. Assume the $\sum_i p_i = k$ for some $k \in [n-1]$. In other words we want to produce a family $X_1, X_2, \ldots, X_n$ of binary random variables such that $P(X_i = 1) = p_i$ for all $i \in [n]$.

Poisson sampling refers to the method of doing this by simply letting each individual be included with its prescribed inclusion probability independently of the other individuals, *i.e.*, taking the X_is to be independent. However, this will produce a sample of random size. If we insist on having exactly k individuals in our sample we can condition on the event $\sum_i X_i = k$. Then we get *conditional Poisson sampling* (CPS). As a special case, with $p_i = k/n$ for every i, we get the (n,k)-Fermi–Dirac statistic.

CPS has many nice properties. For example, it has maximal entropy in the class of all sampling processes with fixed sample size k and the same inclusion probabilities: see [9]. The main drawback is that the conditional inclusion probabilities are in general not equal to the unconditional ones and in order to find out what unconditional inclusion probabilities to use, one has to solve an overwhelming non-linear system of equations.

We now adopt the Feder–Mihail theorem to show that CPS satisfies CNA. Note that a CPS sample conditioned on the value of some of the X_is is another CPS sample with another sample size. Therefore, by the Feder–Mihail theorem, all we need to show is pairwise negative correlations for CPS sampling, *i.e.*, that for $i \neq j$, $P(X_i = 1 | X_j = 1) \leqslant P(X_i = 1 | X_j = 0)$. But the distributions of the X_ls given $X_j = 1$ and $X_j = 0$, respectively, are those of two CPS samples that are identical, except that the sample size is $k-1$ in the first case and k in the latter. Therefore it suffices to establish that the distribution of a CPS sample is stochastically increasing in the sample size. This turns out to be surprisingly difficult by elementary methods and the most efficient way that we know of is by a coupling argument. We will shortly describe it here and refer to [9] for more details.

The concept of coupling makes use of the following observation for two probability distributions μ and ν on the same partially ordered space $\mathscr{P}$. If one can define a pair (X, Y) of $\mathscr{P}$-valued random variables, defined on a common probability space, such that X has distribution μ, Y has distribution ν and $X \leqslant Y$ a.s., then μ is stochastically smaller than ν; $\mu \leqslant_d \nu$. The pair (X, Y) is called a coupling of μ and ν that witnesses the stochastic domination $\mu \leqslant_d \nu$. In fact $\mu \leqslant_d \nu$ if and only if there exists such a coupling; this is Strassen's famous theorem (see, *e.g.*, [11, Section IV.1]).

https://doi.org/10.1017/S0963548306007772 Published online by Cambridge University Press

In the present situation we first note that the CPS distribution is the stationary distribution of the aperiodic irreducible Markov chain on all subsets of size k with the following behaviour.

Let S_t denote the state of the chain at time t. Given S_t, the state S_{t+1} is generated by (i) picking an individual i from S_t uniformly at random, (ii) picking an individual j from the whole population according to the probabilities $\{\theta_l/c\}$ where $\theta_l = p_l/(1-p_l)$ and $c = \sum_{m=1}^{n} \theta_m$.

If $j \in S_t$, then let $S_{t+1} = S_t$ otherwise replace i with j, *i.e.*, put $S_{t+1} = (S_t \setminus \{i\}) \cup \{j\}$. That this Markov chain indeed has the CPS sampling distribution, given that the sample size is k, is straightforward to verify using the detailed balance equations.

Second, we couple the Markov chain with a corresponding one, $\{S'_t\}$, for sample size $k-1$ starting from a state $S'_0 \subset S_0$. Since the p_is are the same for the second chain it is easily seen that the transitions of this chain can then be coupled to the ones of the $\{S_t\}$-chain in such a way that $S'_t \subset S_t$ for all t. Hence we also have the desired stochastic inequality for the two stationary distributions.

4.3. Example 3: The uniform spanning tree

In the examples above, the conditional pairwise negative correlation property was rather easy to check. However, sometimes this can be quite non-trivial. One such example is the *uniform spanning tree*, *i.e.*, a tree T chosen uniformly at random from all possible spanning trees of a graph $G = (V, E)$. Let $X_e, e \in E$ be the indicator variables corresponding to whether an edge e is picked in T. Since a tree has exactly $|V| - 1$ edges, this is a level set distribution and the positive influence property comes for free. To establish CNA, we need to verify the conditional pairwise negative correlation property. However, the only known proof of the conditional pairwise negative correlation property in this case is, very surprisingly, via a transfer-impedance theorem of Burton and Pemantle [2] that exploits an analogy with electrical networks.

5. Srinivasan's sampling process

Let us now look at what we shall call *Srinivasan's sampling process* (SSP). SSP works by *fixing* the values of variables to 0 or 1 one by one. It does so by repeatedly applying the following rules to pairs of unfixed variables in some pre-defined order. As before, p_i denotes the inclusion probability for individual i. Note that $\sum_i p_i = k$, the desired sample size.

Initially start with each variable *unfixed*, and with value $X_i = p_i$. At each step, pick two unfixed variables X_i and X_j. There are now two cases.

Case 1: $p_i + p_j \leqslant 1$. With probability $p_i/(p_i + p_j)$, let $X_j = 0$ and $X_i = p_i + p_j$. The variable X_j is *fixed*. With the complementary probability $p_j/(p_i + p_j)$ set $X_i = 0$ and $X_j = p_i + p_j$. In this case, the variable X_i is fixed.

Case 2: $p_i + p_j > 1$. With probability $(1 - p_j)/(2 - p_i - p_j)$, let $X_i = 1$ and $X_j = (p_i + p_j) - 1$. The variable X_i is fixed. With the complementary probability $(1 - p_i)/(2 - p_i - p_j)$ put $X_j = 1$ and $X_i = (p_i + p_j) - 1$ and the variable X_j is fixed.

https://doi.org/10.1017/S0963548306007772 Published online by Cambridge University Press

One of the main virtues of SSP is that it is extremely easy to describe and to implement and still, as we shall now see, it satisfies strong properties of negative dependence.

It should be noted that the exact probabilities specified above with which the probability mass is transferred from one variable to the other, are designed to make the marginal probability that i is included in the sample equal to p_i. However, this is not essential for the proof of negative dependence. For negative dependence purposes, we can view the SSP as carrying out a series of *matches* between a pair of individual unfixed variables X_i and X_j. Each such match m has a *winner* and a *loser* whose identities are determined by a binary random variable Z_m whose distribution (independent of every other match) is given by some predetermined rule involving the values of the contesting variables. Each such match results in a transfer of probability mass from the loser to the winner. The amount transferred depends on whether or not $p_i + p_j$ exceeds 1, as described above. If $p_i + p_j \geqslant 1$ we refer to a *good* match, and otherwise to a *bad* match. Note that in a good match, the winner will be automatically qualified for a place in the final sample and the loser will go on playing for a place, whereas in a bad match the loser will be once and for all ruled out from the sample, and the winner will go on to play for a place. Let us call such a process of setting the variables a *generalized Srinivasan Sampling Process* (GSSP).

Theorem 5.1. *The distribution produced by applying the GSSP to the variables in any predetermined order produces a distribution satisfying NA.*

Proof. We use the conditional covariance formula: if X, Y and Z are three random variables, then

$$\operatorname{cov}(X, Y) = \mathbf{E}[\operatorname{cov}(X, Y|Z)] + \operatorname{cov}(\mathbf{E}[X|Z], \mathbf{E}[Y|Z]).$$

We will use induction on n, the number of variables (*i.e.*, the number of individuals in the population). First note that for $n = 2$, NA is trivial, so we can focus on the induction step.

Let f and g be two increasing functions depending on disjoint sets of variables and let Z be the outcome of the match between the first two individuals, which we may without loss of generality call 1 and 2. That is, let Z be the indicator that 1 wins that match.

By the conditional covariance formula

$$\operatorname{cov}(f, g) = \mathbf{E}[\operatorname{cov}(f, g|Z)] + \operatorname{cov}(\mathbf{E}[f|Z], \mathbf{E}[g|Z]).$$

Given Z we face a GSSP on $n-1$ individuals and so, by the induction hypothesis, the first term on the right is non-positive.

For the second term note that the distribution of $X_3, X_4, \ldots, X_n$ is independent of Z. Hence the second term will be 0 unless 1 and 2 are in the support of different functions, so assume without loss of generality that 1 is in the support of f and 2 is in the support of g.

However, then $\mathbf{E}[f|Z]$ is increasing in Z and $\mathbf{E}[g|Z]$ is decreasing in Z (here we also use that $X_3, \ldots, X_n$ are independent of Z), and so the second term is also non-positive. This establishes NA. □

https://doi.org/10.1017/S0963548306007772 Published online by Cambridge University Press

To establish CNA is less straightforward but we are lucky since the Feder–Mihail theorem tells us that it is sufficient to prove conditional pairwise negative correlation. However, we will still not be able to do this in full generality; we will need to restrict to the case of a *linear match schedule*, *i.e.*, in the first match 1 plays 2, in the second match the variable from match 1 that remains unfixed plays 3, in the third match the variable that remains unfixed from match 2 plays 4, *etc.*

We need to prove that

$$\mathrm{cov}(X_a, X_b | X_A = x_A) \leqslant 0$$

for any $A \subset [n]$, any $x_A \in \{0,1\}^A$, and any pair of individuals a and b ($a < b$). We will do this by showing that

$$P(X_b = 1 | X_a = 1, X_A = x_A) \leqslant P(X_b = 1 | X_a = 0, X_A = x_A).$$

Assume without loss of generality that $\{a, b\} \cap A = \emptyset$; the result is otherwise trivial.

Assign to each individual $i \in [n]$ for a given conditional SSP a *status*: we say that i has status 1 if it has been conditioned that $X_i = 1$, 0 if it has been conditioned that $X_i = 0$, and U (for 'unknown') if i is not in the set A of individuals conditioned on. We linearly order the set $\{1, 0, U\}$ by putting $0 \leqslant U \leqslant 1$.

Let us now start comparing the conditional SSPs where, on the one hand, we condition on the event $\{X_a = 1, X_A = x_A\}$ and, on the other, on the event $\{X_a = 0, X_A = x_A\}$. For convenience we denote the first process 'Process I' and the latter 'Process II'. We want to show that b has at least as great a chance of qualifying for the sample in Process II as in Process I. We will do this by coupling the two processes so that, as soon as b is included in the sample of Process I, then b is also included in the sample of Process II. Start by letting the two processes evolve independently for the first $a-1$ matches. Assume that match $a-1$ (*i.e.*, the first match a plays) is bad. (The case when this match is good is analogously treated and is left to the reader.) In Process I, where a has status 1, a must necessarily win this match. In particular, this means that the opponent of $a+1$ in match number a will have status 1. (Conditioned on the opponent of $a+1$ in match a all that matters as far as far as the rest of the process is concerned is the status of that opponent.) In Process II a may or may not win match $a-1$, but in any case the opponent of $a+1$ in match a will trivially have a status less than or equal to 1, in particular, less than or equal to the status of the opponent of $a+1$ in Process I. For $j = a+1, a+2, \ldots, b$, let V_j denote the status of the opponent of j in its first match, *i.e.*, in match $j-1$, in Process I and let $W_{a+1}, W_{a+2}, \ldots, W_b$ denote the corresponding variables for Process II; we just showed that $V_{a+1} \geqslant W_{a+1}$.

We now claim that we can couple the two processes so that $V_j \geqslant W_j$ for all j. The claim is trivially true if $b = a+1$, so assume that $b > a+1$. We will prove the claim by showing that, for $a+1 \leqslant j < b$, $V_j \geqslant W_j$ implies that we can couple so that $V_{j+1} \geqslant W_{j+1}$. If $V_j = W_j$ then we can take $V_{j+1} = W_{j+1}$, so we may assume that $V_j > W_j$, *i.e.*, $(V_j, W_j) \in \{(U,0), (1,0), (1,U)\}$. We then face three cases depending on the status of j.

Case 1: j has status U. In this case $V_{j+1} \geqslant W_{j+1}$ whatever we do.

https://doi.org/10.1017/S0963548306007772 Published online by Cambridge University Press

Case 2: j has status 0. Then the same goes when $W_j = 0$. When $V_j = 1$ and $W_j = U$, then either match $j-1$ is good, in which case we will necessarily get $W_{j+1} = V_{j+1} = 0$, and if match $j-1$ is bad then we automatically get $V_{j+1} = 1 \geqslant W_{j+1}$.

Case 3: j has status 1. Here the analysis is analogous (but reversed): when $V_j = 1$ things work out whatever we do; when $W_j = 0$ and $V_j = U$, then if match $j-1$ is bad we necessarily get $V_{j+1} = W_{j+1} = 1$, and if match $j-1$ is good then we necessarily get $W_{j+1} = 0 \leqslant V_{j+1}$.

Now that we know that we can couple Process I and Process II so that $V_j \geqslant W_j$ for all j, we know in particular that we will have $V_b \geqslant W_b$. Now, the conditional probability that b wins match $b-1$ given that the status, S, of its opponent is s is decreasing in s. To see this we may as well assume that the opponent of b is $b-1$. Let Z_b be the indicator that b wins the match, assume $s \in \{0,1\}$ and apply Bayes' formula:

$$\Pr(Z_b = 1 | X_{b-1} = s) = \frac{\Pr(X_{b-1} = s | Z_b = 1)}{\Pr(X_{b-1} = s)} \Pr(Z_b = 1).$$

(Here Pr refers to the conditional distribution of $X_{b-1}, \ldots, X_n$ given $X_{A\setminus[b]} = x_{A\setminus[b]}$, which is the same for Process I as for Process II.) The ratio is trivially at most 1 when $s = 1$ and at least 1 when $s = 0$, and the claimed monotonicity follows as the case $s = U$ corresponds to not conditioning on X_{b-1} at all.

Thus we can couple the two processes so that b wins match $b-1$ for Process II as soon as the same thing happens for Process I. If match $b-1$ is bad, this means that either:

- b is in Process I ruled out from the sample whereas it still has a chance to qualify for the sample of Process II, in which case we simply let the two processes evolve independently for the remaining matches,
- b is ruled out from both samples, in which case we do the same thing, or
- b is still in the game for both processes, in which case we can let b meet the same fate for the two processes.

The case when match $b-1$ is good is analogous. We have shown the following.

Theorem 5.2. *The distribution produced by applying the GSSP to the variables in linear order produces a distribution satisfying CNA.*

6. Conclusion and open questions

The extra condition of linear match schedule was necessary for the proof of CNA for GSSP to work out, but we strongly believe that the result also holds for a general match schedule.

Conjecture 6.1. *GSSP satisfies CNA under any match schedule.*

It is clear that even though SSP has CNA it has some undesirably strong dependencies. For example, if $p_1 + p_2 < 1$ then 1 and 2 cannot both be in the resulting sample. One

https://doi.org/10.1017/S0963548306007772 Published online by Cambridge University Press

way to increase the entropy of the process is of course to give the individuals a random order before starting the SSP. Let us call the resulting process the *random order Srinivasan sampling process* (ROSSP). It is natural to believe that ROSSP also has CNA. However, we do not know how to prove this. All we can show here is pairwise negative correlations. The proof goes via the conditional covariance formula. Let R denote the random order of the individuals (*i.e.*, R is an element of S_n, chosen uniformly at random). Then

$$\operatorname{cov}(X_i, X_j) = \mathbf{E}[\operatorname{cov}(X_i, X_j|R)] + \operatorname{cov}(\mathbf{E}[X_i|R], \mathbf{E}[X_j|R]).$$

The first term is negative by CNA for SSP and the second term is 0 as the individual inclusion probabilities are not affected by R.

Note that in order to prove CNA it is sufficient, by the Feder–Mihail theorem, to prove conditional pairwise negative correlations and one could easily believe that the same argument goes for that too. However, the fact that we are conditioning means that the second term is no longer unaffected by R.

Conjecture 6.2. *ROSSP satisfies conditional negative association.*

The general problem of characterizing (f, ψ) pairs for which positive influence variables exist remains open. It would be interesting and useful if one could establish positive influence lemmas for other broad families of (f, ψ) pairs.

References

[1] Bourgain, J., Kahn, J., Kalai, G., Katznelson, Y. and Linial, N. (1992) The influence of variables in product spaces. *Israel J. Math.* **77** 55–64.
[2] Burton, R. and Pemantle, R. (1993) Local characteristics, entropy and limit theorems for spanning trees and domino tilings via transfer-impedances. *Ann. Probab.* **21** 1329–1371.
[3] Dubhashi, D. and Ranjan, D. (1998) Balls and bins: A study in negative dependence. *Random Struct. Alg.* **13** 99–124.
[4] Eaton, M. L. (1987) *Lectures on Topics in Probability Inequalities*, Vol. 35 of *CWI Tracts in Mathematics*, Amsterdam.
[5] Feder, T. and Mihail, M. (1992) Balanced matroids. In *Proc. 24th Annual STOC*, pp. 26–38.
[6] Fortuin, C. M. and Kasteleyn, P. W. and Ginibre, J. (1971) Correlation inequalities on some partially ordered sets. *Comm. Math. Phys.* **22** 89–103.
[7] Gandhi, R., Khuller, S., Parthasarathy, S. and Srinivasan, A. (2002) Dependent rounding in bipartite graphs. In *Proc. 43rd IEEE Symposium on Foundations of Computer Science* (FOCS), pp. 323–332.
[8] Joag-Dev, K. and Proschan, F. (1983) Negative association of random variables, with applications. *Ann. Statist.* **11** 286–295.
[9] Jonasson, J. and Nerman, O. (1997) On maximum entropy πps-sampling with fixed sample size. Report, Chalmers University of Technology.
[10] Liggett, T. (1977) *The Stochastic Evolution of Infinite Systems of Interacting Particles*, Vol. 598 of *Lecture Notes in Mathematics*, Springer.
[11] Lindvall, T. (1992) *Lectures on the Coupling Method*, Wiley.
[12] Lyons, R. and Peres, Y. (2005) *Probability on Networks and Trees*, Cambridge University Press, book to appear. `http://www.math.gatech.edu/~rdlyons/prbtree/prbtree.html`

https://doi.org/10.1017/S0963548306007772 Published online by Cambridge University Press

[13] Newman, C. (1982) Asymptotic independence and limit theorems for positively and negatively associated random variables. In *Inequalities in Statistics and Probability* (Y. L. Tong, ed.), Vol. 5 of *IMS Lecture Notes, Monograph Series*, pp. 127–140.
[14] Pemantle, R. (2000) Towards a theory of negative dependence. *J. Math. Phys.* **41** 1371–1390.
[15] Port, S. C. (1994) *Theoretical Probability for Applications*, Wiley.
[16] Raghavan, P. and Thompson, C. D. (1987) Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica* **7** 365–374.
[17] Srinivasan, A. (2001) Distributions on level sets with applications to approximation algorithms. In *Proc. 42nd IEEE Symposium on Foundations of Computer Science* (FOCS), pp. 588–597.

https://doi.org/10.1017/S0963548306007772 Published online by Cambridge University Press