

Bundling Attacks in Judgment Aggregation

Noga Alon¹ and Dvir Falik² and Reshef Meir³ and Moshe Tennenholtz⁴

¹Tel-Aviv University and Microsoft Research, Herzliya, nogaa@tau.ac.il

²Queen Mary, University of London, dvir.falik@gmail.com

³The Hebrew University of Jerusalem, Israel and Microsoft Research, Herzliya
reshef.meir@mail.huji.ac.il

⁴Technion-Israel Institute of Technology and Microsoft Research, Herzliya
moshet@microsoft.com

Abstract

We consider judgment aggregation over multiple independent issues, where the chairperson has her own opinion, and can try to bias the outcome by bundling several issues together. Since for each bundle judges must give a uniform answer on all issues, different partitions of the issues may result in an outcome that significantly differs from the “true”, issue-wise, decision.

We prove that the bundling problem faced by the chairperson, i.e. trying to bias the outcome towards her own opinion, is computationally difficult in the worst case. Then we study the probability that an effective bundling attack exists as the disparity between the opinions of the judges and the chair varies. We show that if every judge initially agrees with the chair on every issue with probability of at least $1/2$, then there is almost always a bundling attack (i.e. a partition) where the opinion of the chair on all issues is approved. Moreover, such a partition can be found efficiently. In contrast, when the probability is lower than $1/2$ then the chair cannot force her opinion using bundling even on a single issue.

Introduction

Consider a committee that has gathered to vote over several issues (say, applications for tax reductions by retailers). The chairperson running the meeting decides that issues will not be voted on sequentially or independently. Rather, she insists that the committee will hold a single voting round, thereby either approving or denying all of the applications.

As a result, members who only support the approval of some applications, may now prefer to approve all of them (or to deny all). The Ostrogorski Paradox presented in the next section shows that there are cases where there is a majority for denying every single application, yet by bundling the issues for a single voting round all of them will be approved (Levmore 1999). This example demonstrates that the chairperson setting the agenda can have significant power to bias the outcome towards her own opinion.

This power becomes even stronger if the chairperson can apply more intricate partitions of the issues to

“bundles”, coercing the voters to vote on each bundle separately. By carefully bundling and partitioning the issues, the chairperson may be able to achieve any outcome she has in mind. Throughout this work we assume the goal of the chairperson is to approve all issues.¹ We henceforth refer to such a partitioning as a *bundling manipulation* or *bundling attack*. A *perfect bundling* is a partition that guarantees the approval of all issues.

In this paper we study the limits of the chairperson’s manipulation power from two perspectives. First, we ask how hard is it to find a bundling attack given a preference profile, i.e. what is the worst case complexity of this problem. Second, we study the likelihood that such a partition exists when the preference profile is sampled from a parametrized distribution. The latter question is also tightly related to the *average case complexity* of finding a bundling attack.

Multi-issue voting usually raises conceptual and technical problems due to interdependencies between issues (Farquharson 1969; Lacy and Niou 2000; Xia, Lang, and Ying 2007). Even if voters try to be “truthful” it is not always clear what they should vote. We consider *independent judgment aggregation*, where the preferences of voters over issues are independent. In such a setting the behavior of a judge is straightforward: the judge will approve exactly those issues he supports. Further, if he treats all issues as equally important, then this behavior can be naturally extended to any given agenda (i.e. partition to bundles): the judge will approve bundles where he supports at least half the issues, and will reject the other bundles.

Other than the committee example given above, similar situations may occur in economic and political settings. Consider a vendor bundling together groups of features or products, when the decision what to buy is the responsibility of a board or committee representing the buyer. Thus the vendor may “trick” the buyer to buy all features/products, where none of them would be approved separately.

A more intricate example is a political party wrapping together multiple issues, thereby making voters choose whether to support it as a whole.

¹This is w.l.o.g. in most problems we consider.

Related work

Baumeister et al. (2012) considered several types of control by the chair in judgment aggregation, namely by adding, removing, or replacing some of the judges. Following Endriss et al. (2010a; 2010b), they assume an underlying agenda defined by a set of logical formulae. That is, the opinion of each agent over an issue may depend on his opinions over other issues on the agenda. Baumeister et al. study the computational complexity of these problems under a particular aggregation rule.

Conitzer, Lang and Xia (2009) studied several variations of agenda manipulation problems in voting with multiple binary issues. They considered issues that are interdependent, and thus setting the *order* by which issues are voted on can significantly change the outcome.

Partitioning was studied in voting scenarios, where the chair has the power to partition either the voters or the candidates into two distinct sets, thereby affecting the outcome (Bartholdi, Tovey, and Trick 1992; Hemaspaandra, Hemaspaandra, and Rothe 2007).

In all three settings described above, much of the conceptual and computational complexity stems from the intricate dependency between the issues. In our setting issues are independent, which corresponds to *premise-only* agendas preference aggregation (Endriss, Grandi, and Porello 2010b), or to *separable* preferences in multi-issue voting (Conitzer, Lang, and Xia 2009). Thus in the setting we focus on some of the previously studied problems become trivial (for example the order of issues is irrelevant), and the difficulty arises due to the rich set of possible partitions.

Bundling of products has been studied in the context of auctions (Palfrey 1983). In such settings the bidders can make individual decisions, and the usefulness of bundling to the auctioneer is derived from her uncertainty regarding bidders' valuations.

Our contribution

We show that finding whether a bundling attack exists is NP-hard under two different variations of the problem. First, it is hard to find whether a single bundle can be used to approve a large fraction of the issues. Second, it is hard to find a *perfect bundling*, i.e. a partition that results in approving all issues.² Interestingly, the first problem turns out to be related to the OPTIMAL-LOBBYING problem (Christian et al. 2007), albeit in a non-intuitive way.

We then focus on the frequency of successful bundling attacks, when each voter approves every issue (i.e. agrees with the chair) with independent probability p . Our main result shows that for $p = 1/2$ (and clearly for any larger value), a perfect bundling attack almost always exists. In contrast, when $p < 1/2$ the probability that even a single issue can be approved using bundling goes to zero (when the number of voters is not extremely small). As a corollary, finding whether

²Note that checking whether a perfect outcome can be attained with single bundle is trivial - just bundle all issues and check if all issues are approved.

a perfect bundling attack exists is easy *on average* (for any value of p).

Some proofs are omitted or replaced with a proof sketch due to lack of space, and appear in the full version of this paper.³

Preliminaries

We use bold characters to denote column vectors, as in $\mathbf{a} = (a_1, a_2, \dots)$, upper case letters A, B, C, \dots to denote sets, and X, Y, Z to denote random variables. Row vectors are denoted with an overbar, e.g. $\bar{\mathbf{a}}$. Matrices are denoted with calligraphic letters \mathcal{A}, \mathcal{B} , etc. For binary vectors, we denote by $|\mathbf{a}|$ the number of '1' entries in \mathbf{a} (i.e. the Hamming weight of \mathbf{a}). We denote the set $\{1, \dots, m\}$ by $[m]$.

We denote $\bar{\mathbf{1}} = (1, 1, \dots, 1)$ and $\bar{\mathbf{0}} = (0, 0, \dots, 0)$. For a binary vector \mathbf{a} , let $\text{maj}(\mathbf{a}) \in \{0, 1\}$ be the majority of entries. As a convention, we break ties toward 1 (other tie-breaking methods require slight modifications).

An instance of *Independent judgment aggregation* (IJA) is composed of n judges J (also called voters), m issues I , and a $n \times m$ binary matrix $\mathcal{A} = (a_{ji})_{j=1}^n_{i=1}^m$, where $a_{ji} = 1$ when judge j is in favor of approving issue i , and $a_{ji} = 0$ otherwise. Denote the i 'th column (issue) of \mathcal{A} by \mathbf{a}_i , and the j 'th row (judge) by $\bar{\mathbf{a}}_j$.

The outcome of aggregating the opinions on each issue separately, is denoted by $\bar{\mathbf{s}} = \text{maj}(\mathcal{A}) = (\text{maj}(\mathbf{a}_1), \dots, \text{maj}(\mathbf{a}_m))$. Intuitively, the outcome of a bundling attack (a partition) on IJA \mathcal{A} is as if every judge answers the same on all issues in each bundle, according to the majority of issues in the bundle. Thus the final outcome can be very different from $\text{maj}(\mathcal{A})$.

Formally, let the submatrix $\mathcal{A}|_C$ contain columns $i \in C$ of the matrix \mathcal{A} (i.e. \mathcal{A} restricted to the bundle $C \subseteq I$). The decision on bundle C is simply $d_C(\mathcal{A}) = \text{maj}(\text{maj}(\bar{\mathbf{a}}_1|_C), \dots, \text{maj}(\bar{\mathbf{a}}_n|_C))$. That is, each voter selects 1 or 0 according to the number of issues in C on which he has an affirmative opinion, and then the majority of voters decides on a uniform decision for all issues. We denote by $\bar{\mathbf{s}}(C, \mathcal{A}) \in \{0^{|C|}, 1^{|C|}\}$ the row vector where every entry equals to $d_C(\mathcal{A})$.

Bundling attacks

A *Bundling attack* on IJA instance \mathcal{A} is a partition P of the m issues. Let $P = (C_1, \dots, C_{t^*})$. The outcome of the attack is accepted by concatenating the outcomes of all bundles. Formally,

$$\bar{\mathbf{s}}' = \bar{\mathbf{s}}(P, \mathcal{A}) = (\bar{\mathbf{s}}(C_1, \mathcal{A}), \bar{\mathbf{s}}(C_2, \mathcal{A}), \dots, \bar{\mathbf{s}}(C_{t^*}, \mathcal{A})).$$

We allow to specify partial partitions, where the unspecified columns are assumed to be singletons.

Suppose that the goal of the chair is to approve all issues. Then the success of a bundling attack (i.e. a partition P to bundles) is determined by the number of positive issues in the outcome, i.e. by $|\bar{\mathbf{s}}(P, \mathcal{A})|$. In particular, we say that P is a *perfect bundling* attack if $\bar{\mathbf{s}}(P, \mathcal{A}) = \bar{\mathbf{1}}$, i.e. if the goal of the chair is achieved on all issues.

³Available from <http://tinyurl.com/dx9cuh6>.

Remark 1. *As long as there are no assumptions on the distribution, w.l.o.g. the goal vector is $\bar{\mathbf{1}}$.*

This is since we can replace $\bar{\mathbf{1}}$ with any other target vector $\bar{\mathbf{t}}$, and replace \mathcal{A} with a modified matrix where the names 0 and 1 are flipped in the appropriate columns, formally $\mathcal{A}' = \mathcal{A} + \mathbf{1} \times (\bar{\mathbf{1}} - \bar{\mathbf{t}}) \pmod{2}$. Then $|\{i : \bar{s}(P, \mathcal{A})_i = t_i\}|$ equals to $|\bar{s}(P, \mathcal{A}')|$.

As for the power of bundling attacks, there are simple examples where $\text{maj}(\mathcal{A}) = \bar{\mathbf{0}}$, but $\text{maj}(P, \mathcal{A}) = \bar{\mathbf{1}}$ (for a single bundle $P = \{I\}$). If the fraction of 1's in \mathcal{A} is less than $1/4$, then it is not hard to see that a perfect bundling attack is impossible.

The Ostrogorski paradox.

The following IJA instance appears in (Levmore 1999):

If we aggregate the opinions of the five judges on each issue separately, we get 0 is all issues, i.e. $\bar{s} = \text{maj}(\mathcal{A}) = (0, 0, 0)$.

However, if we bundle all issues together then the three last voters will vote 1, and hence the biased outcome will become $\bar{s}' = \bar{s}(I, \mathcal{A}) = (1, 1, 1)$. This example shows that there are cases where the chair can exploit disagreement among judges to achieve a perfect bundling attack even when it completely reverses the original outcome.

Hardness of Bundling

We define the following variations of the bundling problem. The first is achieving as many positive issues as possible with a single bundle. The second problem is to guarantee that *all* issues will reach a positive decision, allowing as many bundles as needed.

ANYBUNDLING

input AN IJA INSTANCE \mathcal{A} .

question IS THERE A BUNDLE $C \subseteq I$, S.T. $\bar{s}(C, \mathcal{A}) = \bar{\mathbf{1}}$?

k -BUNDLING

input AN IJA INSTANCE \mathcal{A} , A NUMBER $k \leq m$.

question IS THERE A BUNDLE $C \subseteq I$ OF SIZE k , S.T. $\bar{s}(C, \mathcal{A}) = \bar{\mathbf{1}}$?

PERFECTBUNDLING

input AN IJA INSTANCE \mathcal{A} .

question IS THERE A PARTITION P , S.T. $\bar{s}(P, \mathcal{A}) = \bar{\mathbf{1}}$?
All three problems are clearly in NP.

Bundling with a single bundle is hard

We claim that the first two problems are NP-hard. To that end, we want to use a reduction from OPTIMAL-LOBBYING (OL) (Christian et al. 2007). However, we need to slightly modify the original OL problem.

We define the function $\text{maj}_r : \{0, 1\}^{l^*} \rightarrow \{0, 1\}$ as follows. $\text{maj}_r(\mathbf{a})$ returns 1 if $|\mathbf{a}| \geq (l^* + r)/2$ ones, and otherwise returns 0 (i.e. $\text{maj}_0 \equiv \text{maj}$). We extend maj_r to operate on matrices by operating on each column independently, as we did with maj .

OPTIMALLOBBYING' (OL')

input AN IJA INSTANCE \mathcal{A} WITH n^* VOTERS, m^* ISSUES, A NATURAL NUMBER $r \leq n^*/2$.

question IS THERE $\mathcal{A}' \in \{0, 1\}^{n^* \times m^*}$ S.T. $\mathcal{A}, \mathcal{A}'$ DIFFER IN AT MOST r ROWS, AND $|\text{maj}_r(\mathcal{A}')| \geq m^*/2$?

Note that w.l.o.g. modified rows in matrix \mathcal{A}' are $\bar{\mathbf{1}}$. Recently, Nehama (2013) proved that a wide parametrized family of problems that include OL, are all NP-hard. It can be shown that OL', which is similar in spirit to these variations, is also NP-hard.

We are now ready to prove the hardness of our own problem. The reduction from OL' demonstrates the reversed roles of the elements: the voters of the OL' problem become issues in k -BUNDLING and vice versa.

Proposition 1. *k -BUNDLING is NP-complete.*

Proof sketch. Given an OL' instance (\mathcal{A}, r) , we define an instance of k -BUNDLING where the IJA matrix is $\mathcal{B} = \mathcal{A}^T$ (i.e. the transpose of \mathcal{A}), and $k = n^* - r$. Note that in the k -BUNDLING instance we have $n = m^*$ voters and $m = n^*$ issues. Let $K \subseteq [m]$ of size k , and $R = [m] \setminus K$. It can be verified that K is a successful bundling attack on \mathcal{B} iff R is a successful lobbying attack on \mathcal{A} . \square

Note that if k is a constant, then k -BUNDLING can be solved in polynomial time: just try all $\binom{m}{k}$ possible bundles, and check if the attack succeeds.

We emphasize that the hardness of ANYBUNDLING does not follow directly from k -BUNDLING, since in our proof we have to allow values of k above $m/2$. However, ANYBUNDLING has been independently shown to be NP-hard in (Alon et al. 2013).

Perfect bundling by partitioning is hard

Consider the following problem.

ISTRIPARTITE

input A 3-UNIFORM HYPERGRAPH (V, E) .

question ARE THERE V_1, V_2, V_3 S.T. $V_1 \uplus V_2 \uplus V_3 = V$ (I.E. FORM A PARTITION OF V), AND EVERY $e \in E$ INTERSECTS ALL THREE SETS?

We are not aware of any previous complexity analysis of the ISTRIPARTITE problem, however it can be easily shown to be NP-hard by a reduction from graph 3-coloring (Karp 1972).

Proposition 2. *PERFECTBUNDLING is NP-complete.*

Proof sketch. We prove by a reduction from ISTRIPARTITE. Let $t = |V| - 3$. We define a column (an issue) for each vertex $v \in V$. For every hyperedge $e = (v, v', v'')$, we define 5 voters as follows. Two voters (denoted by the set A_e) approve all three issues, and each of the 3 issues is also approved by another single voter (forming the set B_e).

Next, we add t template columns, that are approved by all voters. Finally, we add $|E| - 1$ dummy voters, with rows that are all zero. Thus we have an IJA instance with $n = 6|E| - 1$ and $m = t + |V| = 2|V| - 3$.

“ \Rightarrow ” Assume first that (V, E) is tripartite. Then we can use the following partition to three bundles. We bundle every set V_i , $i \in \{1, 2, 3\}$ with $|V_i| - 1$ template columns. In each such bundle, for every edge e , both voters in A_e and exactly one voter from B_e have more ones than zeros. Thus the bundle is approved.

“ \Leftarrow ” Assume that there is a successful bundling attack $P = (C_1, \dots, C_k)$. Then in every bundle $C_i \in P$ there is a majority of approved rows, i.e. at least $3|E|$ non-dummy voters.

By a careful examination of the partition of template columns, it can be shown that there are exactly 3 bundles. Then, we show that exactly $|E|$ voters of type B approve each bundle, which means that all of the $2|E|$ type A voters must approve all bundles. A voter in $A_e = (v, v', v'')$ will approve all three bundles only if each of his approved columns belongs to a different bundle. Thus every bundle must intersect every triplet e . In other words, the partition P projected on non-template columns is a 3-partition of (V, E) . \square

Frequency of Bundling Attacks

Suppose we are given a random IJA instance. A natural question is “does an efficient bundling attack exist?”

Let $\mathcal{D}_p = \mathcal{D}_p(n, m)$ be the distribution over all $n \times m$ matrices where each entry is (i.i.d.) 1 w.p. p . Intuitively, since we assume the vector $\bar{\mathbf{1}}$ reflects the opinion of the chair, p is the probability that a judge agrees with the chair on a certain issue, where agreement is decided independently for each judge and every issue. $p = 1/2$ is the special case of uniform distribution over all $n \times m$ matrices. $p > 1/2$ means that judges tend to agree with the chair, whereas lower values reflect cases where the chair is trying to achieve a goal contrary to the common opinion.

We denote by $R_p(n, m)$ the probability that there is some bundling attack P guaranteeing at least one positive entry. I.e.

$$R_p(n, m) = Pr_{\mathcal{A} \sim \mathcal{D}_p(n, m)}[\exists P = P(\mathcal{A}) \text{ s.t. } |\bar{\mathbf{s}}(P, \mathcal{A})| \geq 1].$$

Similarly, we denote by $Q_p(n, m)$ the probability that a perfect bundling attack P exists, i.e. where $|\bar{\mathbf{s}}(P, \mathcal{A})| = m$. Clearly $Q_p(n, m) \leq R_p(n, m)$.

By considering the complement of every matrix \mathcal{A} , we observe that for any p , $R_p(n, m) + Q_{1-p}(n, m) \geq 1$ (see appendix in full version).

We think of p as fixed, and take n, m to infinity. Without bundling, the expected fraction of columns with outcome ‘1’ is $1/2$ when $p = 1/2$. For other values, the expected fraction of ‘exceptional’ columns (e.g. positive when $p < 1/2$) decreases exponentially in n .

Our main result shows that there is a phase transition around $p = 1/2$: for values of p lower than $\frac{1}{2}$, no bundling attack can succeed in changing the outcome of even a single issue (when n is not too small), whereas for $p \geq 1/2$ there is almost always an attack that sets all issues to ‘1’. The most interesting case is $p = 1/2$, which is shown last.

High Agreement ($p > 1/2$)

When there is some tendency among judges to agree with the chair, it is easy to achieve a perfect outcome.

Proposition 3. *If $p > 1/2$ and either n or m tend to infinity, then $Q_p(n, m)$ goes to 1.*

Proof. It is easy to see that with high probability (w.h.p.) a single bundle of all issues will work. Indeed, every row has a majority of ones w.p. t_p , where $1 - t_p = Pr_{z \sim Bin(m, p)}[z < m/2] < e^{-2m(p-1/2)^2} \leq e^{-\Omega(m)}$.

Thus w.p. t_p (which goes to 1 with m tending to infinity) a single voter has a majority of ones, and will therefore approve the full bundle.

For any fixed value of m , t_p is some fixed value strictly greater than $1/2$. Thus we can do the same calculation again (replacing p for t_p and m for n) to show that a majority of voters will vote ‘1’ almost always. \square

Low Agreement ($p < 1/2$)

When there are very few voters and many issues (say, $n = o(\ln m)$), then even if $p < 1/2$ there will be some column (w.h.p.) that is approved by a majority of voters. However unless the number of voters is extremely low w.r.t. the number of issues, there will be no such column – and not even a single bundle – that can cause judges to approve any issue.

Proposition 4. *If $p < 0.5$, n tends to infinity and $n \gg \ln m$, then $R_p(n, m)$ goes to 0.*

Proof sketch. We consider separately “small” and “large” bundles (size threshold depends on p). For each size, we bound the probability of every single bundle to succeed using the Chernoff bound twice (once for the columns of a single voter, and once for the rows). Then we use the union bound to bound the total probability of all bundles. \square

Uniform distribution ($p = 1/2$)

Our main theorem for the stochastic setting shows that for $p = 1/2$, we can almost always find a perfect bundling attack. Note that under the uniform distribution, selecting $\bar{\mathbf{1}}$ to be the goal of the chair is w.l.o.g. Thus our result implies that for any goal of the chair, there is almost always a partition that achieves this goal.

Theorem 5. *For $p = 1/2$, $Q_p(n, m)$ goes to 1 when m tends to infinity. Moreover, a successful bundling attack can be found efficiently w.h.p.*

Our algorithm works as follows. First, we take columns whose density (of ‘ones’) is slightly less than $1/2$ (see exact selection criterion in the proof), then arbitrarily partition those columns to bundles of size 3.

Although each such bundle has a low probability to get a positive decision, w.h.p. some of the bundles will be positive. Then, we bundle together all the remaining columns, including those bundles who failed, and the density of this large bundle will be slightly above $1/2$.

Finally, we apply the second moment method to show that the large bundle also succeeds w.h.p., which results in a perfect partition.

We will need the following combinatorial lemma.

Lemma 6. *Let \mathcal{A}^* be a $n \times m$ matrix with integer values bounded by $a^* \in \mathbb{N}$, and let $\mathcal{A} = (a_{ji})_{ji}$ be a random matrix that is obtained by taking a random permutation of every column in \mathcal{A}^* . Denote by $Z_j = \sum_{i=1}^m a_{ji}$ the sum of the j 'th row, and $\mu = \frac{1}{nm} \sum_{ji} a_{ji}^*$. then*

1. $\text{Cov}[Z_j, Z_{j'}] = -\frac{1}{n-1} \text{Var}[Z_j] < 0$, for all $j \neq j'$.
2. In the limit $m \rightarrow \infty$, $\frac{1}{m} Z_j \sim N(\mu, \Theta(1/m))$.

Intuitively, the first part holds since when the sum of a particular row is high, this is an indication that in many columns low entries appear in *other* rows. The second part follows directly from the central limit theorem.

Proof of Theorem 5. Given an input matrix \mathcal{A} , we describe a partition as follows. Let R be the set of columns whose number of ones is $n/2 - \delta$, where $\sqrt{n}/4 \leq \delta \leq \sqrt{n}/2$ i.e., between half and one standard deviations (STDs) from the mean. With probability that goes to 1, R contains a constant fraction α (about 15%) of the columns. Let $1/2 - \mu_R$ be the average density of a column \mathbf{x} in R (i.e. $|\mathbf{x}|/n$), then the expected density of a column in $S = I \setminus R$ is $1/2 + \frac{\alpha}{1-\alpha} \mu_R$.

We divide R to triplets (e.g., by lexicographic order). For every triplet we generate a random permutation of the rows. For every column \mathbf{x} , we apply the respective permutation, and then randomly sample exactly $n/2 - \sqrt{n}/2$ 'ones'. Let \mathbf{x}' be the column obtained from \mathbf{x} after applying the permutation and reduction. For every \mathbf{x}, \mathbf{y} , the intersection of \mathbf{x} and \mathbf{y} is the subset of rows where both have ones. We now say that a triplet $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3$ is *successful*, if the following occur:

- The intersection size of pair $\mathbf{x}'_1, \mathbf{x}'_2$ is between the mean $n/4 - \frac{1}{2}\sqrt{n}$, and $n/4$ (the mean plus one STD).
- The intersection size of pair $\mathbf{x}'_3, \mathbf{y}'$ is strictly over $n/4 + \frac{1}{2}\sqrt{n}$, where $\mathbf{y}' = \mathbf{x}'_1 \oplus \mathbf{x}'_2$ (bit-wise xor).

A short calculation shows that a constant fraction (over 0.5%) of the triplets in R are successful. Note that since we use \mathbf{x}' and not \mathbf{x} to determine success, the following properties apply. First, if \mathbf{x}, \mathbf{y} belong to different triplets then x_j, y_j are independent, even conditioned on success/failure (due to the independent random permutations). Second, the success/failure probability is independent of the actual size of \mathbf{x} (due to the reduction to subsets of the same size). We bundle together every successful triplet, and add all other columns (including failed triplets) to a single large bundle L . We next show that w.h.p. all bundles are approved.

First, in each successful triplet, there are at least $n/4 - \sqrt{n}/2$ rows (voters) that agree with issues $\mathbf{x}'_1, \mathbf{x}'_2$ (and thus also with $\mathbf{x}_1, \mathbf{x}_2$ that contain them), and therefore vote 1. There are also over $n/4 + \sqrt{n}/2$ voters that agree with issue \mathbf{x}_3 and exactly one other issue. In

total, over $n/2$ voters vote 1, and all small bundles are approved w.p. 1.

As for the large bundle, note that it contains all columns, except a small constant fraction (about 1/1000) of columns that are a subset of R . The columns of L either arrive from failed triplets (denote these columns by R') or from singletons S , thus $L = S \cup R'$. Note that the expected density of a column in R' equals to the density in R , which is $1/2 - \mu_R$, and that $|R'| < 0.995|R| = 0.995\alpha m$.

We next show that most rows in L have a majority of ones with probability that tends to 1 as m tends to infinity. Consider first the columns of S . Since all columns are independent, the expected density of each column in S is $1/2 + \frac{\alpha}{1-\alpha} \mu_R$.

The expected density of a column in R' is $1/2 - \mu_R$. Thus, the expected total density of every row j over the large bundle L is

$$\begin{aligned} \mu_L &= \frac{|R'|}{|L|} (1/2 - \mu_R) + \frac{|S|}{|L|} (1/2 + \frac{\alpha}{1-\alpha} \mu_R) \\ &> \frac{0.995\alpha(1/2 - \mu_R)}{0.995\alpha + 1 - \alpha} + \frac{(1-\alpha)(1/2 + \frac{\alpha}{1-\alpha} \mu_R)}{0.995\alpha + 1 - \alpha} \\ &= 1/2 + \frac{0.005\alpha}{1 - 0.005\alpha} \mu_R \geq 1/2 + \Omega(1/\sqrt{n}). \end{aligned}$$

For each $1 \leq j \leq n$ let X_j be an indicator random variable whose value is 1 iff $\text{maj}(\bar{\mathbf{a}}_j|_L) = 1$ (i.e. voter i approves L), and let $X = \sum_{j=1}^n X_j$. We want to show that w.h.p. $X > n/2$.

We replace every three entries in each row that belongs to a failing triplet (i, i', i'') with a single random variable $y_j = a_{j,i} + a_{j,i'} + a_{j,i''}$. Thus $y_j \in \{0, 1, 2, 3\}$, and $E[y] = 3(1/2 - \mu_R)$. Denote the set of merged (non-binary) columns by \tilde{R} . Clearly $|R'| = 3|\tilde{R}|$

Let $r' = |R'|, \tilde{r} = |\tilde{R}|, s = |S|$, and $z = r' + s$. Every row has $\tilde{r} + s$ independent variables, whose values are bounded by 3. Note that since we know only the sum of every column, our $n \times (\tilde{r} + s)$ matrix is equivalent to a fixed matrix with independent random permutation on every column. Denote the sum of the j 'th row by Z_j (note that merging columns does not change the value of Z_j). Then by Lemma 6, Z_j is distributed according to a normal distribution. We know that the expected density of each row is $\mu_L \geq 1/2 + \Omega(1/\sqrt{n})$ by the above calculation, thus $Z'_j = Z_j/z \sim N(1/2 + c/\sqrt{n}, c'/z)$ for some constants c, c' . By the definition of X_j , $\text{Pr}[X_j = 0]$ equals to

$$\text{Pr}_{Z'_j \sim N(\frac{1}{2} + \frac{c}{\sqrt{n}}, \frac{c'}{z})} [Z'_j < \frac{1}{2}] = \text{Pr}_{Z''_j \sim N(0,1)} [Z''_j < -\frac{c''\sqrt{z}}{\sqrt{n}}].$$

Suppose that $z/n = o(1)$, which is the hardest case. This means that $E[X]$ is just slightly above $n/2$. In what follows, we will show that w.h.p. X is sufficiently close to $E[X]$ so that it is still above $n/2$. For row j ,

$$E[X_j] = \text{Pr}[X_j = 1] = 1 - \text{Pr}_{Z''_j \sim N(0,1)} [Z''_j < -c''\sqrt{z/n}].$$

This term equals to $1/2 + \Omega(\sqrt{z/n})$, since $c''\sqrt{z/n}$ is close to 0, and the cumulative distribution function of the normal distribution is roughly linear around 0.

By Lemma 6, $Z_j, Z_{j'}$ are negatively correlated Normal random variables, and therefore $X_j, X_{j'}$ are also negatively correlated. We next bound the variance of X . X_j is binary and thus $\text{Var}[X_j] \leq 1/4$, and $\text{Var}[X] \leq n/4 + \sum_{j \neq j'} \text{Cov}[X_j, X_{j'}] \leq n/4$.

Recall that $E[X] = nE[X_j] = n(1/2 + \Omega(\sqrt{z/n}))$. Thus, by Chebyshev,

$$\Pr[X \leq n/2] \leq \frac{\text{Var}[X]}{(E[X] - n/2)^2} \leq \frac{n/4}{\Omega(n\sqrt{z/n})^2} = O(1/z),$$

which tends to 0 as m (and thus z) grows to infinity.

It remains to prove for values of z above n . However by increasing z/n we only increase $E[X_j]$ and thereby decrease $\Pr[X < n/2]$. Thus our result still holds. \square

Average case complexity

The above results imply that while the PERFECT-BUNDLING problem is NP-hard for some instances, it is easy on average given a distribution $\mathcal{D}_p(n, m)$. Indeed, if we know the parameters then we can answer in $O(1)$ whether a given matrix sampled from this distribution has a perfect bundling attack (i.e. without even looking at the matrix), by answering ‘yes’ iff $p \geq 1/2$. This answer will be true w.h.p. whenever p is fixed, n goes to infinity, and $n \gg \ln m$. In the special case where $n = O(\ln m)$ bounding the error probability may require a more careful analysis of the relations between p, n , and m .

Moreover, given an IJA instance $\mathcal{A} \sim \mathcal{D}_p(n, m)$ where $p \geq 1/2$, we can return a successful bundling attack in polynomial time. If $p > 1/2$ then just return the bundle $[m]$. For $p = 1/2$ we can run the procedure described in the proof of Theorem 5: Columns of R can be easily selected just by counting the number of 1’s. Checking if each triplet is a successful one is also easy. Then, we know that the remaining large bundle succeeds w.h.p. We note that the randomization steps are only required to facilitate the proof, and that our algorithm can be easily derandomized.

Discussion

We studied the problem of bundling in judgment aggregation. We showed that even under the simplifying assumption that opinions over issues are independent, computing whether an optimal bundling attack exists is hard. While the hardness of this combinatorial problem is not very surprising, it reveals an interesting connection with the related problem of optimal lobbying where the roles of voters and issues are reversed.

We also studied the probability that a bundling attack exists in the average case, under various parameters, and demonstrated a sharp threshold phenomenon. In particular, under the uniform distribution, for *any goal vector* \mathbf{t} of the chair, it is almost always possible to construct a perfect bundling attack (we just need to relabel each a_{ij} to ‘1’ if $t_j = 1$ and to ‘0’ otherwise).

Moreover, such a perfect partition can be found in polynomial time, in contrast to the worst case behavior.

The dual problem. Proposition 1 shows that the k -BUNDLING and the OPTIMALLOBBYING problems are “dual” to one another, in the sense that issues and voters play reversed roles.

It turns out that the partitioning problem also has a dual interpretation, where voters and issues change roles. Consider a population of voters in some state that are participating in a referendum about multiple political issues included in a new proposed legislation. In each issue, they can either support or object the policy of the state government, and each district can either adopt the legislation or forgo it entirely, according to whether the majority of issues are supported by its population. This formulation is equivalent to voting over bundles of issues, only we mapped voters, issues and bundles of the original IJA instance, to issues, voters and districts (respectively) of the referendum. More specifically, a single issue supported in a district corresponds to a single voter approving a bundle. A district approving the entire legislation now corresponds to an approved bundle in the original problem. The number of citizens for which the legislation will take effect thus depends not only on their opinions, but also on how they are divided between districts in the state.

The paradoxes above show that it is possible in principle that while *every* voter objects most issues (and thus objects the new legislation), most issues will be supported by the majority of voters in a given partition, and thus the legislation will be approved by all districts. As a result, governors, who can change the segmentation of their states to districts, have substantial power in their hands. This is closely related to the concept of *Gerrymandering*.⁴ We believe that further research can uncover deeper links between studies on Gerreymandering (Lublin 1999; Tasnádi 2011) and the bundling problem studied in this work.

Future work. An important challenge is to better understand the power of bundling in settings where issues are interrelated, for example in the models of Endriss et al. (2010b) or Conitzer et al. (2009). Naturally, computational complexity can only increase when moving to a more general model.

Another venue is to explore bundling attacks when the set of allowed partitions is itself restricted, for example due to constraints on the number of bundles, their size, or relevance of bundled issues.

Acknowledgments

The authors thank Uri Feige and Omer Lev for their useful comments.

⁴From Wikipedia: “In the process of setting electoral districts, gerrymandering is a practice that attempts to *establish a political advantage* for a particular party or group *by manipulating geographic boundaries* to create partisan or incumbent-protected districts.” (our emphasis)

References

- Alon, N.; Bredereck, R.; Chen, J.; Kratsch, S.; Niedermeier, R.; and Woeginger, G. J. 2013. How to (ab)use a multiple referendum for your hidden agenda. manuscript.
- Bartholdi, J.; Tovey, C.; and Trick, M. 1992. How hard is it to control an election? *Mathematical and Computer Modelling* 16(8):27–40.
- Baumeister, D.; Erdélyi, G.; Erdélyi, O.; and Rothe, J. 2012. Control in judgment aggregation. In *STAIRS'12*. 23–34.
- Christian, R.; Fellows, M.; Rosamond, F.; and Slinko, A. 2007. On complexity of lobbying in multiple referenda. *Review of Economic Design* 11:217–224.
- Conitzer, V.; Lang, J.; and Xia, L. 2009. How hard is it to control sequential elections via the agenda. In *IJCAI'09*, 103–108.
- Endriss, U.; Grandi, U.; and Porello, D. 2010a. Complexity of judgment aggregation: Safety of the agenda. In *AAMAS'10*, 359–366.
- Endriss, U.; Grandi, U.; and Porello, D. 2010b. Complexity of winner determination and strategic manipulation in judgment aggregation. In Conitzer, V., and Rothe, J., eds., *COMSOC'10*.
- Farquharson, R. 1969. *Theory of Voting*. Yale Uni. Press.
- Hemaspaandra, E.; Hemaspaandra, L.; and Rothe, J. 2007. Hybrid elections broaden complexity-theoretic resistance to control. In *IJCAI'07*, 1308–1314.
- Karp, R. M. 1972. *Reducibility among combinatorial problems*. Springer.
- Lacy, D., and Niou, E. 2000. A problem with referendums. *Journal of Theoretical Politics* 12(1):5–31.
- Levmore, S. 1999. Voting paradoxes and interest groups. *The Journal of Legal Studies* 28(2):259–281.
- Lublin, D. 1999. *The paradox of representation: Racial gerrymandering and minority interests in Congress*. Princeton University Press.
- Nehama, I. 2013. Complexity of optimal lobbying in threshold aggregation (extended abstract). In *AAMAS'13*. to appear.
- Palfrey, T. 1983. Bundling decisions by a multiproduct monopolist with incomplete information. *Econometrica: Journal of the Econometric Society* 463–483.
- Tasnádi, A. 2011. The political districting problem: A survey. *Society and Economy* 33(3):543–554.
- Xia, L.; Lang, J.; and Ying, M. 2007. Sequential voting rules and multiple elections paradoxes. In *TARK'07*, 279–288. ACM.