The Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21)

Restricted Domains of Dichotomous Preferences with Possibly Incomplete Information

Zoi Terzopoulou,[1] Alexander Karpov,[2] and Svetlana Obraztsova[3]
[1]Institute for Logic, Language and Computation, University of Amsterdam, The Netherlands
[2]HSE University, Institute of Control Sciences of Russian Academy of Sciences, Russia
[3]Nanyang Technological University, Singapore

Abstract

Restricted domains over voter preferences have been extensively studied within the area of computational social choice, initially for preferences that are total orders over the set of alternatives and subsequently for preferences that are dichotomous—i.e., that correspond to approved and disapproved alternatives. This paper contributes to the latter stream of work in a twofold manner. First, we obtain forbidden subprofile characterisations for various important dichotomous domains. Then, we are concerned with incomplete profiles that may arise in many real-world scenarios, where we have partial information about the voters' preferences. We tackle the problem of determining whether an incomplete profile admits a completion within a certain restricted domain and design constructive polynomial algorithms to that effect.

1 Introduction

Individual preferences on the one hand, and the aggregation of these preferences into one collective choice on the other hand, constitute central elements of AI research (Domshlak et al. 2011). With applications ranging from recommender systems to electronic voting and automated personal assistants, the problem of choosing suitable preference models and aggregation methods becomes evident. But well-behaved aggregation mechanisms are not always easy to find, notably because determining the outcome of the aggregation is often an intractable task (e.g., Procaccia, Rosenschein, and Zohar, 2007). Luckily, good news come to light under the assumption that the agents' preferences conform to a certain structure, also known as a *domain restriction*.

For the above reason, domain restrictions over the preferences of voters have received increasing attention within the field of *computational social choice* (see Elkind, Lackner, and Peters (2017) for a recent survey). On a conceptual level, restricted domains represent structures that arise as sensible preference models in many real-life settings; on a technical level, they allow for the efficient application of several voting mechanisms, the utilisation of which is in general a computationally hard problem.

More specifically, restricted domains of preferences that are total orders over the set of alternatives are well-studied. In contrast, domains of dichotomous preferences where the voters hold an approved and a disapproved set of alternatives—although very natural—were developed only recently. Elkind and Lackner (2015) introduced various domains of this kind (later generalized by Yang, 2019); they showed that a number of structures of dichotomous preferences admit polynomial algorithms in the context of two popular approval-based multiwinner rules, for which determining the winning committee is known to be NP-hard: Proportional Approval Voting (PAV) and Maximin Approval Voting (MAV), described by Kilgour and Marshall (2012) and Brams, Kilgour, and Sanver (2007), respectively.

This paper consists of two parts. The first part builds on Elkind and Lackner's work with the following contribution:

- We prove characterisation theorems for restricted dichotomous domains on which the results of Elkind and Lackner (2015) rely, by identifying the patterns that prevent a preference profile from exhibiting a certain structure.

In particular, introducing an order of the alternatives and introducing an order of the voters are the two main approaches when defining structured dichotomous preferences. But a structure of the voters is dual to a structure of the alternatives. So, we present all our results for domains resting on voter structures, but these results can be directly translated and assumed to hold for structures of alternatives as well.

The literature on domains of total orders contains characterization results using forbidden patterns, reminiscent to ours. Ballester and Haeringer (2011) characterised the single-peaked and the group-separable domains, while Bredereck, Chen, and Woeginger (2013) performed the task for the single-crossing domain, and Peters and Lackner (2020) worked on the domain of single-peaked preferences on a circle. We now know that all these domains are characterised by a finite number of forbidden patterns—but this is not true for other domains of total orders. For instance, Chen, Pruhs, and Woeginger (2017) proved that finitely many forbidden patterns are not enough for the one-dimensional euclidean domain. It is also worth stressing here that, for any domain characterised by a finite number of forbidden patterns, it is computationally easy to check whether a preference profile conforms to it. Bartholdi III and Trick (1986) designed the original algorithm for detecting whether a profile is single-peaked, while Elkind, Faliszewski, and Slinko (2012) and Bredereck, Chen, and Woeginger (2013) solved the same exercise for single-crossing preferences.

However, one important aspect has not been considered in the literature to date, namely the fact that our information about the exact dichotomous preferences of the voters will often be *incomplete*. Either because it is costly for the voters to report all their preferences, or because they have not yet formed full preferences when they are asked to express them, we may have no access to the complete preference profile. Yet, it is crucial to know whether a certain structure can potentially be manifested in a given incomplete profile, most importantly to understand whether an aggregation method has the chance to be efficiently applied.

The second part of this paper, which contains our main focus, aims to close the aforementioned gap in the literature. We point at the general case of incomplete information on dichotomous preferences and ask:

- Given an incomplete profile, is it *possible* to complete it in a way that complies with a certain restriction? Is it *necessary* that a completion will conform to the restriction of our interest? If the answer is positive, can we efficiently discover an appropriate completion? We design polynomial algorithms that provide constructive answers for the relevant dichotomous domains.

Our original algorithms subsume and extend the algorithms for the complete case of Elkind and Lackner (2015).

Regarding incomplete profiles of total orders, the work of Lackner (2014) (recently elaborated upon by Fitzsimmons and Lackner, 2020) was the first to address the problem of extending partial preferences to full preferences that respect a given restriction, specifically investigating the domain of single-peaked preferences on total orders. Following up, Elkind et al. (2015) explored single-crossing domains.

Along similar lines, researchers in the area of *preference elicitation* (Walsh 2008; Conitzer 2009) are specifically concerned with scenarios where the voters are not able to report their full preferences and we thus need to perform a limited number of queries, for example asking for comparisons between two alternatives at a time. Knowing whether an incomplete profile admits a completion that complies with a desirable structure is an essential part of preference elicitation, which currently only involves preferences that—when fully elicited—are total orders.

Lastly, a different—yet intuitively related—task on domain restrictions, pioneered by Faliszewski, Hemaspaandra, and Hemaspaandra (2014) and so far only explored on domains of total orders, is about recognising profiles that *nearly* enjoy a given structure (Elkind and Lackner 2014; Erdélyi, Lackner, and Pfandler 2017; Jaeckle, Peters, and Elkind 2018), according to some distance metric. This is very natural area to investigate for dichotomous preferences as well, but this is a topic for another paper.[1]

The remainder of this paper is organised as follows. Section 2 introduces our model and reviews a number of domain restrictions on dichotomous preferences of practical relevance. In Section 3, we characterise complete profiles in these domains via forbidden patterns. In Section 4, we present the principal contributions of this paper: We study settings of incomplete information and design polynomial algorithms that detect whether a given incomplete profile can possibly (and analogously, necessarily) have a completion with a certain structure. Then, Section 5 concludes.

2 The Model

This section presents our basic notation and terminology and defines domain restrictions on dichotomous preferences.

Preliminaries

In our model, a finite set of *voters* $\mathcal{N} = \{v_1, \ldots, v_n\}$, with $n \geq 2$, hold *dichotomous preferences* over a finite set of *alternatives* $\mathcal{A} = \{a_1, \ldots, a_m\}$, with $m \geq 2$. That is, a voter v_j either approves or disapproves an alternative a_i, denoted by $p_{i,j} = 1$ and $p_{i,j} = 0$, respectively. The dichotomous preferences of all voters are captured by a *profile* P, which is an $m \times n$ binary matrix.[2] We may also have incomplete information about that matrix, corresponding to cells of unknown value "?". Let $\mathcal{M}_{m\times n}$ be the set of all *complete* $m \times n$ matrices with entries "0" or "1", and let $\mathcal{I}_{m\times n}$ be the set of all *incomplete* matrices with entries "0", "1", or "?". So, $\mathcal{M}_{m\times n} \subseteq \mathcal{I}_{m\times n}$. Given matrices $X \in \mathcal{I}_{m\times n}$ and $Y \in \mathcal{M}_{m\times n}$, we say that Y is a *completion* of X if every cell of known value in X has the same value in Y.

For a number $k \in \mathbb{N}$, we denote by $[k]$ the set $\{1, \ldots, k\}$ and by S_k the set of all permutations on $[k]$. For two matrices $X, Y \in \mathcal{I}_{m\times n}$, we say that X and Y are *equivalent* if X equals Y after some permutation of rows and columns:

$$X \equiv Y \text{ if } x_{i,j} = y_{\sigma(i),\tau(j)} \text{ for some } \sigma \in S_m, \tau \in S_n.$$

We say that the matrix $X \in \mathcal{I}_{k\times\ell}$ *occurs as a pattern* in the matrix $Y \in \mathcal{I}_{m\times n}$ if for some submatrix $Z \in \mathcal{I}_{k\times\ell}$ of Y it is the case that $X \equiv Z$. If X does not occur as a pattern in Y, we say that Y *avoids* X. For any class of matrices $\mathcal{X}$, we write $Av(\mathcal{X}) = \{Y \mid \text{ for all } X \in \mathcal{X}, Y \text{ avoids } X\}$ for the set of matrices that avoid all matrices in $\mathcal{X}$. For simplicity, we say that a matrix $Y \in Av(\mathcal{X})$ *avoids* the class $\mathcal{X}$.

For $x \in \{0, 1\}$, we denote by $P[p_{i,j}|x]$ the new matrix obtained from the matrix P by placing the value "x" in the cell $p_{i,j}$ (that may previously have a known or an unknown value). Then, the operation that we call *cancellation* applies in a row of a matrix and changes all its elements from "0" to "1" and from "1" to "0". For a subset of alternatives $A \subseteq \mathcal{A}$, we denote by $P[\overline{A}]$ the profile obtained from P by cancelling all rows corresponding to the alternatives in A.

Finally, given a profile $P \in \mathcal{I}_{m\times n}$, we define P's *consecutive order graph*, a novel object that will be a very useful tool for our technical results. We construct P's consecutive order graph as follows: We have n nodes, one for each voter, and we have a directed edge from node v_j to node v_ℓ if and only if for some alternative a_i it is the case that v_j approves a_i and v_ℓ disapproves a_i (i.e., $p_{i,j} = 1$ and $p_{i,\ell} = 0$). See Figure 1 for an example and note that for the remainder, we will not indicate the voters and the alternatives on the various profiles—it should be clear that voters correspond to columns and alternatives to rows of the matrix.

[1] Completing incomplete matrices in order to satisfy certain desirable properties is studied in many different contexts as well. For example, Ganian et al. (2018) examined completions that minimize the rank, or the number of distinct rows of a matrix.

[2] We will use the terms "profile" and "matrix" interchangeably.

	v_1	v_2	v_3	v_4
a_1	1	?	0	?
a_2	?	?	0	1
a_3	1	1	1	1
a_4	0	0	?	?
a_5	?	0	1	0

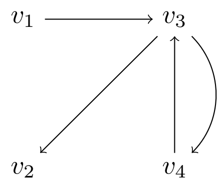

Figure 1: An incomplete profile and its representation via the consecutive order graph.

Domain Restrictions

We are interested in *domain restrictions*, i.e., in dichotomous preferences with a special structure, which are significant with respect to the application of prominent approval-based rules. To this end, we study the following three domain restrictions that we formally define below: VI, VEI, and PART. These particular restrictions on complete preference profiles have been found to be of major interest to the computational social choice community, since Elkind and Lackner (2015) showed that they admit efficient algorithms for the popular—yet computationally hard—rules PAV and MAV.[3]

In addition to the above, we consider a novel restriction to which we refer as SVEI. SVEI is logically stronger than VEI, and hence it also allows for efficiency regarding both PAV and MAV. The reason why we include SVEI in our analysis is twofold: First and foremost, analysing the properties of SVEI and obtaining reliable results for it is very useful technically, as our relevant results for the more standard property of VEI heavily depend on the former. Second, SVEI is not unreasonable in voting scenarios. For example, assume that the alternatives constitute the candidates of a specific political party. We may then order the voters from the most loyal party supporter to the most adversarial one so that each candidate firstly wins the support of the most loyal voter, then possibly also gets a vote from the next most loyal voter (depending on how convincing she is), and so on. Note also that in a profile that satisfies SVEI, most common approval-based multiwinner voting rules (like Chamberlin-Courant, PAV, and MAV) will select the same, optimal winning committee; that is, the committee with the most votes (Faliszewski et al. 2017).

Definition 1. *We say that a profile $P \in \mathcal{M}_{m\times n}$ satisfies*

- ***Voter Interval (VI)*** *if the voters can be reordered so that for every alternative a_i, the voters that approve it form an interval of the ordering.*
- ***Voter External Interval (VEI)*** *if the voters can be reordered so that for every alternative a_i, the voters that approve it form a prefix or a suffix of the ordering.*
- ***Single-sided Voter External Interval (SVEI)*** *if the voters can be reordered so that for every alternative a_i, the voters that approve it form a prefix of the ordering.*

[3]We also know by Elkind and Lackner (2015) that the property *Candidate Interval* (CI) is equivalent to many other properties, viz. *Dichotomous Euclidean*, *Possibly Single-Peaked*, and *Possibly Euclidean*. Hence, by studying VI (which, as we mentioned in the introduction, is dual to CI), we have answers for all these domains.

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

(a) VI

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

(b) VEI

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

(c) SVEI

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

(d) PART

Figure 2: Ordered profiles satisfying different properties.

- ***Partition (PART)*** *if the set of alternatives $\mathcal{A}$ can be partitioned into subsets $A_1, \ldots, A_\ell$ such that each voter approves exactly one of the sets $A_1, \ldots, A_\ell$ and each such set is approved by at least one voter.*

Clearly, each of SVEI, PART, and VEI implies VI. Figure 2 provides examples for the above properties.

3 Complete Profiles

In this section, we characterise profiles in restricted domains of dichotomous preferences via *forbidden patterns*: For each property, we identify a class of matrices $\mathcal{Y}$ such that a profile will satisfy the property if and only if it avoids $\mathcal{Y}$.

Although no characterisations of dichotomous preference domains exist in the computational social choice literature, Tucker (1972), working in combinatorics, has obtained a forbidden pattern characterisation for VI that contains infinitely many forbidden patterns. To be specific, Tucker obtained forbidden patterns that characterise matrices with the *consecutive-ones* property in columns, which are exactly transposed matrices with the VI property.

Proposition 1 (Tucker, 1972)**.** *A profile $P \in \mathcal{M}_{m\times n}$ satisfies the VI property if and only if it avoids the (infinite) class $\mathcal{T} = \bigcup_{k=1}^{3} T_k \cup \{T_4, T_5\}$, where $T_k = \bigcup_{\ell=1}^{\infty}\{T_k^\ell\}$ and for all $\ell \geq 1$, $T_1^\ell, T_2^\ell,$ and T_3^ℓ are such that:*

$$T_1^\ell = \begin{pmatrix}1&1&&&&\\&1&1&&&\\&&\cdot&\cdot&0&\\&0&&\cdot&\cdot&\\&&&&1&1\\&&&&&1&1\\1&0&&\ldots&&0&1\end{pmatrix} \in \mathcal{M}_{(\ell+2)\times(\ell+2)},$$

$$T_2^\ell = \begin{pmatrix}1&1&&&&&0\\&1&1&&&&0\\&&\cdot&\cdot&0&&0\\&0&&\cdot&\cdot&&\vdots\\&&&&1&1&0\\&&&&&1&1&0\\1&1&1&\ldots&1&1&0&1\\0&1&1&\ldots&1&1&0&1\end{pmatrix} \in \mathcal{M}_{(\ell+3)\times(\ell+3)},$$

$$
T_3^\ell = \begin{pmatrix} 1 & 1 & & & & & & 0 \\ & 1 & 1 & & & & & 0 \\ & & \ddots & & & 0 & & \vdots \\ & 0 & & \ddots & & & & \\ & & & & 1 & 1 & & 0 \\ & & & & & 1 & 1 & 0 \\ 0 & 1 & 1 & \dots & 1 & 1 & 0 & 1 \end{pmatrix} \in \mathcal{M}_{(\ell+2)\times(\ell+3)},
$$

$$
T_4 = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 & 0 & 1 \end{pmatrix}, T_5 = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 & 1 \end{pmatrix}.
$$

Next, we present original characterisations for the properties of VEI, SVEI, and PART.

Proposition 2. *A profile $P \in \mathcal{M}_{m\times n}$ satisfies the VEI property if and only if it avoids the class*

$$
\mathcal{Z} = \bigcup_{k=1}^{5} \{Z_k\}, \text{ where}
$$

$$
Z_1 = \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix}, Z_2 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, Z_3 = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix},
$$

$$
Z_4 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 1 & 1 \end{pmatrix}, Z_5 = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 \end{pmatrix}.
$$

Proof. Trivially, matrices containing patterns from the class $\mathcal{Z}$ do not satisfy VEI. Then, we need to prove that $\mathcal{Z}$ contains all forbidden patterns for VEI. Recall that VEI is a logically stronger property than VI. This means that in every matrix $T \in \mathcal{T}$ (where $\mathcal{T}$ is the class of forbidden patterns for VI as defined in Proposition 1), there must occur a subpattern Z that is forbidden for VEI.

Now, suppose that a profile P violates VEI. There are two cases. *Case 1*: P violates VI. Then, P must contain a forbidden pattern Z for VEI that is a subpattern of a forbidden pattern T for VI. *Case 2*: P satisfies VI. Then, P contains a forbidden pattern Z for VEI since it violates this property (and clearly forbidden patterns for VEI exist—we just do not yet know which they are). Now, if we cancel some rows in Z, then P will violate VI (because if that weren't the case, then P would satisfy VEI, contradicting our hypothesis).

So, a class of matrices $\mathcal{Y}$ includes all forbidden patterns for VEI if (i) every matrix in $\mathcal{T}$ contains a pattern from $\mathcal{Y}$, and (ii) when closed under cancellation for any subset of rows in any of its matrices, $\mathcal{Y}$ gives rise to the same patterns.

For the class $\mathcal{Z}$ of our statement, it can be easily checked that condition (ii) holds. We show that condition (i) holds as well: For $\ell = 1$, Tucker's matrix T_1^ℓ coincides with Z_1. For $\ell \geq 2$, the matrix T_1^ℓ contains the pattern Z_5 in the last two rows. For $\ell \geq 1$, the matrices T_2^ℓ and T_3^ℓ contain the pattern Z_5 in the first two rows. Lastly, Tucker's matrices T_4 and T_5 contain the pattern Z_5 in the last two rows. □

Proposition 3. *A profile $P \in \mathcal{M}_{m\times n}$ satisfies the SVEI property if and only if it avoids the pattern X, where*

$$
X = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.
$$

Proof. Note that a profile satisfying SVEI can be uniquely reconstructed from its row and column sums. Matrices with this property are characterized by the above forbidden pattern X (Ryser 1957). □

Proposition 4. *A profile $P \in \mathcal{M}_{m\times n}$ satisfies the PART property if and only if (i) it avoids the pattern W, where*

$$
W = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}, \text{ and}
$$

(ii) *every row and every column of P has at least one "1".*

Proof. Obviously, if W occurs as a pattern in P, then P cannot satisfy PART; if condition (ii) does not hold, then P cannot satisfy PART either, by definition. Suppose now that condition (ii) holds and W does not occur as a pattern in P. Then, for every voter v_j we can define the set of alternatives that she approves as $A_j = \{a_i \in \mathcal{A} \mid p_{i,j} = 1\}$, and for any two voters v_j, v_ℓ it will be the case that either $A_j = A_\ell$ or $A_j \cap A_\ell = \emptyset$. So, we have a partition of $\mathcal{A}$ as prescribed by the definition of PART. □

Notably, the characterisation results of this section play an essential role in settings of incomplete information too, when we are looking for profiles that *possibly* or *necessary* satisfy a given property. They imply that we can check whether the aforementioned condition holds in polynomial time. Most importantly, in Section 4 we will additionally see how we can address this issue in a constructive fashion.

4 Incomplete Profiles

Before proceeding to the second part of this paper, some additional terminology is in order. Given an incomplete profile $P \in \mathcal{I}_{m\times n}$, we say that P *possibly* (respectively, *necessarily*) satisfies a specific property if *some* (respectively, *all*) completions of P satisfy that property.

In this section, we address the following questions: Given an incomplete profile $P \in \mathcal{I}_{m\times n}$, can we detect easily whether P admits a completion that conforms to a specific structure? And if such a completion exists, can we find it?

We know that the problem of detecting whether an incomplete $m \times n$ profile can be completed in a way such that VI is satisfied is NP-complete (Klinz, Rudolf, and Woeginger 1995), while Golumbic (1998) showed that the analogous problem for PART can be solved in polynomial, $O(mn)$, time.[4] We will next design polynomial algorithms for both SVEI and VEI (with the latter building on the former).

[4]Since Proposition 4 provides a characterisation of PART via finitely many forbitten subprofiles, it is not hard to see that checking whether an incomplete profile *necessarily* has a completion that satisfies this property can also be done in polynomial time.

SVEI

The algorithm SVEI-INCOMPLETE is based on three sub-algorithms, viz., FILLING, GRAPH, and ORDERING. In what follows, $P \in \mathcal{I}_{m\times n}$ is a profile and L is a linear order over $\mathcal{N}$.

FILLING(P): First set P_f to P. For an arbitrary cell of unknown value $p_{i,j}$ in P_f, check whether the matrices $P_f[p_{i,j}|0]$ and $P_f[p_{i,j}|1]$ contain the forbidden pattern X for SVEI. If they both do, announce "invalid" and exit; if neither does, continue to the next cell of unknown value in P_f; if only one—the matrix $P_f[p_{i,j}|x]$ for $x \in \{0,1\}$—does, then set P_f to $P_f[p_{i,j}|1-x]$ and continue to the next cell of unknown value in P_f. When no other cell of unknown value remains to be considered, return the profile P_f. Repeat this process nm times—each repetition ensures that at least one cell that can be filled will indeed be filled, and the thus order of the selected cell does not matter overall.

GRAPH(P): Construct the consecutive order graph of P.

ORDERING(P, L): Order the voters in P according to L by defining the permutation $\sigma \in S_n$ such that $\sigma(i) < \sigma(j)$ whenever $(v_i, v_j) \in L$. Return the ordered profile P_o.

Lemma 1 ensures that the algorithm SVEI-INCOMPLETE, which we will shortly construct, will be well-defined.

Lemma 1. *If the forbidden configuration X for SVEI does not occur as a pattern in the profile P_f (after FILLING(P) has been applied), then the consecutive order graph of P_f is acyclic and can thus be extended to a linear order.*

Proof. We prove the contrapositive. Suppose—without loss of generality on the names of the voters—that there is a cycle $v_1 \to v_2 \to \ldots \to v_k \to v_1$ in the consecutive order graph of P. Then, some alternative x_1 is approved by v_1 and rejected by v_2, some alternative x_2 is approved by v_2 and rejected by v_3,..., and some alternative x_k is approved by v_k and rejected by v_1. Without loss of generality, assume that $x_1 = a_1$, $x_2 = a_2$, ..., and $x_k = a_k$. Then:

$$\begin{array}{llll} p_{1,1} = 1 & (1a) & p_{1,2} = 0 & (1b) \\ p_{2,2} = 1 & (2a) & p_{2,3} = 0 & (2b) \\ \vdots & & \vdots & \\ p_{k,k} = 1 & (ka) & p_{k,1} = 0 & (kb) \end{array}$$

If $p_{1,3} = 1$, we obtain a forbidden configuration for SVEI from Equations ($1b$), ($2a$), and ($2b$). Otherwise, it must be the case that

$$p_{1,3} = 0. \qquad (1')$$

Continuing with the same reasoning for $p_{1,4}$, $p_{1,5}$, etc., we either reach a step where we locate a forbidden configuration for SVEI, or deduce that

$$p_{1,k} = 0. \qquad (k')$$

But from Equations ($1a$), (ka), (kb), and (k'), we have that

$$\begin{array}{ll} p_{1,1} = 1 & p_{1,k} = 0 \\ p_{k,1} = 0 & p_{k,k} = 1, \end{array}$$

which is a forbidden configuration for SVEI. ☐

$$\begin{pmatrix} 0 & ? & ? & 1 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & ? \\ ? & ? & ? & 0 \end{pmatrix} \qquad \begin{pmatrix} 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 \\ 1 & 0 & 1 & ? \\ 0 & 0 & ? & 0 \end{pmatrix}$$

(a) original profile (b) filling

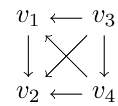

$v_3 \to v_4 \to v_1 \to v_2$

(c) graph (d) linearisation

$$\begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & ? & 1 & 0 \\ ? & 0 & 0 & 0 \end{pmatrix} \qquad \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & ? & 1 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}$$

(e) ordering (f) completing

Figure 3: SVEI-INCOMPLETE: An example.

We are now ready to define our main algorithm.

SVEI-INCOMPLETE(P): Apply FILLING(P). If "invalid" is announced, exit with failure. Otherwise, apply GRAPH(P_f) and extend the resulting graph to a linear order L. Then, obtain an ordered profile P_o by calling ORDERING(P_f, L). Exit and return P_o.

Next, we design an algorithm that completes an incomplete profile so that SVEI will hold, provided that this is possible.

S-COMPLETING(P): Repeat the steps below for all rows $i \in \{1,\ldots,m\}$. Take the smallest $j \in \{1,\ldots,n\}$ such that $p_{i,j} \notin \{0,1\}$. If there is no $\ell < j$ with $p_{i,\ell} = 0$, then set P to $P[p_{i,j}|1]$. Otherwise, set P to $P[p_{i,j}|0]$. Repeat for the next cell of unknown value in this row.

Proposition 5. *SVEI-INCOMPLETE detects in polynomial time whether a profile of dichotomous preferences possibly satisfies SVEI. If it does, SVEI-INCOMPLETE also finds an appropriate order of the voters and S-COMPLETING finds a suitable completion in polynomial time.*

Proof. Indeed, SVEI-INCOMPLETE terminates in polynomial time: FILLING takes $O(m^3n^3)$ time. We can also construct the consecutive order graph of P in $O(mn)$ time and extend it to a linear order L in $O(n+n^2)$ time (Kahn 1962). Finally, ORDERING and S-COMPLETING can be done in $O(n^2)$ and $O(mn)$ time, respectively.

Suppose now that FILLING did not announce "invalid", which means that a forbidden pattern for SVEI was not detected in P_f. Then, intuitively, GRAPH(P_f) draws an edge from v_j to v_ℓ when v_j has to appear before v_ℓ in the final ordering, and the linear order L that extends the consecutive order graph preserves this property. So in the ordered profile, for every alternative a_j, the voters that approve it appear in the ordering before those that disapprove it, and this also holds after S-COMPLETING. Hence, SVEI-INCOMPLETE will work correctly and the final completion will satisfy SVEI. ☐

$$\begin{pmatrix}1 & 0 & 0 & ?\\ 1 & ? & 0 & 1\end{pmatrix} \qquad \begin{pmatrix}1 & 0 & 0 & ?\\ 0 & 0 & 1 & 1\end{pmatrix}$$

Figure 4: Rows corresponding to alternatives that match (on the left) and that contradict (on the right) each other.

Let us illustrate SVEI-INCOMPLETE with an example.

Example 1. Consider a profile P of dichotomous preferences as depicted in (a) of Figure 3. Then, FILLING(P) in step (b) can be easily done: For instance, we see that the cell with unknown value in the first row and third column of the matrix has to be assigned with "1" in order to prevent the creation of a forbidden subprofile with the corresponding cells of the fourth column and second row. Then, GRAPH(P) in step (c) is constructed according to the relevant definition. For step (d), note that v_3 has no incoming edges in the consecutive order graph, so it should be placed first in the linear order, and so on. The ordering and completing algorithms in steps (e) and (f) follow quite straightforwardly once we know the right sequence of the voters. △

Next, recall that we know exactly what the conditions that can prevent a complete profile from satisfying SVEI are, from Proposition 3. Consequently, we can easily detect whether a profile *necessarily* satisfies SVEI. The proof of Lemma 2 is immediate and, as such, omitted.

Lemma 2. *An incomplete profile of dichotomous preferences P necessarily satisfies SVEI if and only if none of the following matrices occurs as a pattern in P:*

$$\begin{pmatrix}0 & 1\\ 1 & 0\end{pmatrix}, \begin{pmatrix}? & ?\\ 1 & 0\end{pmatrix}, \begin{pmatrix}0 & ?\\ 1 & ?\end{pmatrix}, \begin{pmatrix}? & 1\\ 1 & ?\end{pmatrix}, \begin{pmatrix}0 & ?\\ ? & 0\end{pmatrix},$$

$$\begin{pmatrix}0 & 1\\ 1 & ?\end{pmatrix}, \begin{pmatrix}0 & ?\\ 1 & 0\end{pmatrix}, \begin{pmatrix}? & ?\\ ? & 0\end{pmatrix}, \begin{pmatrix}? & ?\\ ? & 1\end{pmatrix}, \begin{pmatrix}? & ?\\ ? & ?\end{pmatrix}.$$

Proposition 6. *We can check whether an incomplete profile of dichotomous preferences necessarily satisfies SVEI in polynomial time.*

Proof. The statement is deducible from Lemma 2, since we can check whether a relevant subprofile occurs as a pattern in a given incomplete profile in polynomial time. □

VEI

Relying on our algorithm for SVEI, we proceed to design an algorithm for VEI. The key intuition is the following: A profile P satisfies VEI if and only if the profile obtained from P by cancelling certain rows satisfies SVEI. But which rows should be cancelled, if any? Loosely speaking, our algorithm VEI-INCOMPLETE answers that question.

Given a profile P, we say that an alternative a_i *contradicts* the alternative a_k if there exist $j, \ell \in [n]$ such that $p_{i,j} = 1 = 1 - p_{i,\ell}$ and $p_{k,j} = 0 = 1 - p_{k,\ell}$. Analogously, an alternative a_i *matches* the alternative a_k if there exist $j, \ell \in [n]$ such that $p_{i,j} = p_{k,j} = 1$, and $p_{i,\ell} = p_{k,\ell} = 0$. See Figure 4 for an example, and note that two alternatives may simultaneously match and contradict each other.

Below we present the algorithm EXPANDING, used in the algorithm COARSENING that we call in VEI-INCOMPLETE. The main function of the algorithm EXPANDING is to expand every set of alternatives A (within a partition $\mathcal{P}$ of $\mathcal{A}$) into a new set $A' \supseteq A$ by adding to it alternatives that either match or contradict the elements of A—in case of a contradiction, the rows corresponding to alternatives that are added to A are cancelled. The algorithm COARSENING repeats this process until we know exactly which rows must be cancelled.

In what follows, $P \in \mathcal{I}_{m\times n}$ is a profile, $\mathcal{P}$ is a partition of the set of alternatives, and $A \in \mathcal{P}$ is a set of alternatives.

EXPANDING($A, \mathcal{P}, P$): Start with setting $P' = P$ and $\mathcal{P}' = \mathcal{P}$. Then, repeat the following steps for all sets $Y \in \mathcal{P}$ (including A). Consider first some $B \in \mathcal{P}$.
Check if some $a_i \in A$ matches some $a_j \in B$.

(1) If it does, mark B as used; continue to the next $C \in \mathcal{P}$.

(2) Otherwise, check if some $a_i \in A$ contradicts an $a_j \in B$.

  (2a) If it does, set P' to $P[\overline{B}]$, mark B as used, and continue to the next $C \in \mathcal{P}$.

  (2b) Otherwise, continue to the next $C \in \mathcal{P}$.

When there is no other set $C \in \mathcal{P}$ to consider, return the sets that are marked as used, the profile P', and the partition

$$\mathcal{P}' = \{ \bigcup_{\substack{B\in\mathcal{P}\\ B \text{ used}}} B\} \cup \{C \in \mathcal{P} \mid C \text{ not used }\}.$$

COARSENING($\mathcal{P}, P$): First, mark all sets in $\mathcal{P}$ as not used. Take a set of alternatives $A_1 \in \mathcal{P}$ and apply EXPANDING($A_1, \mathcal{P}, P$), obtaining a new profile P_1 and a new partition $\mathcal{P}_1$. Then, take a set $A_2 \in \mathcal{P}$ that is still not marked as used and apply EXPANDING($A_2, \mathcal{P}_1, P_1$), obtaining a new profile P_2 and a new partition $\mathcal{P}_2$. Repeat until you find the last set $A_k \in \mathcal{A}$ that is not marked as used. Apply EXPANDING($A_k, \mathcal{P}_{k-1}, P_{k-1}$), obtaining a new profile P_k and a new partition $\mathcal{P}_k$. Return P_k and $\mathcal{P}_k$.

VEI-INCOMPLETE(P): Apply COARSENING($\mathcal{P}, P$) first for the finest partition $\mathcal{P} = \{\{a_1\}, \ldots, \{a_m\}\}$. If there are alternatives in the same set of the output partition $\mathcal{P}'$ that contradict each other, then exit with failure. Otherwise, apply COARSENING($\mathcal{P}', P'$). If in the new output partition $\mathcal{P}''$ there are alternatives in the same set that contradict each other, then exit with failure. Otherwise, continue by applying COARSENING($\mathcal{P}'', P''$), and so on, until COARSENING($\mathcal{X}, Y$) outputs the partition $\mathcal{X}$ and the profile Y for some $\mathcal{X}$ and Y. Finally, apply SVEI-INCOMPLETE(Y) and order P in accordance with the obtained output.

COMPLETING(P): Repeat the steps below for all rows $i \in \{1, \ldots, m\}$. Take the smallest $j \in \{1, \ldots, n\}$ such that $p_{i,j} \notin \{0, 1\}$ and check whether there exists $\ell > j$ with $p_{i,\ell} \in \{0, 1\}$. If it does, then find the smallest such ℓ and set P to $P[p_{i,j}|p_{i,\ell}]$. Otherwise, check whether there exists $\ell < j$ with $p_{i,\ell} \in \{0, 1\}$. If it does, then find the largest such ℓ and set P to $P[p_{i,j}|p_{i,\ell}]$. Otherwise, set P to $P[p_{i,j}|0]$. Repeat for the next cell of unknown value in this row.

Proposition 7. *`VEI-INCOMPLETE` detects in polynomial time whether a profile of dichotomous preferences possibly satisfies VEI. If it does, `VEI-INCOMPLETE` also finds an appropriate order of the voters and `COMPLETING` finds a suitable completion in polynomial time.*

Proof. Clearly, `VEI-INCOMPLETE` will terminate in polynomial time: One application of `EXPANDING` takes at most $O(m^2n^2)$ time, and thus one application of `COARSENING` takes $O(m^3n^2)$ time. So, the coarsest partition will be obtained in $O(m^4n^2)$ time. Then, `SVEI-INCOMPLETE` will be applied, which we know takes $O(m^2n^2)$ time. It is also easy to see that `COMPLETING` takes $O(mn^2)$ time.

Moreover, it follows from the relevant definitions that a profile P satisfies VEI if and only if the profile P_s obtained from P by cancelling certain rows satisfies SVEI. Suppose that, after some application of `COARSENING`, two alternatives a_i and a_k are in the same set of the partition in the output. This means that in P_s we must have either both or neither of the rows i and k cancelled. Now, if a_i and a_k contradict each other, then the forbidden configuration for SVEI will occur as a pattern in P_s, and thus P_s will fail SVEI and P will fail VEI. But suppose that, via `VEI-INCOMPLETE`, we obtain the coarsest possible partition of the set of alternatives and no alternatives from two different sets contradict each other; if moreover no alternatives in the same set contradict each other, then the resulting profile P_s will not contain any forbidden configuration for SVEI as a pattern. So P_s will satisfy SVEI, and thus P will satisfy VEI. Then, in the order obtained by `SVEI-INCOMPLETE` applied in P, for every alternative a_i, the voters that approve it appear in the ordering either before or after the voters that disapprove it. This property will be preserved after `COMPLETING`. □

Example 2 demonstrates `VEI-INCOMPLETE` in practice.

Example 2. Consider a profile P of dichotomous preferences as depicted in (a) of Figure 5. Starting with step (b), we apply the coarsening algorithm for P on the finest partition $\{\{a_1\},\ldots,\{a_5\}\}$. We first examine the singleton set $\{a_1\}$ and expand it into $\{a_1, a_4\}$ because the rows 1 and 4 of the matrix match. In the second iteration of expansion, we examine the set $\{a_2\}$ that is not used yet and create the new set $\{a_2, a_3\}$ because the rows 2 and 3 match. Then, the set $\{a_5\}$ remains unused. But the rows 5 and 3 contradict each other, so the set that a_3 belongs to, namely $\{a_2, a_3\}$, will be cancelled and the singleton $\{a_5\}$ will be expanded into $\{a_2, a_3, a_5\}$. After that round, the first application of coarsening will be terminated; its second application will involve the partition $\{\{a_1, a_4\}, \{a_2, a_3, a_5\}\}$. Inspecting the set $\{a_1, a_4\}$, we see that row 4 matches the cancelled row 3 (since they were contradicting each other in the original matrix). So, all alternatives in the set $\{a_2, a_3, a_5\}$ will now join $\{a_1, a_4\}$, giving rise to the coarsest possible partition. Step (d) follows from the algorithm `SVEI-INCOMPLETE`, while steps (e) and (f) are straightforward. △

Finally, Proposition 8 is implied by Proposition 2. As explicitly shown in Lemma 2 and Proposition 6, the check in fact concerns many more profiles than the five from Proposition 2 to accomodate all possibilities with unknown cells.

$$\begin{pmatrix} ? & ? & ? & ? & 1 & 0 \\ ? & ? & 0 & 1 & ? & ? \\ 0 & 1 & 0 & 1 & ? & ? \\ 1 & 0 & ? & ? & 1 & 0 \\ 1 & 1 & ? & 0 & ? & ? \end{pmatrix}$$

(a) original profile

$$\begin{array}{l} \{\{a_1\},\{a_2\},\{a_3\},\{a_4\},\{a_5\}\} \\ \{\{a_1,a_4\},\{a_2\},\{a_3\},\{a_5\}\} \\ \{\{a_1,a_4\},\{a_2,a_3\},\{a_5\}\} \\ \{\{a_1,a_4\},\{a_2,a_3,a_5\}\} \\ \{\{a_1,a_2,a_3,a_4,a_5\}\} \end{array}$$

(b) coarsening (via expanding four times)

$$\begin{pmatrix} ? & ? & ? & ? & 1 & 0 \\ ? & ? & 1 & 0 & ? & ? \\ 1 & 0 & 1 & 0 & ? & ? \\ 1 & 0 & ? & ? & 1 & 0 \\ 1 & 1 & ? & 0 & ? & ? \end{pmatrix}$$

(c) new profile (rows 2 and 3 cancelled)

$$\begin{pmatrix} ? & 1 & ? & ? & ? & 0 \\ ? & ? & 1 & ? & 0 & ? \\ 1 & ? & 1 & 0 & 0 & ? \\ 1 & 1 & ? & 0 & ? & 0 \\ 1 & ? & ? & 1 & 0 & ? \end{pmatrix}$$

(d) ordering (via `SVEI-INCOMPLETE`)

$$\begin{pmatrix} ? & 1 & ? & ? & ? & 0 \\ ? & ? & 0 & ? & 1 & ? \\ 0 & ? & 0 & 1 & 1 & ? \\ 1 & 1 & ? & 0 & ? & 0 \\ 1 & ? & ? & 1 & 0 & ? \end{pmatrix}$$

(e) ordering of original profile

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

(f) completing

Figure 5: `VEI-INCOMPLETE`: An example.

Proposition 8. *We can check whether an incomplete profile of dichotomous preferences necessarily satisfies VEI in polynomial time.*

5 Conclusion

We have studied structured dichotomous preferences from two angles. First, in settings with *complete* information, we have obtained forbidden subprofile characterisations that expand previous work on the topic. Second, in cases with *incomplete* information, our main challenge concerned detecting whether an incomplete profile can admit a *possible* or a *necessary* completion within certain restricted domains. Our questions were answered by locating results in the relevant literature, as well as by designing new constructive algorithms. Yet, one of them remains open: What is the complexity of determining whether the completion of an incomplete dichotomous preference profile necessarily satisfies VI?

While the beginning has been made, restricted domains of dichotomous preferences are certainly worthy of further investigation. For complete profiles that violate a given property, it would be interesting to know the minimum number of values that need to be swapped to make the property hold. Such a result would be relevant for applications with *noisy* inputs—some pertinent work has been recently conducted by Rani, Subashini, and Jagalmohanan (2019). Then, regarding cases of incomplete information, one could also study the probability of a profile to comply with a certain structure after being randomly completed, or limit the number of cells with unknown values and conduct a parametrised complexity analysis. All these problems should of course be explored with respect to different domain restrictions as well.

Acknowledgments

Support from the Basic Research Program of the National Research University Higher School of Economics is gratefully acknowledged. We also thank Ronald de Haan for his valuable feedback during the writing process of this paper.

References

Ballester, M. A.; and Haeringer, G. 2011. A characterization of the single-peaked domain. *Social Choice and Welfare* 36(2): 305–322.

Bartholdi III, J.; and Trick, M. A. 1986. Stable matching with preferences derived from a psychological model. *Operations Research Letters* 5(4): 165–169.

Brams, S. J.; Kilgour, D. M.; and Sanver, M. R. 2007. A minimax procedure for electing committees. *Public Choice* 132(3-4): 401–420.

Bredereck, R.; Chen, J.; and Woeginger, G. J. 2013. A characterization of the single-crossing domain. *Social Choice and Welfare* 41(4): 989–998.

Chen, J.; Pruhs, K. R.; and Woeginger, G. J. 2017. The one-dimensional Euclidean domain: Finitely many obstructions are not enough. *Social Choice and Welfare* 48(2): 409–432.

Conitzer, V. 2009. Eliciting single-peaked preferences using comparison queries. *Journal of Artificial Intelligence Research* 35: 161–191.

Domshlak, C.; Hüllermeier, E.; Kaci, S.; and Prade, H. 2011. Preferences in AI: An overview. *Journal of Artificial Intelligence Research* 19(4): 1037–1052.

Elkind, E.; Faliszewski, P.; Lackner, M.; and Obraztsova, S. 2015. The complexity of recognizing incomplete single-crossing preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, 865–871.

Elkind, E.; Faliszewski, P.; and Slinko, A. M. 2012. Clone structures in voters' preferences. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, 496–513.

Elkind, E.; and Lackner, M. 2014. On detecting nearly structured preference profiles. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, 661–667.

Elkind, E.; and Lackner, M. 2015. Structure in dichotomous preferences. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 2019–2025.

Elkind, E.; Lackner, M.; and Peters, D. 2017. Structured preferences. In Endriss, U., ed., *Trends in Computational Social Choice*, 187–207. AI Access.

Erdélyi, G.; Lackner, M.; and Pfandler, A. 2017. Computational aspects of nearly single-peaked electorates. *Journal of Artificial Intelligence Research* 58: 297–337.

Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. A. 2014. The complexity of manipulative attacks in nearly single-peaked electorates. *Artificial Intelligence* 207: 69–99.

Faliszewski, P.; Skowron, P.; Slinko, A.; and Talmon, N. 2017. Multiwinner voting: A new challenge for social choice theory. *Trends in Computational Social Choice* 74: 27–47.

Fitzsimmons, Z.; and Lackner, M. 2020. Incomplete preferences in single-peaked electorates. *Journal of Artificial Intelligence Research* 67: 797–833.

Ganian, R.; Kanj, I.; Ordyniak, S.; and Szeider, S. 2018. Parameterized algorithms for the matrix completion problem. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*.

Golumbic, M. C. 1998. Matrix sandwich problems. *Linear algebra and its applications* 277(1-3): 239–251.

Jaeckle, F.; Peters, D.; and Elkind, E. 2018. On recognising nearly single-crossing preferences. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, 1079–1086.

Kahn, A. B. 1962. Topological sorting of large networks. *Communications of the ACM* 5(11): 558–562.

Kilgour, D. M.; and Marshall, E. 2012. Approval balloting for fixed-size committees. In *Electoral Systems*, 305–326. Springer.

Klinz, B.; Rudolf, R.; and Woeginger, G. J. 1995. Permuting matrices to avoid forbidden submatrices. *Discrete applied mathematics* 60(1-3): 223–248.

Lackner, M. 2014. Incomplete preferences in single-peaked electorates. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, 742–748.

Peters, D.; and Lackner, M. 2020. Preferences single-peaked on a circle. *Journal of Artificial Intelligence Research* 68: 463–502.

Procaccia, A. D.; Rosenschein, J. S.; and Zohar, A. 2007. Multi-winner elections: Complexity of manipulation, control and winner-determination. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*, 1476–1481.

Rani, M. R.; Subashini, R.; and Jagalmohanan, M. 2019. Simultaneous consecutive ones submatrix and editing problems: Classical complexity and fixed-parameter tractable results. *Theoretical Computer Science* .

Ryser, H. J. 1957. Combinatorial properties of matrices of zeros and ones. *Canadian Journal of Mathematics* 9: 371–377.

Tucker, A. 1972. A structure theorem for the consecutive 1's property. *Journal of Combinatorial Theory, Series B* 12(2): 153–162.

Walsh, T. 2008. Complexity of terminating preference elicitation. In *Proceedings of the 7th international joint conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 967–974.

Yang, Y. 2019. On the tree representations of dichotomous preferences. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*.