2015 IEEE 27th International Conference on Tools with Artificial Intelligence

Multi-Round Vote Elicitation for Manipulation under Candidate Uncertainty

Manel Ayadi* and Nahla Ben Amor†

LARODEC, Institut Supérieur de Gestion, Université de Tunis
41 Avenue de la liberté, 2000 Le Bardo, Tunisie
**Email: manel.ayadi@hotmail.com*
†*Email: nahla.benamor@gmx.fr*

***Abstract*—Manipulation becomes harder when manipulators are uncertain about the preferences of sincere voters. Elicitation may communicate information, of sincere voters' votes, to a manipulator, allowing him to vote strategically. In this paper, a multi-round elicitation process, of sincere voters' preferences, is derived that yields to an optimal manipulation with minimal information elicited. Through in-depth experimental study, this paper answers the question: *How many candidates, per sincere voter, are needed to be known for an optimal manipulation?* Probabilistic models such as *IC* and *SP-IC* are used to complete preference profiles.**

***Keywords*-manipulation; vote elicitation; partial preferences; uncertainty.**

I. Introduction

Voting is a standard mechanism of reaching a joint decision. In this context, voting rules are used to aggregate the voters' preferences over a set of candidates standing for election to make a socially desirable decision. A significant problem in social choice is that there is generally one or more voters that can obtain a more desirable outcome by misreporting their preferences. This is called *manipulation* or *strategic voting*. Manipulation is an undesirable phenomenon leading to fairness issues. A seminal negative result stated by Gibbard-Satterthwaite [1], [2], shows that voting rules that are not manipulable do not exist.

Analysis of manipulation suffers from two significant deficiencies. The first one is related to the assumption behind manipulator's knowledge. In fact, most of existing work [3], [4], is confined to case in which manipulators know exactly the non-manipulators' votes. Clearly, this assumption does not cope with real world situations where manipulators are generally uncertain about sincere voters' votes, or even completely ignorant about their votes. Recently, this assumption was relaxed in [5], where authors extended coalitional manipulation to incomplete knowledge. Second, manipulation analysis underlines the probability of manipulation, ignoring its impact on social welfare. However, characterizing the impact of a manipulating coalition's action only in terms of its probability of succeeding can sometimes be ambiguous, leading to a possible gap between the social welfare of the optimal alternative and the social welfare of the one that is ultimately elected under manipulation. To the best of our knowledge, the unique work that considers the impact of manipulation on social welfare is [5]. This paper proposes to overcome these two problems by studying coalitional manipulation problem where the manipulators are uncertain about the non-manipulators' votes, and assessing the loss on social welfare.

Another important problem, related to manipulators' knowledge under uncertainty, is the amount of information revealed by sincere voters. In this context, eliciting the pertinent information from sincere voters, allows the manipulator to ensure the victory of a specific candidate. While preference elicitation and manipulation are closely related, unfortunately, previous investigations have mainly focused on each field separately. This paper emphasizes the connection between strategic voting and preference elicitation. To this end, we consider preference elicitation to determine when to stop eliciting information from sincere voters as the manipulation becomes tractable. We propose a new manipulation strategy that yields to an optimal manipulation by solving the restricted manipulators' knowledge using a multi-round elicitation process. Uncertainty regarding the voter preferences is facilitated by using probabilistic preference models such as *IC* and *SP-IC*. This paper answers the question: *How many candidates, per sincere voter, are needed to be known for a successful manipulation?*

This paper is organized as follows: In Section 2, we give some basic background on voting procedures and manipulation model. In Section 3, after briefly discussing deficiencies in preference elicitation, we propose a new multi-round elicitation process that allows the voters to incrementally send their preferences, converging to a winning candidate. In Section 4, we describe a new manipulation strategy, which solves manipulators' restricted knowledge, using the proposed multi-round elicitation process. Finally, Section 5 is dedicated to the experimental study.

II. Background on voting procedures and manipulation

A. Voting procedures

We define a *partial* vote by an election $E' = (N, A, \Pi)$ where: $N = \{1, ..., n\}$, is the set of *voters*; $A = \{a, b, ...\}$, is the set of *candidates* such that $|A| = m$; and $\Pi =$

1082-3409/15 $31.00 © 2015 IEEE
DOI 10.1109/ICTAI.2015.113

$\{\pi_1, ..., \pi_n\}$, is the *partial preference profile* of voters in N. Let π_i denote the partial preference order of voter i. Voter i's preferences are presented by a partial ranking $\succ_i$ over A. When $a \succ_i b$ for some $a, b \in A$, said that voter i prefers a to b.

A *completion* of π_i is any vote $\succ_i$ that extends π_i. Let $C(\pi_i)$ denote the set of *completions* of π_i , i.e. the set of all complete votes $\succ_i$ that extends π_i. Let $C(\Pi) = C(\pi_1) \times ... \times C(\pi_n)$ be the set of completions of Π. We refer to $E = (N, A, \succ^N)$, an election obtained after completion, where $\succ^N$ denote the complete voters' votes. A *voting rule* is a procedure for making a choice from the set of candidates. Formally, given an election $E = (N, A, \succ^N)$ as input, a *voting rule* f is a function $f : E \rightarrow S$ that outputs a non-empty subset $S \subseteq A$. The elements of S are called the winners of the election E under f. If $|f(E)| > 1$ for any election E, the mapping f is called a *voting correspondence*.

A *positional scoring function* $s : \{1, ..., m\} \mapsto \Re_{\geq 0}$ maps ranks onto scores such that $s_1 \geq ... \geq s_m$ defines a scoring rule over a set of candidates of size m. A candidate receives s_j points from each voter who ranks him in the j^{th} position, and the score of a candidate is the total number of points she receives from all voters. In the remaining, we focus on the Borda score to evaluate the outcome of the election where $s = (m-1, m-2, ..., 1, 0)$.

B. Manipulation

According to [1], [2], *manipulation* refers to one or more voters that may declare preferences, that are not their true ones, with the aim of obtaining a better outcome for themselves. Formally, in a *coalitional manipulation problem*, the N voters are divided into two groups, the manipulators and the non-manipulators (aka. *sincere* voters). We define our manipulation problem as follows: For any voting rule f, an instance $I = (H, M, A, \succ^N, p)$ is given by an election $E = (N, A, \succ^N)$ where: $H = \{1, ..., h\}$ is the set of *sincere voters*; $M = \{1, ..., (n-h)\}$ is the set of *manipulators*; $A = \{a, b, ...\}$ is the set of *candidates* where $|A| = m$; $\succ^N = (\succ^1, ..., \succ^h)$ presents the *preference profile* of voters in H; and a distinguished candidate $p \in A$. The goal is to answer whether the M manipulators, knowing the preferences of the H sincere voters, can provide a preference such that p will be chosen by f.

Studying manipulation problem under incompleteness or uncertainty regarding the votes of sincere voters has received little attention. In fact, most of existing work [3], [4] assume that the manipulating coalition has a complete knowledge of the sincere voters' votes, however, this assumption is unrealistic in the real world setting where manipulators have rarely access to such information. In this paper, we relax this assumption by studying unweighted constructive coalitional manipulation problem, in the setting in which all manipulators have identical preferences, and they are uncertain about the non-manipulators' votes. We represent this uncertainty by a *top-k* most preferred candidates over the sincere voters' votes.

The impact of manipulation on social welfare provides a different analysis of manipulability of several voting rules. While the probability of manipulation can inform us on the chances given to manipulators to change the outcome of the election, however, alternatives with higher success probability can have an undesirable impact on social welfare causing less societal satisfaction. In this context, social welfare for alternative p chosen by the manipulators must be close to that of the optimal (non-manipulated) alternative, which in turn means that the *damage* in social welfare caused by manipulation will itself be limited. Formally, we refer to $SW(s, s^*)$ as the difference between the true score s^* and the approximate score s associated to a set of candidates under a given scoring rule and propose analyzing rules in this light. In this paper, we consider both the probability of manipulation and its impact on social welfare, by computing the *optimal manipulation* strategy which refers to minimize the loss on social welfare and ensure a successful manipulation where manipulators increase their chances of winning.

Characterizing manipulating coalition's action by studying the impact of manipulation on social welfare represents a solution concept to minimize manipulation's side effect. In a different direction, manipulation's analysis tends to focus on the problem of equilibrium selection (e.g. Nash Equilibrium) as a potential solution concept for preference aggregation scenarios [6].

III. VOTE ELICITATION

Efficient vote elicitation concerns eliciting only pertinent information from the voter because expressing complete preferences can be arduous. In this context, *top-k* voting is an especially natural form of partial vote elicitation in which only length k prefixes of rankings are elicited. Authors in [7] show that eliciting partial order preferences from the voters can allow the voting protocol to determine the outcome well before all of the preferences have been elicited. In fact, much work has mainly focused on theoretical analysis of preference elicitation, reporting the upper and lower bounds of the required communication with the voters [8]. While these analyses are efficient to determine the right outcome, however, they may sometimes elicit more information than necessary in practice. Therefore, preference elicitation that are driven by more practical considerations are required. However, to the best of our knowledge, only two studies propose practical algorithms for effective vote elicitation [9], [10], still the study of practical frameworks for elicitation has received little attention.

A. Probabilistic elicitation models

Given partial information about voters' preferences, several approaches can be used to select the winner of an election [11]. Within these approaches, we can mention probabilistic ranking models, which require the specification of some prior distribution over voter preferences. The most adopted probabilistic model in social choice, is *impartial culture model (IC)*, which assumes the preference of any voter are drawn from the uniform distribution over the set of candidates. A related model is the *impartial anonymous culture (IAC)*, in which each voting situation is equally likely.

Among probabilistic analysis used in social choice, there is also an increasing focus on identifying tractable special cases using domain restrictions [12]. One of the most common domain restriction considered in social choice theory is that of *single peaked preferences*. The *single peaked impartial culture (SP-IC)* model generates singled peaked votes, consistent with a given social axis, uniformly from the IC model. A variety of other models have been proposed that reflect different interpretations of the ranking process (e.g., Plackett-Luce, Bradley-Terry, Mallows, etc.). In this paper, we investigate IC and SP-IC models empirically below.

B. Multi-round vote elicitation process

We propose a *multi-round elicitation process*, in the form of *top-k* preferences, interested in: (1) allowing the voter to incrementally send his partial preferences in rounds, and (2) aiming to reduce the number of interactions with the voter, as well as the amount of information elicited, to ensure that either a winning or a high quality candidate can be determined with high probability. The key idea is that the voters incrementally send their preferences in rounds, one preference each round, in a decreasing order of their preferred candidates. After each round, the voters wait for the voting center to decide whether sufficient information has been received to determine a winner, or additional voters' preferred candidates are needed to be sent in the next round. In each round the response query of each voter is added to all his previous votes and the center evaluates the outcome of the election using a specific scoring rule, and based on the available information. The multi-round process stops eliciting preferences once the approximate winner is determined with high probability, and from that point the voters no longer send preference values.

Formally, we refer to a *query* by a single request for information from a voter. We consider type of query in the form of *top-k* most preferred candidates (e.g., *who is your top-k preferred candidate?*). Given a particular class of queries Q, a *multi-round* voting protocol selects, at each round, a subset of voters, and one query per selected voter. Let I_{t-1} be the *information set* available at round *t*. I_{t-1} represents the voters' responses to queries for rounds $t \in \{1, ..., (t-1)\}$, where the elicitation of *top-k* preferences is simulated for $k \in \{1, ..., (m-1)\}$; allowing the voters to rank $k = m-1$ candidates in the worst case, since the complete order ranking is not considered. The information set at round t is added to all his previous responses. Let π^{t-1} be a partial profile of *top-k* preferred candidates for each voter. Then, given an election $E' = (N, A, \Pi)$, we propose a protocol based on three functions:

1) *Querying function* ψ, represents a sequence of mappings $\psi_t : \Pi \longmapsto (N \longmapsto Q)$, selecting for each voter a single query at stage t given the current information set I_{t-1}.
2) *Completion function* $C : \Pi \longmapsto V$, given a partial profile Π, $C(\Pi)$ is the set of consistent extensions of Π to obtain a full ranking profile. The output of C, is a complete ranking preference profile V over the set of candidates. Probabilistic models are used to deal with uncertainty regarding the voter preference and to complete the profile.
3) *Winner selection function* $\omega : V \longmapsto A$, where $\omega(V)$ denotes the approximate winner under a given complete profile V. Borda scoring rule is used to evaluate the outcome given the voters' complete votes V, and to determine the approximate winner over the set of candidates A. If the approximate winner is far from the true one, the protocol proceeds to round $t+1$, otherwise the protocol terminates with the chosen winner at round t.

Given a protocol $\Omega = (\psi, C, \omega)$, in order to evaluate the effectiveness of the multi-round elicitation process, we propose to use three properties, namely:

- *Quality of the winner*: Determine the quality of the approximate winner of the election, using partial information elicited in each round. While the Borda rule is used, the output of Ω is a scoring vector over the set of candidates. The approximate winner is evaluated by considering how far from the true one could be.
- *Amount of elicited information*: Determine the necessary amount of elicited information from voters in each round, in order to ensure a high winner quality. The amount of information will be measured by counting the number of k preference elicited in each round.
- *Number of rounds*: Determine the number of rounds needed, in order to obtain a high quality of the winner.

Algorithm 1 presents the multi-round elicitation process. The *querying function* is performed in lines (1-5) where partial preference orders are presented in ${\pi_i}^k$, which denotes the *top-k* ranking of voter i over a set of m-1 candidates. ψ_t contains the partial profile elicited from voters in each round. The response query of each voter is added to all his previous votes in I_t. The *completion function* is presented in line 6 where the partial profile is extended to a complete one V. Then, the *winner selection function* is introduced in lines (7-10) where the Borda rule is applied. If the approximate winner $\omega(V)$, obtained with the predicted profile V, is far from the true one ω^*; then, the algorithm proceeds to the next round (line 8) by eliciting the next most preferred candidate of each voter (line 9); otherwise, the approximate winner is declared the winner of the election (line 10).

Algorithm 1. Multi-round elicitation process

Input: N $voters$, m $candidates$, *Top-k votes:* $\left(\pi_1{}^k, ..., \pi_n{}^k\right)$
Output: $\omega\left(V\right)$ $(approximate\ winner)$
Initialization: t=1, k=1, I_t= $\emptyset$, ω^* $(true\ winner)$

1: **while** ($(\omega\left(V\right) \neq \omega^*)$ OR $(t < m)$) **do**
2:   **for** $i \in N$ **do**
3:     $\psi_t.add\left(\pi_i{}^k\right)$
4:   **end for**
5:   $I_t.add\left(\psi_t\right)$
6:   $V \leftarrow C\left(I_t\right)$
7:   **if** $(\omega\left(V\right) \neq \omega^*)$ **then**
8:     $t = t + 1$
9:     $k = k + 1$
10:   **else**
    $return\ \omega\left(V\right)$
11:   **end if**
12: **end while**

IV. New manipulation model

In this section, we propose to study the connection between strategic voting and preference elicitation. Elicitation may reveal information about honest voters' votes to a manipulator, which allows him to vote strategically. To this end, we introduce a *new manipulation strategy with uncertain knowledge*, to solve restricted manipulators' knowledge, regarding the non-manipulators' votes, using the proposed multi-round elicitation process. In this context, the multi-round elicitation process will elicit the *top-k* preferences of non-manipulators. This amount of information elicited presents the manipulators' knowledge, which will allow them to change the outcome to their favor. Using this new approach, we aim to:

- Determine the minimal value of k needed to be known by the manipulators to achieve their goal, with the minimal number of rounds.
- Decide when to stop eliciting sincere voters' preferences, as an optimal manipulation is guaranteed.

While manipulation and elicitation are two closely related fields, unfortunately, previous investigations have mainly focused on each field separately. To the best of our knowledge, only two studies introduce the connection between them [7], [12], but no practical process has been proposed. The *new manipulation strategy with uncertain knowledge* will allow us to overcome existing weaknesses. This new approach uses the three functions defined in the multi-round elicitation process, namely: *querying function*, *completion function* and *winner selection function*; in order to elicit preferences from sincere voters and to determine the outcome of the election by considering the manipulators' votes.

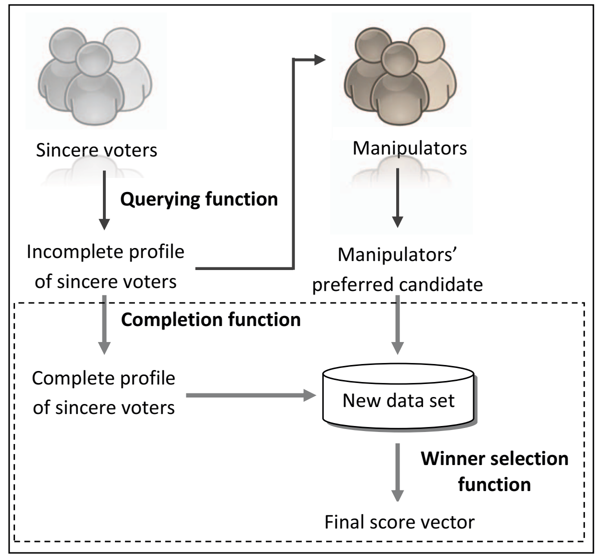

Figure 1. Proposed approach description

The whole process is illustrated in Figure 1 and proceeds as follows: In the first round (*t=1*), the *querying function* will elicit *top-1* sincere voters' preferences. Each sincere voter expresses his most preferred candidate over the available ones. Then, the information set elicited is sent to the manipulators, which constitutes their knowledge at the first round. Based on these incomplete preferences, the manipulators will cast a vote that makes their preferred candidate p win the election. In order to complete the sincere votes, we consider an internal process where the *completion function* is performed. In this step, sincere votes are extended using probabilistic models. A complete preference profile V over the whole set of candidates is provided in output. Now, given the complete predicted profile V and the manipulators' desired candidate p, gathered on the *new data set*; the *winner selection function* is performed, where Borda rule is applied. The output of this step is an approximate winner $\omega\left(V\right)$ which will be compared to the true one ω^*, in order to evaluate the impact of manipulation i.e. verify if the manipulators have succeeded to change the outcome of the election to their favor. In order to evaluate the performance of the protocol to elicit the minimal value of k from the sincere voters, we measure both, the probability of manipulation and its impact on social welfare. If no *optimal manipulation* is found, the process proceeds to the next round by eliciting the next most preferred candidate of the sincere voters, until an optimal strategy is detected.

The use of the multi-round elicitation process on the new manipulation strategy, will allow us to investigate the relation between the amount of information revealed by sincere voters and the probability of manipulation; by studying how a reduction or increase in uncertainty may change a strategic vote. To this end, our goal is to answer the question: *How many candidates, per sincere voter, are needed to be known for an optimal manipulation?*

V. Experimental study

To explore the effectiveness of the proposed approach, we run a suit of experiments with top-k preference distributions, as well as real data sets. Our interest is two-folds: (i) evaluate the effectiveness of the multi-round elicitation process for determining the optimal value of k with quality guarantees. We focus on IC and SP-IC models to complete the preference order, however, our framework is not limited to such models; (ii) evaluate the performance of the new manipulation strategy by measuring both the probability of manipulation and its impact on social welfare. Experiments are performed using two data sets, namely: *Sushi data* [13] from *Preflib* [14], having 10 alternatives and 5000 rankings. Since this latter is not a single peaked data, we have randomly generated a data set so called *Ran-Gen data* with 10 alternatives and 1000 voters having a single peaked preferences.

In order to test the performance of the multi-round elicitation process, our first set of experiments consider *Sushi* data (resp. *Ran-Gen* data), with 10 preference profiles where in each profile we simulate the elicitation of *top-k* preference orders $k \in \{1, ..., 10\}$ with N=5000 rankings (resp. N=1000). Figure 2 summarizes the experiments performed with *Sushi* data, by showing, in plot (a) (resp. (b)), the margin of victory of 7 preference profiles, by increasing the number k of alternatives elicited $k \in \{1, ..., 7\}$ in each profile; using IC model (resp. SP-IC). The horizontal axis presents the set of different candidates, and the vertical axis shows the score associated to each one of them. The peaks in each preference profile, in a decreasing order, present the order of the winners using the Borda rule. For instance, $k = 1$ means that in the first profile, *top-1* queries are elicited from each voter over a set of 10 alternatives. We refer to $k = 10$, by the complete profile where *top-10* alternatives are elicited, which will be used to evaluate the quality of the winner.

Clearly, under IC model (plot a), the curves' shape is the same, which means that even with partial preferences, the scoring rule ranks the alternatives, almost, in the same order. Results with *Sushi* data using IC model, show that it is always possible to determine the winner with incomplete information i.e. the highest peak with different values of k, is always located in alternative 8 which represents the true winner (compared to $k = 10$). More precisely, *top-1* queries are sufficient to determine an approximate winner with high quality. More interestingly, when we consider the rank of different alternatives, results suggest that eliciting preferences from voters in 2 rounds (top-2), is always sufficient

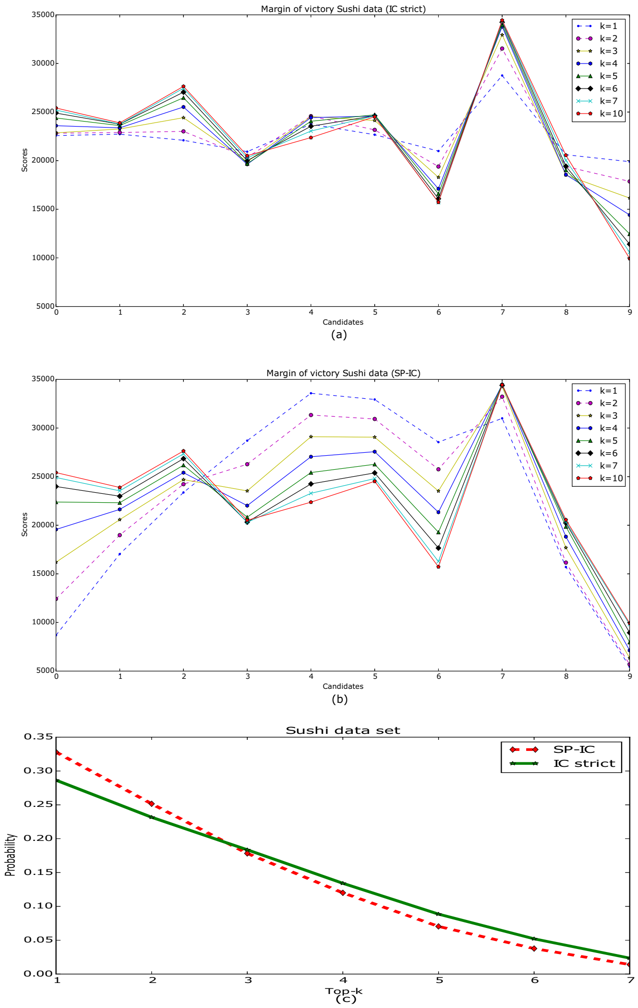

Figure 2. (a) Sushi results with IC model, (b) Sushi results with SP-IC, (c) Probability of the approximate winner quality.

to guarantee not only a correct outcome, but also a correct ranking of *4/10* alternatives, which minimize the damage on social welfare. *Top-4* queries are a very good approximate to rank all alternatives in the same order. Results with *Sushi* data using SP-IC model (plot b), show that the curves' shape is not always the same, which means that the margin between the different scores is sparse. For small value of $k = 1$, the approximate winner deviates from the true one i.e. with $k = 1$, the approximate winner is alternative 5 while the true winner is alternative 8. Results with SP-IC model, suggest that, *top-2* queries, are usually enough to obtain the desired outcome. However, *top-6* queries are sufficient to guarantee less damage on social welfare.

To conclude, the IC model performs well with *Sushi* data with small value of $k = 1$ than the SP-IC model. This is shown in the 3^{rd} plot (plot c), where the probability of deviation of the true ranking is of 27% using IC model against 34% with $SP\text{-}IC$ model. With top-3 queries, the true ranking is determined in the most cases with a slight deviation of 17%. Intuitively, while increasing the values of k, the process converges to the correct prediction, and is near perfect with top-3 using the two models (deviation of 15%).

Similarly to the above experiments, Figure 3 summarizes experiments performed on *Ran-Gen* data. The use of the IC model on this data (plot a), shows that the true winner is always determined except with $k = 2$ i.e. with *top-2* queries, the approximate winner is alternative 6 while the true one is the alternative 5. The curves' shape, with different values of k, is almost the same except with $k =$ 1-2 i.e. an additional alternative (top-3 queries) is sufficient to determine, not only the true winner, but also a correct ranking of different alternatives. A very good approximation with less damage on social welfare, is obtained with top-4 queries. Results with $SP\text{-}IC$ are even more illuminating (plot b), where the curves' shape is always the same with all values of k i.e. not only the approximate winner is determined with high quality, but also the scores are the same with any incomplete profile. In other words, any value of $k \in \{1, ..., 9\}$, is sufficient to determine the true winner with the lowest loss on social welfare, which shows the performance of $SP\text{-}IC$ model with a single peaked data set. This is also clear in (plot c), where the probability of deviation from the true ranking with $SP\text{-}IC$ model is evaluated to 14% with all the incomplete preferences $k \in \{1, ...7\}$ i.e. a negligible impact on social welfare. However, with the IC model, the probability of deviation from the true ranking is evaluated to 30% with *top-1* queries and decreases by increasing the number of queries.

Now, let us turn our attention to the normal form of the distribution generated by *Ran-Gen* data. Since this latter is a single peaked data, our results suggest that using IC and $SP\text{-}IC$ models, to complete an original single peaked data, generate always a single peaked vote able to determine the correct ranking of alternatives in almost all cases. For instance, plots (a) and (b), present a local maximum on the most preferred alternative of the voters i.e. alternatives drew along the horizontal axis present a local maximum on alternative 4 which represents the winner with the highest score; and then the curve goes down with the other alternatives in a decreasing order.

In order to test the effectiveness of the new manipulation approach, we measure both the probability of manipulation and its impact on social welfare. In this context, our second set of experiments consider *Sushi data* by drawing 10 profiles such that, each one of them contains *top-k* preference orders $k \in \{1, ..., 10\}$. We vary the number of manipulators $M \in \{750, 1000\}$ with *N= 5000* voters in each profile; where *M= 750* manipulators (resp. *M= 1000*), represents 15% (resp. 20%) of the total number of voters. We use

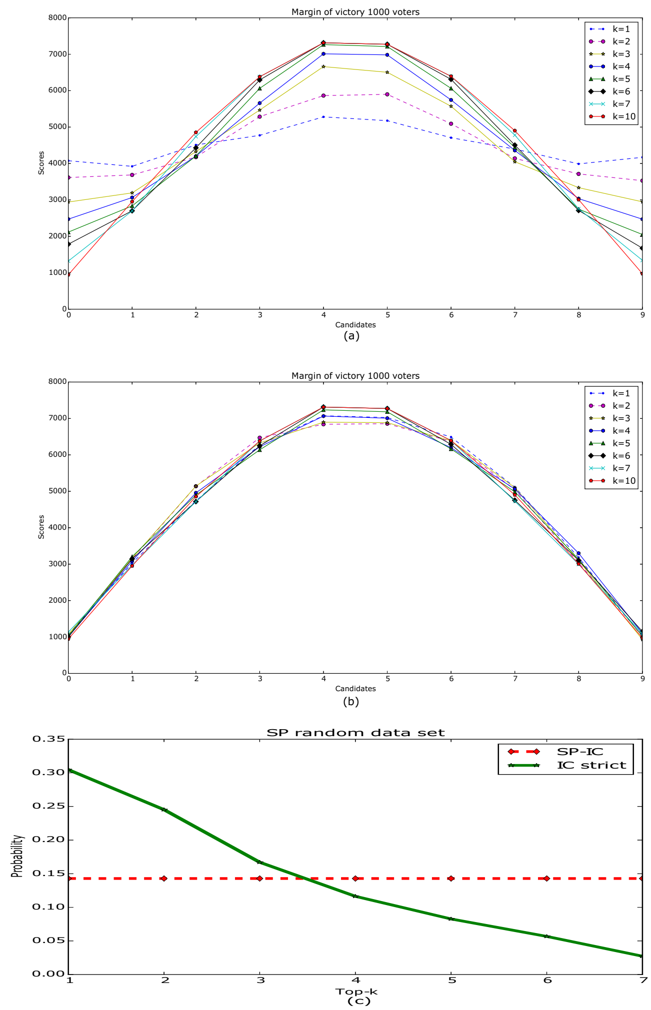

Figure 3. (a) Ran-Gen with IC model, (b) Ran-Gen results with $SP\text{-}IC$, (c) Probability of the approximate winner quality.

the IC model to complete non-manipulators' votes, since its effectiveness performed with *Sushi data* in the above experiments. Since manipulability varies greatly with the preferred alternative chosen by the manipulators, we show results for the alternatives whose expected ranks in different profiles are second, third, and fourth i.e. *P=2, P=3, P=4*.

Figure 4 shows results of three different manipulators' strategies with *M=750*. In plot (a) (resp. plot (b), plot (c)), the manipulators' preferred alternative is the one whose expected rank in different profiles is second, *P=2* (resp. *P=3, P=4*). As in the above experiments, the horizontal axis presents the different alternatives and the vertical axis illustrates the score associated to each one. We show results of 5 incomplete preference profiles by varying the value of $k \in \{1, ..., 5\}$ in each profile. For instance, in plot (a), *K=5* means that the manipulators know only top-5 preferences of sincere voters; based on these incomplete information, their aim is to ensure the victory of the alternative whose expected rank is second (*P=2*). In order to evaluate the effectiveness of different manipulators' strategies; in each plot, *'opt'* refers to the true ranking of alternatives, with complete preferences and without manipulation, which is: *8, 3, 1, 6, 2, 5, 4, 9, 7, 10*. The different ranks are presented by the peaks in the graph.

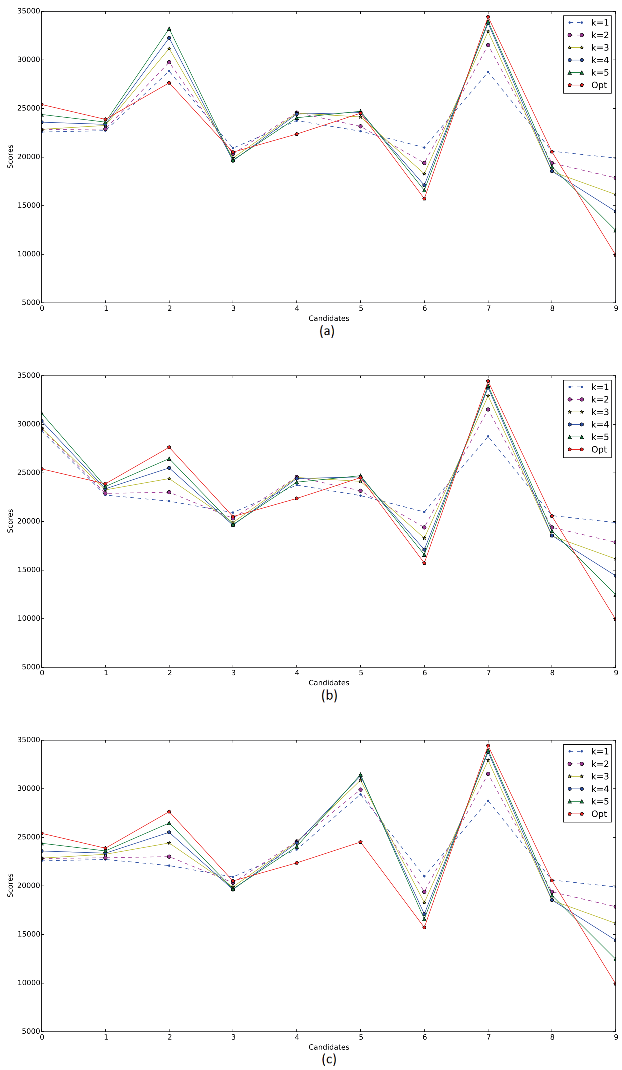

Figure 4. (a) Manipulation strategy with P=2 and $M = 750$, (b) Manipulation strategy with P=3 and $M = 750$, and (c) Manipulation strategy with P=4 and $M = 750$

Interesting results are presented when we consider *M=750* manipulators (Figure 4) with *top-1* queries, where, with different manipulation strategies and knowing only one preference of the non-manipulators' votes, the manipulators can ensure the victory of their most preferred alternative i.e. when *P=2* and with *top-1* queries, the winner is candidate 3 which represents the manipulators' desired alternative. However, an additional knowledge about the non-manipulators' votes (*top-2* queries), allows the manipulators to only increase their chances of winning, by raising their most preferred alternative to the second rank. For instance, when the manipulators strategy is *P=3* (plot c), with *top-1* queries, the winner is alternative 1; however, with *top-2, top-3 and top-4* queries, the winner is always the alternative 8 (true winner without manipulation), while alternative 1 is ranked second. Results with *M=750* manipulators, suggest that an optimal manipulation is derived when the manipulators have restricted knowledge i.e. with *top-1* queries of the non-manipulators' votes, the chances of the manipulators' preferred alternative increase by ranking him first, while the true winner is ranked second; which minimize the loss on social welfare.

In order to evaluate the effectiveness of the multi-round elicitation process, we propose to compare our results with those obtained in [10], where authors theoretically outlined a general framework for the design of a multi-round elicitation protocol, however they have empirically dealt only with one-round elicitation of *top-k* candidates. In this context, authors have estimated the expected minimax regret using the *MMR* solution in order to determine the winners given partial profile. As in our study, they have used Borda scoring rule to evaluate the outcome of the election. Their results with *Sushi data* suggest that with *top-5* queries one can usually find the true winner. While our goal is to minimize the information elicited as well as the number of rounds, results are more efficient when we consider probabilistic models to complete profiles. Our results suggest that *top-1* (resp. *top-2*) queries, are sufficient to determine the correct winner in one round (resp. 2 rounds), using the IC (resp. $SP\text{-}IC$) probabilistic model. More interesting results are obtained with single peaked data, where the winner of the election is always determined with any value of k, using the two probabilistic models.

More recent work proposed by authors in [15], addressing the problem of determining the outcome of an election with only partial information. In this context, authors propose a novel application with stronger connections to machine learning, where they used classification algorithms to predict the missing preferences of a ballot via latent patterns in the partial information provided. To this end, an imputation based approach to social choice was derived relying on the ability of machine learning algorithms to provide reasonable imputations of user's preferences in order to deal with incompleteness. As in this work, authors consider partial elicitation process in the form of top-k style preferences and Borda scoring rule to select the winner. However, our work differs from this recent context insofar as it uses conventional social choice techniques and recommends a particular elicitation strategy for voters i.e. the voter is asked to incrementally send his partial preferences in rounds, one preference each round; but instead works with the partial preferences to achieve the same goal.

While the proposed approach in [15] provides good results with incomplete data expressed by a low error rate using the SVM classifier, however, there are a number of reasons to suppose that our multi-round vote elicitation process should be preferred, at least in some situations. Chief among them is that in the case where the ballots have exceptionally high incompleteness (e.g. voters provide only top-1 preferences), classification becomes a difficult problem with the imputation method requiring additional refinement to perform well, thus, the correct winner is not predicted. In the similar situation, our results show that even with high incompleteness; not only the correct winner is determined but also the damage on social welfare is limited, which guarantee a high approximation of the entire ordering using a scoring rule.

To the best of our knowledge, this is the first work that investigates the relation between the amount of information needed to be known by the manipulators and the success of manipulation, empirically.

VI. Conclusion

This paper has addressed the problem of uncertainty in manipulators' knowledge, using an efficient multi-round elicitation process. Our contribution is of two-fold: First, we propose a *multi-round elicitation process* for choosing the ideal threshold k with the minimal number of rounds. Second, we propose a *new manipulation strategy with uncertain knowledge*, where the incomplete manipulators' knowledge is solved using the proposed multi-round elicitation process. Experiments with two different data sets prove the practical viability and advantages of the proposed multi-round elicitation process. Interesting results are derived when we consider a single peaked data, where the winner is determined with high quality using only *top-1* preferences from voters. Moreover, experiments performed with the new manipulation strategy indicate that our approach is quite tractable. Specifically, it allows the determination of an optimal manipulation using only a small fraction of sincere voters' preferences and with less impact on social welfare. Furthermore, results show that with restricted knowledge, the manipulators are always able to change the outcome to their favor, however when we increase the number of queries elicited, the manipulators can only raise the rank of their preferred candidate to the second place.

For future research, we are interested in examining the implication of single-peakedness data on our manipulation strategy i.e. study if it is sufficient to eliminate the possibility of profitable manipulation by using single peaked preferences. Moreover, extending our analysis to a richer class of probabilistic models, such as the Mallows model [16], is an important next step, by considering the appropriate form of queries that fits this model. For example, Irish electoral data is not adopted to such model since it is single peaked.

References

[1] A. Gibbard, *Manipulation of voting schemes: a general result*, Econometrica: journal of the Econometric Society, pp. 587-601, 1973.

[2] M. A. Satterthwaite, *Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions*, Journal of economic theory, vol. 10, no. 2, pp. 187-217, 1975.

[3] V. Conitzer, T. Sandholm, and J. Lang, *When are elections with few candidates hard to manipulate?*, Journal of the Association for Computing Machinery, vol. 54, no. 3, p. 14, 2007.

[4] N. Narodytska and T. Walsh, *The computational impact of partial votes on strategic voting*, in Proc. of the 21st European Conference on Artificial Intelligence, pp. 657-662, 2014.

[5] T. Lu, P. Tang, A. D. Procaccia, and C. Boutilier, *Bayesian vote manipulation: Optimal strategies and impact on welfare*, in Proc. of the 28th Conference on Uncertainty in Artificial Intelligence, pp. 543-553, 2012.

[6] Z. Rabinovich, S. Obraztsova, O. Lev, E. Markakis, and J. S. Rosenschein, *Analysis of equilibria in iterative voting schemes*, in Proc. of the 29th National Conference on Artificial Intelligence, pp. 1007-1013, 2015.

[7] V. Conitzer and T. Sandholm, *Vote elicitation: Complexity and strategy -proofness*, in Proc. of the 17th National Conference on Artificial Intelligence, pp. 392-397, 2002.

[8] V. Conitzer and T. W. Sandholm, *Communication complexity of common voting rules*, in Proc. of the 6th ACM conference on Electronic commerce, pp. 78-87, 2005.

[9] M. Kalech, S. Kraus, G. A. Kaminka, and C. V. Goldman, *Practical voting rules with partial information*, Journal of Autonomous Agents and Multi-Agent Systems, vol. 22, no. 1, pp. 151-182, 2011.

[10] T. Lu and C. Boutilier, *Vote elicitation with probabilistic preference models: Empirical estimation and cost tradeoffs*, in Proc. of the 2nd International Conference on Algorithmic Decision Theory, pp. 135-149, 2011.

[11] K. Konczak and J. Lang, *Voting procedures with incomplete preferences*, in Proc. of Multidisciplinary IJCAI Workshop on Advances in Preference Handling, vol. 20, pp. 124-129, 2005.

[12] T. Walsh, *Uncertainty in preference elicitation and aggregation*, in Proc. of the National Conference on Artificial Intelligence, vol. 22, no. 1, pp. 3-8, 2007.

[13] T. Kamishima, H. Kazawa, and S. Akaho, *Supervised ordering-an empirical survey*, in Proc. of the 5th IEEE International Conference on Data Mining, pp. 673-676, 2005.

[14] N. Mattei and T. Walsh, *Preflib: A library for preferences http://www. preflib. org*, in Proc. of the 3rd International Conference on Algorithmic Decision Theory, pp. 259-270, 2013.

[15] J. A. Doucette, K. Larson, and R. Cohen, *Approximate winner selection in social choice with partial preferences*, in Proc. of the 12th International Conference on Autonomous Agents and Multiagent Systems, 2014.

[16] C. L. Mallows, *Non-null ranking models. I*, Biometrika, pp. 114-130, 1957.