ARTICLE IN PRESS
YINCO:4208

Information and Computation ••• (••••) •••–•••

Contents lists available at ScienceDirect

Information and Computation

www.elsevier.com/locate/yinco

Prices matter for the parameterized complexity of shift bribery [☆]

Robert Bredereck [a], Jiehua Chen [a,*], Piotr Faliszewski [b], André Nichterlein [a], Rolf Niedermeier [a]

[a] Institut für Softwaretechnik und Theoretische Informatik, TU Berlin, Berlin, Germany
[b] AGH University of Science and Technology, Krakow, Poland

ARTICLE INFO

Article history:
Received 3 February 2015
Received in revised form 5 August 2016
Available online xxxx

Keywords:
Computational social choice
Voting and rank aggregation
Multivariate complexity analysis
W-hardness
Fixed-parameter tractability (FPT)
FPT approximation scheme

ABSTRACT

In the Shift Bribery problem, we are given an election, a preferred candidate p, and a budget. The goal is to ensure p's victory by shifting p higher in some voters' preference orders. However, each such shift request comes at a price and we must not exceed the given budget. We study the parameterized computational complexity of Shift Bribery for a number of parameters and several classes of price functions: For the number of affected voters, Shift Bribery is W[2]-hard for Borda, Maximin, and Copeland. For the number of positions by which p is shifted in total, the problem is fixed-parameter tractable for Borda and Maximin, and is W[1]-hard for Copeland. For the budget, the results depend on the price function class. Finally, Shift Bribery tends to be tractable when parameterized by the number of voters, but the results for the number of candidates are more enigmatic.

1. Introduction

Rank aggregation and election-winner determination are of key importance in various economical and political settings. For instance, there are product rankings based on comparing their prices, their features, and different tests (performed by various institutions such as foundations, journals, etc.); universities are judged based on multiple different criteria (e.g., the number of students per faculty member, availability of particular facilities, the number of Nobel prize winners employed etc.); sport competitions involve multiple rankings (for example, a Formula 1 season consists of about twenty races, each resulting in a ranking of the drivers); and political elections require members of the society to express preferences regarding the participating candidates. In each of these cases the provided rankings are aggregated into the final one, often of significant importance (for example, customers decide on their purchases based on product rankings, students pick the best-ranked universities, the Formula 1 world champion is the driver who comes out first in the aggregated ranking, and the most appreciated candidate becomes the country's president). A sophisticated way of dealing with rankings based on multiple different criteria is to compute a consensus ranking using preference-based rank aggregation methods.[1] In order to affect the outcome of the rank aggregation one has to influence the component rankings obtained from different sources (different

☆ An extended abstract appeared in the Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence 2014, AAAI Press, pages 1398–1404.
* Corresponding author.
E-mail addresses: robert.bredereck@tu-berlin.de (R. Bredereck), jiehua.chen@tu-berlin.de (J. Chen), faliszew@agh.edu.pl (P. Faliszewski), andre.nichterlein@tu-berlin.de (A. Nichterlein), rolf.niedermeier@tu-berlin.de (R. Niedermeier).

[1] For example, the German website idealo.de aggregates different product tests by first translating the test results into a unified rating system and then taking the "average" of all the ratings. Various university rankings are prepared in a similar way. It would be very interesting, however, to utilize

http://dx.doi.org/10.1016/j.ic.2016.08.003
0890-5401/

product tests, different judgment criteria, different races, different voters). Clearly, the cost of influencing may differ from source to source and, indeed, can sometimes be quite high. Nonetheless, the effect of improved position in the final ranking can be very beneficial.

In this work, we study the computational complexity of affecting the outcome of the rank aggregation by "bribing" specific agents to change their rankings. Moreover, replacing "bribery" with "product development," or "university expansion," or "training," or "political campaigning" we see that our work is relevant to all the settings mentioned above; the particular entities (companies offering their products, universities, drivers, politicians) can find out how much effort they need to invest in order to achieve a better position in the aggregated ranking (or maybe even become the winner). A natural and simple model in this context, using the formalisms of voting theory, is Shift Bribery as introduced by Elkind et al. [32]. We extend their studies in terms of charting the border of computational worst-case tractability, herein putting particular emphasis on the voter-specific "shifting prices" (how expensive it is to shift a candidate by x positions "up").

Informally (see Section 3 for a formal definition), Shift Bribery is the following decision problem:

Shift Bribery
Input: An election, that is, a set of candidates and a set of voters, each with a linear preference order over the candidate set, some preferred candidate p, and some budget.
Question: Can we make p win by bribing voters to shift p higher in their preference orders by "paying" no more than the given budget?

We assume that we have the knowledge of the voters' preference orders (for example, from preelection polls). Further, in our example settings often the full rankings are known. For example, a driver preparing for a new Formula 1 season has full knowledge of the results from the previous one. Our Shift Bribery problem models the situation where we approach each of the voters, one-on-one, and try to convince[2] him or her to rank p higher. Naturally, the effect (the number of positions by which p is shifted in each voter's preference order) depends on the voter's character and situation, and on the amount of effort we invest into convincing the voter. This "effort" could, for example, mean the amount of time spent, the cost of implementing a particular change, or, in the bribery view of the problem, the payment to the voter. Thus, the computational complexity of the problem depends on the voting rule used in the election, on various election parameters such as the numbers of candidates and voters, and on the type of price functions describing the efforts needed to shift p up by a given number of positions in the voters' preference orders. Our goal is to unravel the nature of these dependencies.

Related work The computational complexity of bribery in elections was first studied by Faliszewski et al. [35]. They considered the Bribery problem, where one asks if it is possible to ensure that a given candidate is an election winner by changing at most a given number of votes. Its priced variant, $Bribery, is the same except that each voter has a possibly different price for which we can change his or her vote. These problems were studied for various election rules, including Borda [22,35], Maximin [38], and Copeland [36] (see Section 2 for exact definitions of these rules). Recently, Gertler et al. [42] studied the bribery problem for linear ranking systems. Notably, the destructive variant of the Bribery problem (known under the name Margin of Victory), where the goal is to ensure that a despised candidate does not win (and which was studied, e.g., by Magrino et al. [49] and Xia [58]) has a surprisingly positive motivation—it can be used to detect fraud in elections.

The above problems, however, do not take into account that the price of bribing a voter may depend on what vote we wish the "bribed" voter to cast. For example, a voter might be perfectly happy to swap the two least preferred candidates but not the two most preferred ones. To model such situations, Elkind et al. [32] introduced the Swap Bribery problem. They assumed that each voter has a swap-bribery price function which gives the cost of swapping each two candidates (provided they are adjacent in the voter's preference order; one can perform a series of swaps to transform the voter's preference order in an arbitrary way). They found that Swap Bribery is both NP-hard and hard to approximate for most well-known voting rules (essentially, because the Possible Winner problem [4,6,47,59], which is NP-hard for almost all natural voting rules, is a special case of Swap Bribery with each swap costing either zero or infinity). Motivated by this, Dorn and Schlotter [28] considered the parameterized complexity of Swap Bribery for the case of k-Approval (where each voter gives a point to his or her top k candidates). In addition, Elkind et al. [32] also considered Shift Bribery, a variant of Swap Bribery where all the swaps have to involve the preferred candidate p. They have shown that Shift Bribery remains NP-hard for Borda, Maximin, and Copeland but that there is a 2-approximation algorithm for Borda and a polynomial-time algorithm for the k-Approval voting rule. Shift Bribery was further studied by Elkind and Faliszewski [31], who viewed it as a political campaign management problem (and whose view we adopt in this paper), and who gave a 2-approximation algorithm for all scoring rules (generalizing the result for Borda) and other approximation algorithms for Maximin and Copeland. Then, Schlotter et al. [57] have shown that Shift Bribery is polynomial-time solvable for the case of Bucklin and Fallback voting rules.

the rankings themselves, instead of the ratings, for the aggregation. Moreover, Formula 1 racing (and numerous similar competitions) use pure ranking information (e.g., Formula 1 uses a very slightly modified variant of the Borda election rule).

[2] What "to convince" means can vary a lot depending on the application scenario. On the evil side we have bribery, but it can also mean things such as product development, hiring more faculty members, training on a particular racing circuit, or explaining the details of one's political platform. Clearly, different ranking providers may appreciate different efforts, which is modeled by the individual price functions.