Control in Computational Social Choice

Jiehua Chen[1], **Joanna Kaczmarek**[2], **Paul Nüsken**[2], **Jörg Rothe**[2], **Ildikó Schlotter**[3] and **Tessa Seeger**[2]
[1]Institute of Logic and Computation, TU Vienna, A-1040 Vienna, Austria
[2]Institut für Informatik, MNF, Heinrich-Heine-Universität Düsseldorf, D-40225 Düsseldorf, Germany
[3]Institute of Economics, HUN-REN Centre for Economic and Regional Studies, 1097 Budapest, Hungary
jiehua.chen@ac.tuwien.ac.at, {joanna.kaczmarek, paul.nuesken, rothe, tessa.seeger}@hhu.de, schlotter.ildiko@krtk.hun-ren.hu

Abstract

We survey the notion of control in various areas of computational social choice (COMSOC) such as voting, fair allocation, cooperative game theory, matching under preferences, and group identification. In all these scenarios, control can be exerted, for instance, by adding or deleting agents with the goal of influencing the outcome. We conclude by briefly covering control in some other COMSOC areas including participatory budgeting, judgment aggregation, and opinion diffusion.

1 Introduction

Computational social choice was founded by three seminal papers of Bartholdi *et al.*, and the founding fathers of this area—at that time new, but now a key topic at all large AI conferences—focused on control of elections [Bartholdi III *et al.*, 1992], manipulation, and winner determination. We survey some central models and results about control in computational social choice since its beginnings. Electoral control means that a (usually external) agent (called the election chair) modifies the structure of an election by, e.g., adding or deleting voters or candidates with the goal of either making a favorite candidate win (in the constructive case) or preventing a despised candidate's victory (in the destructive case).

Along with manipulation and bribery, control attacks on single-winner elections were the main focus of attention in the early days of computational social choice. Since then the study of control has spread like a wildfire over various other subfields of computational social choice.

Our survey covers control not only in single-winner and multiwinner voting but also in fair allocation, cooperative game theory, matching under preferences, and group identification. In each of these fields, we describe the underlying models and scenarios and explain how control can be exerted in them, for instance, by adding or deleting agents with the goal of influencing the outcome. We give an overview of some of the main results on control in each of these fields and highlight a number of open questions and challenges for future research. Finally, we briefly cover control in participatory budgeting, judgment aggregation, and opinion diffusion.

2 Control in Voting

An election is given as a pair (C, V) with a set C of candidates and a list V of votes over C. We will assume that votes are linear orders (but note that there are also other ways of representing voter preferences, e.g., approval ballots). In order to determine the winner(s) of an election (respectively, its winning committee(s) of a given size), many single-winner (respectively, multiwinner) voting rules have been proposed.

A very important class of single-winner voting rules are the *positional scoring protocols* where candidates score points based on their positions in the votes, and whoever scores the most points wins. Only top-ranked candidates score a point in *plurality*, and in the *Borda* rule each of m candidates score $m - i$ points when ranked in a vote's i-th position. For instance, in the election shown in Figure 1, d with a score of 5 is the plurality winner (whereas a, b, and c score only 1, 3, and 3 points), and b and d with a score of 19 are the Borda winners (whereas a and c score only 17 points).

Other voting rules are based on pairwise comparisons of candidates—among those, especially important are the *Condorcet-consistent* rules, which elect the Condorcet winner whenever there is one. A *Condorcet winner* is a candidate who beats all other candidates by a majority of votes in pairwise comparison. Condorcet winners do not always exist, but if so, they are unique. For example, the *Schulze* rule is Condorcet-consistent. Being widely used in practice and celebrated for its many useful properties, it is based on the strength of paths in the *weighted majority graph* (WMG) of an election (C, V): There is a vertex for each candidate, and there is an edge from x to y exactly if the edge weight, defined as the difference $D_V(x, y)$ of how many voters prefer x to y minus how many prefer y to x, is positive (see the WMGs in Example 1). Define the *path strength* $\mathrm{str}(p)$ as the weight $D_V(c, d)$ of the weakest edge (c, d) on p. For each pair of distinct candidates $c, d \in C$, define the *strength of a strongest path between c and d* as $P_V(c, d) = \max\{\mathrm{str}(p) \mid p \text{ is a path from } c \text{ to } d\}$. Now, $c \in C$ is a *Schulze winner of* (C, V) if $P(c, d) \geq P(d, c)$ for each $d \in C \setminus \{c\}$.

Example 1. Anna (a), Belle (b), Chris (c), and David (d) run for president of the renowned Association for Advancing Anonymous Ideas (AAAI). The 12 current AAAI members eligible to vote cast the ballots shown in Figure 1 (left), where candidates are ordered from left (most preferred) to right

$v_1 : b\ a\ c\ d$	$v_2 : d\ c\ b\ a$
$v_3 : c\ a\ b\ d$	$v_4 : d\ b\ c\ a$
$v_5 : a\ d\ b\ c$	$v_6 : c\ b\ a\ d$
$v_7 : b\ c\ a\ d$	$v_8 : d\ a\ b\ c$
$v_9 : d\ b\ a\ c$	$v_{10} : c\ a\ d\ b$
$v_{11} : d\ c\ a\ b$	$v_{12} : b\ a\ d\ c$

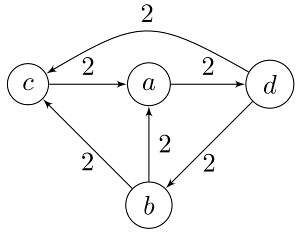

Figure 1: An election (left) and its WMG (right)

(least preferred). The corresponding WMG (right) shows that there is no Condorcet winner (as no vertex has only outgoing edges) and all candidates are Schulze winners. Evil Eve, though, is not happy about this. Being the election chair, she has the power to add new voters (whose preferences she knows). Wishing to make her favorite candidate d the *unique* Schulze winner, she adds the (boldfaced) voters $v_{13}, \ldots, v_{18}$, and we obtain the new election and WMG shown in Figure 2.

$v_1 : b\ a\ c\ d$	$v_2 : d\ c\ b\ a$
$v_3 : c\ a\ b\ d$	$v_4 : d\ b\ c\ a$
$v_5 : a\ d\ b\ c$	$v_6 : c\ b\ a\ d$
$v_7 : b\ c\ a\ d$	$v_8 : d\ a\ b\ c$
$v_9 : d\ b\ a\ c$	$v_{10} : c\ a\ d\ b$
$v_{11} : d\ c\ a\ b$	$v_{12} : b\ a\ d\ c$
$\boldsymbol{v_{13}} : b\ a\ c\ d$	$\boldsymbol{v_{14}} : d\ c\ b\ a$
$\boldsymbol{v_{15}} : b\ c\ a\ d$	$\boldsymbol{v_{16}} : d\ a\ b\ c$
$\boldsymbol{v_{17}} : d\ b\ a\ c$	$\boldsymbol{v_{18}} : c\ a\ d\ b$

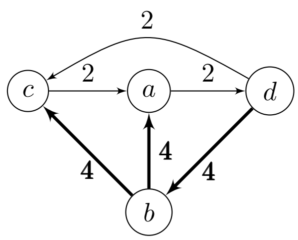

Figure 2: Evil Eve's control by adding voters

The WMG above shows that Eve has reached her goal: d alone wins. Fraudulent Frodo, however, is not amused. By making Eve's control attack public (thus causing her impeachment), he becomes the new election chair. Unlike Eve's constructive goal, his goal is purely destructive: He doesn't care who wins as long as d is *not* the only Schulze winner. Since he doesn't want to delete the voters just added, he exerts control by deleting b and obtains the election and WMG shown in Figure 3. Now, Frodo has reached his goal: Each of a, c, and d win, so d is not a unique Schulze winner. ⌟

$v_1 : a\ c\ d$	$v_2 : d\ c\ a$
$v_3 : c\ a\ d$	$v_4 : d\ c\ a$
$v_5 : a\ d\ c$	$v_6 : c\ a\ d$
$v_7 : c\ a\ d$	$v_8 : d\ a\ c$
$v_9 : d\ a\ c$	$v_{10} : c\ a\ d$
$v_{11} : d\ c\ a$	$v_{12} : a\ d\ c$
$v_{13} : a\ c\ d$	$v_{14} : d\ c\ a$
$v_{15} : c\ a\ d$	$v_{16} : d\ a\ c$
$v_{17} : d\ a\ c$	$v_{18} : c\ a\ d$

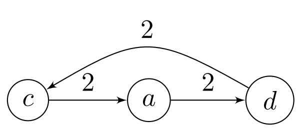

Figure 3: Fraudulent Frodo's control by deleting a candidate

The scenarios described in Example 1 give rise to defining the following problems. For *constructive control by adding voters* (CCAV), we are given a set C of candidates; two lists (V and U) of votes over C, where already registered voters cast the votes in V, and in U are those of as yet unregistered voters; a designated candidate $c \in C$; and a positive integer k. The question is whether there is a sublist $U' \subseteq U$, $|U'| \leq k$, such that c wins the election $(C, V \cup U')$. For *destructive control by deleting candidates* (DCDC), an election (C, V), a designated candidate $c \in C$, and a positive integer k are given, and we ask whether at most k candidates can be deleted from C such that c does not win the resulting election.

Two winner models are distinguished: The *unique-winner model* requires c to be the only winner in the constructive case and not winning alone in the destructive case, whereas the *nonunique-winner model* only requires c to be one (of possibly several) winner(s) in the constructive case and not winning at all in the destructive case. The problems of *constructive control by deleting voters* (CCDV) and *by deleting or adding candidates* (CCDC and CCAC; Bartholdi *et al.* [1992] originally defined the variant CCAUC with an *unlimited* number of candidates to be added) and of *destructive control by adding candidates* (DCAC) and *by adding or deleting voters* (DCAV and DCDV) are defined analogously.

A variety of *control scenarios by partition of voters or candidates* have also been studied. Control by partition of voters, for instance, models "gerrymandering"—ways of redistricting voting districts. We omit them here. We also omit defining and discussing further types of control, such as control by *replacing voters or candidates* [Loreggia *et al.*, 2015] and *multimode control attacks* [Faliszewski *et al.*, 2011], which combine various standard control types.

For some of the control scenarios defined above, the election chair's goal can never be reached. For example, constructive control by adding candidates is never possible for the chair in Condorcet voting: If the designated candidate c is not a Condorcet winner in a given election, c does not beat all other candidates in pairwise comparison, so c can never be made a Condorcet winner by adding more candidates. We then say Condorcet voting is *immune* (I) to this type of control. If a voting rule is not immune to some control type, we say it is *susceptible* (S) to it, and in that case we consider the computational complexity of the corresponding problem. If it can be solved in P, we say the rule is *vulnerable* (V) to this control type; and if it is NP-hard, we say the rule is *resistant* (R) to it. Table 1 gives an overview of the known complexity results for the four rules and all control scenarios defined above. Results marked by ⋆ are due to Bartholdi *et al.* [1992]; by † due to Hemaspaandra *et al.* [2007]; by § due to Russel [2007]; by $ due to Elkind *et al.* [2011]; by ♠ due to Parkes and Xia [2012]; by ‡ due to Menton and Singh [2013]; by ❄ due to Chen *et al.* [2017]; by £ due to Loreggia *et al.* [2015]; by ♡ due to Hemaspaandra and Schnoor [2016]; by ♣ due to Neveling and Rothe [2021]; by ◇ originally claimed by Menton and Singh [2013] whose proof was later shown to be wrong and corrected by Maushagen *et al.* [2024]; and by ¶ originally claimed by Menton and Singh [2012] but later stated as open [Menton and Singh, 2013] and re-established by Maushagen *et al.* [2024]. Question marks indicate open

	CCAV	DCAV	CCDV	DCDV	CCAC	DCAC	CCAUC	DCAUC	CCDC	DCDC
Plurality	V^{*}	$V^{\dagger}$	V^{*}	$V^{\dagger}$	$R^{\dagger}$	$R^{\dagger}$	R^{*}	$R^{\dagger}$	R^{*}	$R^{\dagger}$
Borda	$R^{\S}$	$V^{\S}$	$R^{\heartsuit}$	$V^{\S}$	$R^{\$}$	$V^{\pounds}$	$R^{\clubsuit}$	$V^{\clubsuit}$	$R^{\text{❄}}$	$V^{\pounds}$
Condorcet	R^{*}	$V^{\dagger}$	R^{*}	$V^{\dagger}$	$I^{\dagger}$	$V^{\dagger}$	I^{*}	$V^{\dagger}$	V^{*}	$I^{\dagger}$
Schulze	$R^{\spadesuit}$	$R^{\spadesuit}$	$R^{\spadesuit}$	$R^{\spadesuit}$	$R^{\spadesuit}$	?	$R^{\ddagger}$	?	$R^{\diamondsuit}$	$V^{\P}$

Table 1: Control complexity results for some voting rules

problems and background coloring is used to group each constructive with its destructive variant.

Observe that the destructive case never is harder than the constructive case (note that immunity means that the control problem is trivial and thus in P). Solving the open cases in Table 1 for Schulze voting seems to be a challenging task.

Challenge 1. Solve the open cases in Table 1 for Schulze voting: What is the complexity of DCAC and DCAUC?

The control complexity has also been explored for many other natural voting rules: for approval voting [Hemaspaandra *et al.*, 2007]; Copeland and Llull voting [Faliszewski *et al.*, 2009]; k-approval and k-veto [Lin, 2012]; range voting and normalized range voting [Menton, 2013]; Bucklin and fallback voting [Erdélyi *et al.*, 2015]; veto and maximin voting [Maushagen and Rothe, 2018; Maushagen and Rothe, 2020]; and ranked-pairs voting [Maushagen *et al.*, 2024].

Control in the context of multiwinner elections has been an area of growing interest in recent years. Multiwinner voting rules aim to select a fixed-size subset of candidates, referred to as a *winning committee*, that optimally reflects the preferences of the voters. The objectives of control in such settings can vary, including ensuring the winning committee maximizes the external agent's utility [Meir *et al.*, 2008], guaranteeing the inclusion of specific candidates (constructive control) [Yang, 2023], or preventing any given candidates from being included (destructive control).

The primary control actions investigated typically involve four common types: adding or deleting candidates [Meir *et al.*, 2008; Karh Bet *et al.*, 2024] or voters [Meir *et al.*, 2008]. These investigations span a range of voting systems, including approval-based and proportional representation methods, and typically focus on common control actions such as adding or removing either voters or candidates.

3 Control in Fair Allocation

Dividing resources (or, goods) among agents in a fair and efficient way is a practical problem that has been around since biblical times. The possible settings are widely varied, based on the type of resources, the fairness and efficiency criteria, and the possible additional constraints on the desired allocation. The resources to be allocated can be either *divisible* or *indivisible*, and they are usually nonhomogeneous, i.e., different agents may value a given (part of a) resource differently. In *cake cutting*, each agent has a utility function over a divisible resource called the *cake*, while in *fair division*, each agent has utilities over a set of indivisible *items*, expressed either as a cardinal utility function or a linear preference order.

Example 2. Suppose that Anna (a), Belle (b), Chris (c), and David (d) receive a gift bag from their aunt for Christmas, which contains a kite (K), a toy lion (L), a pair of mittens (M), a jar of nut spread (N), and an oboe (O). (Old aunt forgot how many nephews and nieces she actually has.)

To distribute the gifts in a fair way, their father asks the children to evaluate them, eliciting the values in Table 2.

	K	L	M	N	O
$a:$	6	2	1	10	1
$b:$	0	3	3	10	4
$c:$	5	3	0	10	2
$d:$	3	2	0	10	5

	π_{dad}	π_{mum}
$a:$	K	K
$b:$	M	L, M
$c:$	L	P
$d:$	O	O

Table 2: Children's values (left) and parents' allocations (right)

The children's father, anticipating a calamity, quickly confiscates the jar of nut spread. He assigns the gifts according to the allocation π_{dad} shown above. Pointing out that each child has received a *superproportional share*, i.e., a gift that is worth more than a fourth of the total value of all remaining gifts (i.e., more than $\frac{10}{4}$), he walks away with the nut spread.

Immediately, a skirmish breaks out, because Belle envies David for his oboe, and Chris envies Anna for her kite. The children's mother comes to the rescue brandishing a set of paints (P), valued to 5 by each child, and redistributes the gifts according to allocation π_{mum} above. Peace returns. ⌟

In Example 2, the control action performed by the father was *item removal*, with the aim of achieving an allocation that is (at least) *proportional*, i.e., that allocates to each agent p a bundle (i.e., a subset of the item set I) having value at least $u_p(I)/|A|$, where A is the set of agents and $u_p : I \to \mathbb{N}$ denotes p's valuation function which naturally extends to 2^I by assuming additive valuations. The second control action, performed by the mother, was *item addition*, to facilitate even an *envy-free* allocation, i.e., an allocation $\pi : A \to 2^I$ where $u_p(\pi(p)) \geq u_p(\pi(q))$ holds for each two agents p and q.

The study of control in fair allocation was initiated by Aziz *et al.* [2016]. Besides item removal and addition, they also define *agent removal* and *addition*, as well as *item/agent replacement* and *item/agent partitioning* for achieving fairness. Instead of defining these control actions formally, we focus on the control action considered most often (in fact, almost exclusively) in the literature: item removal. The popularity of this notion is probably due to the fact that donating goods is a natural and practically feasible option in most scenarios.

Caragiannis *et al.* [2019] consider cardinal and additive preferences, and propose an algorithm for finding an allocation that is *envy-free up to any item* (EFX), meaning that no agent envies any other agents for their bundle after the worst item is discarded from it, and is guaranteed to have at least half of the Nash welfare (i.e., of the geometric mean of the agents' utilities) achievable by any allocation. Chaudhury *et al.* [2021] give a method for finding an allocation that is EFX by donating a bundle of at most $|A|$ items such that no agents prefer the donated bundle to their own.

In a setting with ordinal preferences, Brams *et al.* [2014] have devised an efficient algorithm for two agents that produces an envy-free[1] partial allocation with the minimal number of unallocated (or, from a different perspective, donated) items. Aziz *et al.* [2016] show that, unless P = NP, a similar algorithm is not possible for three agents, since even determining whether a *complete* envy-free allocation exists is NP-hard. Under ordinal preferences, deciding the existence of a *proportional*[2] allocation is easy, but deciding whether removing at most a given number of items leaves an instance admitting a proportional allocation is NP-hard already for three agents [Dorn *et al.*, 2021]. Besides obtaining an FPT algorithm for three agents, parameterized by the number of item removals, Dorn *et al.* [2021] also consider a setting where some fixed allocation is given in advance, and the task is to make this allocation proportional by removing items.

Applying control to make an *a priori* fixed allocation fair has been studied by Boehmer *et al.* [2024] for the setting with additive cardinal utilities over indivisible items, and by Segal-Halevi [2022] for cake cutting with geometric constraints.

Finally, to achieve fairness, Hosseini *et al.* [2020] and Bliznets *et al.* [2024] have also considered hiding information, which may also be considered a form of control.

4 Control in Cooperative Game Theory

Coalitional games with a characteristic function are defined as pairs (N, v) with *player set* N and *characteristic function* $v : 2^N \to \mathbb{R}$, where $v(\emptyset) = 0$. Each subset of N is called a *coalition*. A coalitional game is *monotonic* if $v(C) \leq v(C')$ for all C, C' with $C \subseteq C' \subseteq N$, and if additionally $v(C) \in \{0,1\}$ for all $C \subseteq N$, it is said to be *simple*: If $v(C) = 1$, we call C a *winning coalition*, and if $v(C) = 0$, we call it a *losing coalition*. We call a player i *pivotal for a coalition* $C \subseteq N \setminus \{i\}$ if $v(C \cup \{i\}) - v(C) = 1$.

The analysis of simple games includes answering the question of how important a player is in forming winning coalitions, which is measured by *power indices* such as the *Shapley–Shubik index* [Shapley and Shubik, 1954] and the *probabilistic Penrose–Banzhaf index* [Dubey and Shapley, 1979]. The power indices count—each in a different way—the coalitions for which the player is pivotal in the considered game.

A *weighted voting game* (WVG) $\mathcal{G} = (w_1, \ldots, w_n; q)$ is a compactly representable simple coalitional game with player set $N = \{1, \ldots, n\}$, a *quota* $q \in \mathbb{N}$, and nonnegative integer weights, where w_i is the *weight of player* $i \in N$. Let $w_C = \sum_{i \in C} w_i$ for $C \subseteq N$. The characteristic function v of $\mathcal{G}$ is defined by $v(C) = 1$ if $w_C \geq q$, and $v(C) = 0$ otherwise. For a given player in a given WVG, it is #P-complete to compute the Shapley–Shubik index or the probabilistic Penrose–Banzhaf index, where #P is the class of functions that give the number of solutions of NP problems.

Inspired by the idea of control in voting (Section 2), Rey and Rothe [2018] introduced control by either *adding players to* or *deleting them from* WVGs. In both cases, the goal of the control action is to *increase*, *nondecrease*, *decrease*, *nonincrease*, or *maintain* a given player's power. To analyze the corresponding problems in terms of their computational complexity, they define them as follows: For control by deleting players, given a WVG, a distinguished player, and a specified limit, the question is whether it is possible to change or maintain—according to the chosen goal—the distinguished player's power index by deleting no more than the specified number of players. For control by adding players, new players are additionally given (by their weights), and the question is whether the specified goal can be achieved by adding no more than the specified number of new players.

These problems have been studied for the probabilistic Penrose–Banzhaf and the Shapley–Shubik indices first by Rey and Rothe [2018] and later on by Kaczmarek and Rothe [2024b; 2024a], who established $\mathrm{NP}^{\mathrm{PP}}$-completeness for all problems related to control by adding players, where $\mathrm{NP}^{\mathrm{PP}}$ is the class of problems solvable by an NP oracle machine accessing an oracle set from PP—"probabilistic polynomial time." For the problems of control by deleting players, the question of completeness is still open, i.e., there are still huge complexity gaps between the currently known upper bound of $\mathrm{NP}^{\mathrm{PP}}$ and the lower bounds that—depending on the specific goal of the problem—range from hardness for the complexity classes NP, coNP, DP, to Θ_2^p, where DP is the class of differences of NP sets and Θ_2^p is the class of problems solvable by a P oracle machine accessing an NP oracle set at most logarithmically often.

Kaczmarek and Rothe [2024b] have also introduced the model of weighted voting games where a game's quota is not fixed when the player set N changes but is defined to be $r \cdot w_N$ for some specified factor $r \in [0,1]$. This model takes into account the change of total weight when players are added or deleted. Kaczmarek and Rothe [2024b] again establish an upper bound of $\mathrm{NP}^{\mathrm{PP}}$ for all related control problems, and they show hardness for the classes NP, coNP, DP, and PP.

WVGs have also been studied with an additional restriction by an undirected simple graph whose vertices correspond to the players. In these games, a coalition C wins if and only if there exists $C' \subseteq C$ such that $w_{C'} \geq q$, for a given quota q, and C' induces a connected subgraph of the restricting graph. Despite the restriction, computing the two power indices remains #P-complete. With the restricting graph, additional control possibilities arise: Given a graph-restricted WVG, a distinguished player, and a specified limit, the question is whether adding or removing up to the specified number of edges can change or maintain the distinguished player's power. Kaczmarek *et al.* [2025] show that most of these problems are at least NP-hard or coNP-hard (and some even PP-hard), and upper-bounded by $\mathrm{NP}^{\mathrm{PP}}$.

Challenge 2. For the goal of either changing or maintaining a player's power, show completeness (for $\mathrm{NP}^{\mathrm{PP}}$?) of (1) control by deleting players from a WVG; (2) control by adding

[1] An allocation π is *envy-free under ordinal preferences* if for each two agents p and q, there is an injection f from $\pi(q)$ to $\pi(p)$ such that for each item $x \in \pi(q)$, agent p prefers $f(x)$ to x.

[2] An allocation $\pi : A \to 2^I$ is *proportional under ordinal preferences* if for any $i \leq |I|$, all agents get at least $i/|A|$ items among the first i items of their preference lists.

players to or deleting them from a WVG with changing quota; and (3) control by adding edges to or deleting them from the graph underlying a graph-restricted WVG. Further, analyze these control types for other power indices.

5 Control in Matching Under Preferences

Most work on control related to matching under preferences concentrates on the classical *stable marriage* (SM) problem and its generalization, the *college admission* (CA) problem [Gale and Shapley, 1962]. In an instance of SM, we are given a set of agents on a two-sided market, traditionally called *men* and *women*, and a preference list for each agent, which is a strict linear order over a subset of agents from the opposite side of the market. The task is to find a *matching* between men and women that is *stable*, i.e., contains no man–woman pair such that both of them prefer each other to their partner in the matching (called a *blocking pair*).

Example 3. Suppose that Anna (a), Belle (b), Chris (c), and David (d) are attending a dance class, and need to form opposite-sex couples. Their preferences are as follows: $a : c\ d$, $b : c\ d$, $c : a\ b$, $d : b\ a$.

The only stable matching in this instance is $\{(a, c), (b, d)\}$. However, Belle's friend, evil Eve, is among the teachers of the class, and her goal is to match Belle with her top-choice partner, Chris. She considers three options to achieve her goal: (1) declaring that Chris is too short for Anna and hence cannot dance with her, (2) stepping on Anna's toes with her high heels, thereby sending her off to ER, or (3) inviting her attractive friend, Frodo (f), whom Anna prefers to Chris. She decides on option (3), and obtains: $a : f\ c\ d$, $b : c\ d\ f$, $c : a\ b$, $d : b\ a$, $f : b\ a$.

Yet, Eve is not entirely satisfied, as now there are two stable matchings, $M_1 = \{(a, f), (b, c)\}$ and $M_2 = \{(a, c), (b, f)\}$. Hence, she declares that Frodo is too tall for Belle, thus ensuring that M_1 becomes the only stable matching. ⌟

Example 3 highlights some of the settings examined by Boehmer *et al.* [2021] who initiated the study of control problems in relation to stable matchings. Boehmer *et al.* [2021] define five manipulative actions and three different goals, thus obtaining 15 different computational problems. Among these actions are the control actions showcased in Example 3:

- ADDAG: adding agents (e.g., inviting Frodo),
- DELAG: deleting agents (e.g., removing Anna), and
- DELACC: deleting acceptability (e.g., declaring constraints on who can dance with whom);

others involve changing the preference lists of the agents and thus fall into the category of manipulation or bribery.

Example 3 depicts also some of the possible goals that an external controller may want to ensure. These may either focus on a distinguished agent or pair of agents, or aim to obtain a stable matching with some desirable property, e.g., a *perfect* matching where *every* agent is matched. In the generalization of SM where the underlying graph is not necessarily bipartite, called the *stable roommates* (SR) model, a stable matching may not exist, so ensuring the existence of a stable matching becomes a meaningful aim. Below, we summarize known results for problems where the controller's goal is to

	CSM-$\mathcal{A}$-MA	CSR-$\mathcal{A}$-MA	CSM-$\mathcal{A}$-MP	CSR-$\mathcal{A}$-MP	CSM-$\mathcal{A}$-MS	CSR-$\mathcal{A}$-MS	CSM-$\mathcal{A}$-USM	CSR-$\mathcal{A}$-USM	CSR-$\mathcal{A}$-∃SM	CSM-$\mathcal{A}$-∃PSM	CSR-$\mathcal{A}$-∃PSM
ADDAG	R$^{\triangledown}$	R$^{\triangledown}$	R*	R*	V*	R$^{\triangledown}$	R*	R*	R$^{\triangledown}$	R$^{\triangledown}$	R$^{\triangledown}$
DELAG	V$^{\underline{*}}$	V$^{\triangledown}$	V*	V$^{\triangledown}$	R*	R*	R*	R*	V†	V†	V†
DELACC	R$^{\heartsuit}$	R$^{\heartsuit}$	R*	R*	V*	V$^{\underline{*}}$	V*	?	R$^{\diamond}$	R$^{\spadesuit}$	R$^{\spadesuit}$

Table 3: Complexity results for the CSM-$\mathcal{A}$-$\mathcal{G}$ and CSR-$\mathcal{A}$-$\mathcal{G}$ problems. Background coloring is used to group each problem CSR-$\mathcal{A}$-$\mathcal{G}$ and its bipartite variant (where the CSM-$\mathcal{A}$-∃SM column is omitted, as a stable matching always exists in the bipartite case).

- match a given agent in a stable matching (MA),
- make a given pair contained in a stable matching (MP),
- make a given matching stable (MS),
- make a given matching the only stable matching (USM),
- ensure that a stable matching exists (∃SM), or
- ensure that a perfect and stable matching exists (∃PSM).

For each control action $\mathcal{A} \in \{$ADDAG, DELAG, DELACC$\}$ and each goal $\mathcal{G} \in \{$MA, MP, MS, USM, ∃SM, ∃PSM$\}$ discussed above, in CONTROL-IN-STABLE-MARRIAGE-$\mathcal{A}$-$\mathcal{G}$ (or, CSM-$\mathcal{A}$-$\mathcal{G}$), we ask whether at most a given number of control actions achieves the given goal in a given SM instance; we call the nonbipartite variants of these problems CONTROL-IN-STABLE-ROOMMATES-$\mathcal{A}$-$\mathcal{G}$ (or, CSR-$\mathcal{A}$-$\mathcal{G}$).[3]

In Table 3, results by Boehmer *et al.* [2021] and easy consequences thereof are marked with $*$ and with $\underline{*}$, respectively. Results by Chen and Schlotter [2025] are marked with $\triangledown$, by Tan [1990; 1991] with †, by Mnich and Schlotter [2020] with ♡, by Abraham *et al.* [2005] with ◇, and by Biró *et al.* [2010] with ♠. Note that some of the problems CSR-DELACC-$\mathcal{G}$ have been studied in the context of *almost stable* matchings, i.e., matchings with only a few blocking pairs: For example, an instance of SR admits a matching with at most k blocking pairs if and only if we can ensure the existence of a stable matching by k deletions of acceptability. Similarly, the minimum number of blocking pairs in any matching that matches a given agent or contains a given pair is exactly the minimum number of DELACC actions that ensures achieving the corresponding goal $\mathcal{G} \in \{$MA, MP$\}$.

Further intractability and parameterized results are provided by Chen *et al.* [2018] for CSR-DELACC-MP, by Gupta and Jain [2025] for weighted and destructive variants of many of the problems in Table 3, and by Bérczi *et al.* [2024] regarding agent deletion problems with additional constraints. Kamiyama [2025] looks at the problem where preferences can contain ties, and when aiming for a super-stable matching by deleting as few agents as possible.

Challenge 3. What is the complexity of the problem CSR-DELACC-USM, which is open in Table 3?

A prominent line of research has also emerged in con-

[3] For the precise definition of control goals MS and USM for control by changing the agent set, see [Boehmer *et al.*, 2021].

nection to the CA problem, the many-to-one variant of SM where the two sides of the market represent *students* and *colleges*, and each college comes with a *capacity*. The task is then to find a matching of students to colleges that *respects capacities and is stable*, i.e., there exists no student–college pair (s, c) such that s prefers to be matched to c, while either c is unsaturated, or there is a student matched to c to whom c prefers s. The control actions focused on by most researchers in this setting are *capacity increase*, *capacity decrease*, or *capacity modification* (when both increasing and decreasing capacities is allowed). The controller's goal is most often to ensure the existence of a stable matching that fulfills some desirable property such as being perfect or Pareto-optimal, by using as few control actions as possible.

Bobbio *et al.* [2022] have studied the problem of minimizing capacity increase (or decrease) for obtaining a stable matching that minimizes the average college rank to which students are matched, and proved both problems to be NP-complete and hard to approximate; a further study developed mixed integer programs for these problems. Chen and Csáji [2023] initiated the study of determining the existence of a stable matching that is simultaneously perfect or Pareto-optimal through capacity increases, while bounding the sum or the maximum of these modifications. Among these four problems, only one turned out to be polynomial-time solvable (when we bound the maximum capacity increases and aim for a perfect and stable matching), while the other three are NP-hard. They further investigate the parameterized complexity.

Afacan *et al.* [2024] also consider obtaining a Pareto-optimal and stable matching, but they work with a model that includes a lower quota for each college and has an upper bound only on the sum of college capacities. Gokhale *et al.* [2024] give a polynomial-time algorithm for stabilizing a given matching (MS) via capacity increase or decrease, and they prove that the problem of ensuring that some student–college pair is contained in a stable matching (MP) is NP-hard for these two control actions. Nguyen and Vohra [2018] consider a setting where students can form couples and submit joint preference lists. They show that there is a capacity modification yielding a stable matching where each college's capacity is modified by at most two, and the total capacity modification is at most four.

6 Control in Group Identification

Broadly speaking, group identification deals with finding a *socially qualified group* among individuals. To this end, each individual either qualifies or disqualifies all other individuals (and themselves) for inclusion in the group. More formally, given a set A of agents (or, individuals) and a profile $\varphi : A \times A \to \{0, 1\}$, a *group identification rule* F determines a socially qualified group $F(\varphi, A) \subseteq A$. We say *an individual* $a \in A$ *(socially) qualifies an individual* $b \in A$ if $\varphi(a, b) = 1$, and a *(socially) disqualifies* b if $\varphi(a, b) = 0$. In the setting of control, three group identification rules have been studied: the consent rule [Samet and Schmeidler, 2003] and two procedural rules [Dimitrov *et al.*, 2007], namely the liberal-start-respecting-rule and the consensus-start-respecting-rule.

In the *consent rule*, $F^{(s,t)}$, a social qualification is determined by the agents' individual assessments and two thresholds, s and t. If individuals qualify themselves, they are qualified if and only if at least $s - 1$ other individuals qualify them. Vice versa, if individuals do not qualify themselves, they are disqualified if and only if at least $t - 1$ other individuals also disqualify them. The second type of group identification rules are the procedural rules. These rules recursively add individuals who are qualified by the current members to the group until no member qualifies any individual outside the group (i.e., no new member is added during a recursion call). The two rules differ in how they select the initial group. The *liberal-start-respecting-rule*, F^{lib}, starts with the set of individuals who qualify themselves, while in the *consensus-start-respecting-rule*, F^{con}, the initial set is given by individuals who are qualified by everyone (including themselves).

Example 4. After noticing the rigged election from Example 1, AAAI proposes a new format for finding their leadership. They ask each of the four candidates—Anna (a), Belle (b), Chris (c), and David (d)—whom they deem qualified to lead the association in the coming years. Of course, everyone qualifies themselves. In addition, Anna qualifies everyone else; Belle qualifies Chris and David; Chris qualifies no one else; and David qualifies Chris. These qualifications are depicted in the left graph of Figure 4 that contains a directed edge from v to v' if and only if $\varphi(v, v') = 1$ for $v, v' \in A$.

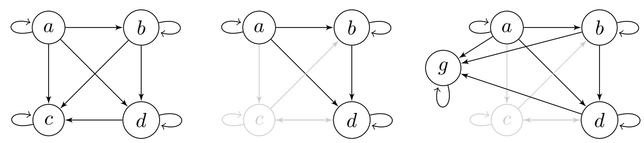

Figure 4: Two control actions for Example 4

Using F^{con}, the association determines Chris as the sole member of the qualified group. Evil Eve, again unhappy with the result, removes Chris from the competition, thus making David the sole qualified individual (middle graph of Figure 4). This, again, enrages fraudulent Frodo who takes action and adds Grace (g) to the pool of individuals (right graph of Figure 4). Grace qualifies only herself and is deemed qualified by everyone. As a result, the new socially qualified group consists solely of Grace, and Frodo is happy. ⌟

Evil Eve and fraudulent Frodo both changed the qualified group by modifying the set of participating individuals. The study of control complexity in group identification was initiated by Yang and Dimitrov [2018], who first studied *constructive group control by adding individuals* (CGCAI), *constructive group control by deleting individuals* (CGCDI), and *constructive group control by partitioning sets of individuals* (CGCPI). In the setting of group identification, the controller's goal is to help a subset $A^+ \subseteq A$ of individuals to become socially qualified (constructive control) [Yang and Dimitrov, 2018], or to prevent a subset $A^- \subseteq A$ of individuals from being included in the socially qualified group (destructive control) [Erdélyi *et al.*, 2020]. Table 4 gives an overview of results for control in group identification for F^{lib} and F^{con}. Results marked by ◇ are due to Yang and Dimitrov [2018], and by ♠ are due to Erdélyi *et al.* [2020].

	CGCAI	DGCAI	CGCDI	DGCDI	CGCPI	DGCPI
F^{lib}	$\mathrm{R}^{\diamond}$	$\mathrm{I}^{\spadesuit}$	$\mathrm{I}^{\diamond}$	$\mathrm{V}^{\spadesuit}$	$\mathrm{I}^{\diamond}$	$\mathrm{V}^{\spadesuit}$
F^{con}	$\mathrm{R}^{\diamond}$	$\mathrm{R}^{\spadesuit}$	$\mathrm{V}^{\diamond}$	$\mathrm{V}^{\spadesuit}$	?	?

Table 4: Control complexity results for group identification

Recently, more attention has been given to restricting to the domain of *consecutive qualifications* by Yang and Dimitrov [2023] who study control in this domain for the consent rule and both procedural rules, and they do so through the lens of parameterized complexity.

Challenge 4. Solve the open cases in Table 4 for F^{con}: What is the complexity of CGCPI and DGCPI?

7 Outlook: Control in Other COMSOC Areas

As demonstrated in the previous sections, the study of control has spread exhaustively through the world of COMSOC. A keen ready may, however, have noticed that some areas of COMSOC remain unmentioned, to which we will turn now.

We start with *participatory budgeting* (PB) where agents collectively decide on project funding under a budget constraint. Notably, approval-based PB is related to multiwinner voting, i.e., if each project is assigned a cost of 1 and the budget is precisely k, the PB process can be modeled as a size-k committee election. To this day, the study of control in participatory budgeting has been relatively sparse. Baumeister *et al.* [2021] introduce a very general control framework, called *manipulative interference*, for both the constructive and destructive settings, in which they study two control strategies: changing the budget and changing an item's cost. Most recently, *candidate control in PB* has been studied by Faliszewski *et al.* [2025].

Next, the task in *judgment aggregation* (JA) is to aggregate individual judgments over logically interconnected propositions by using a JA rule. Focusing on the important class of so-called "uniform premise-based quota rules," Baumeister *et al.* [2020] introduced control by adding, deleting, replacing, and bundling judges and studied the related problems in terms of their complexity. For the Kemeny JA rule (which, roughly speaking, minimizes the cumulative Hamming distance to the individual judgment sets), the complexity of control by adding or deleting issues has been studied by de Haan [2017].

Last but not least, we turn to *opinion diffusion* and *social networks*. Both areas focus on a network of connected agents, each of whom holds an opinion, which can be expressed in many forms—to name just a few: Opinions can be represented in binary, as a ranking over some alternatives, a bit vector, or a state (e.g., activated). Using a certain update rule, the state of the network (i.e., the agents' opinions) can then be updated in an iterative process. Turning to control in opinion diffusion and social networks, we mention two models, which most closely relate to electoral control: influencing the network itself (see, e.g., [Bredereck and Elkind, 2017]) and influencing a social network with the goal of changing an election outcome (see, e.g., [Castiglioni *et al.*, 2021]).

8 Discussion and Conclusions

We have surveyed the central notion of *control* in a variety of COMSOC subareas, starting from voting where Bartholdi *et al.* [1992] were the first to introduce it and to define the standard types of electoral control—including adding or deleting either candidates or voters. Their early work was among the key seminal papers that launched COMSOC as a stand-alone area which due to its many applications in real-world scenarios is now indispensible for any of the large AI conferences.

Since then, an extraordinarily large body of work has explored control in many settings. Novel standard types of control have been introduced and studied, such as *destructive* control due to Hemaspaandra *et al.* [2007] where the goal is to prevent some candidate from winning an election. However, in addition to voting scenarios, control has also been introduced to many other subareas of COMSOC. We have focused on presenting the models of control and (mainly complexity) results about the related problems in fair allocation, cooperative game theory, matching under preferences, and group identification. Further, we have briefly presented some initial control results in participatory budgeting, judgment aggregation, and opinion diffusion—the latter being closely related to social networks where control scenarios (even though often dubbed differently) have been studied for a long time already.

The tasks of a controller in subareas other than voting can be distinct from simply adding or deleting agents; for example, items can be added or deleted in fair allocation or edges of an underlying communication graph can be added or deleted in graph-restricted WVGs. This may call for new techniques to tackle these problems and classify them in terms of their computational or parameterized complexity.

It is also noteworthy that electoral control—which in fact is what one commonly means when speaking of a *"rigged election"*—is much worse in moral or ethical terms than strategic voting (which, in technical terms, is dubbed *"manipulation"*) or bribery in elections. Indeed, voters are justifiably allowed to cast strategic ballots that best serve their own goals and bribery can synonymously be seen as *"campaign management"*—what does it cost to make voters cast certain ballots and what can be achieved within a given campaign budget?

Control, manipulation, and bribery—three key topics at the heart of COMSOC—have been studied in much depth for voting. Much less is known in the other COMSOC subareas we have covered, and we encourage the reader to contribute to exploring especially control in them. In addition, our last and perhaps most demanding challenge goes one step further:

Challenge 5. If your favorite COMSOC area is not among those covered in this survey, introduce to it a suitable model of control and study the complexity of the related problems.

Notwithstanding the obvious importance of worst-case complexity results such as NP-hardness for control, it will also be crucial to advance our knowledge of hardness for typical instances occurring in practice or even on average. Even if that is beyond our control so far, we will never give up hope.

Acknowledgments

This work was supported in part by the Vienna Science and Technology Fund (WWTF) [10.47379/VRG18012], by Deutsche Forschungsgemeinschaft under DFG research grant RO-1202/21-2 (project 438204498), and by the Hungarian Academy of Sciences under its Momentum Programme (LP2021-2) and its János Bolyai Research Scholarship.

References

[Abraham *et al.*, 2005] D. Abraham, P. Biró, and D. Manlove. "Almost stable" matchings in the roommates problem. In *WAOA'05*, pages 1–14, 2005.

[Afacan *et al.*, 2024] M. Afacan, U. Dur, and M. Van der Linden. Capacity design in school choice. *GEB*, 146:277–291, 2024.

[Aziz *et al.*, 2016] H. Aziz, I. Schlotter, and T. Walsh. Control of fair division. In *IJCAI'16*, pages 67–73, 2016.

[Bartholdi III *et al.*, 1992] J. Bartholdi III, C. Tovey, and M. Trick. How hard is it to control an election? *Math. and Comp. Modelling*, 16(8/9):27–40, 1992.

[Baumeister *et al.*, 2020] D. Baumeister, G. Erdélyi, O. Erdélyi, J. Rothe, and A. Selker. Complexity of control in judgment aggregation for uniform premise-based quota rules. *JCSS*, 112:13–33, 2020.

[Baumeister *et al.*, 2021] D. Baumeister, L. Boes, and J. Hillebrand. Complexity of manipulative interference in participatory budgeting. In *ADT'21*, pages 424–439, 2021.

[Bérczi *et al.*, 2024] K. Bérczi, G. Csáji, and T. Király. Manipulating the outcome of stable marriage and roommates problems. *GEB*, 147:407–428, 2024.

[Biró *et al.*, 2010] P. Biró, D. Manlove, and S. Mittal. Size versus stability in the marriage problem. *TCS*, 411(16–18):1828–1841, 2010.

[Bliznets *et al.*, 2024] I. Bliznets, A. Bukov, and D. Sagunov. Fair division with minimal withheld information in social networks. *TCS*, 991:114446, 2024.

[Bobbio *et al.*, 2022] F. Bobbio, M. Carvalho, A. Lodi, and A. Torrico. Capacity variation in the many-to-one stable matching. Technical Report arXiv:2205.01302, 2022.

[Boehmer *et al.*, 2021] N. Boehmer, R. Bredereck, K. Heeger, and R. Niedermeier. Bribery and control in stable marriage. *JAIR*, 71:993–1048, 2021.

[Boehmer *et al.*, 2024] N. Boehmer, R. Bredereck, K. Heeger, D. Knop, and J. Luo. Multivariate algorithmics for eliminating envy by donating goods. *JAAMAS*, 38(2):43, 2024.

[Brams *et al.*, 2014] S. Brams, M. Kilgour, and C. Klamler. Two-person fair division of indivisible items: An efficient, envy-free algorithm. *Notices AMS*, 61(2):130–141, 2014.

[Bredereck and Elkind, 2017] R. Bredereck and E. Elkind. Manipulating opinion diffusion in social networks. In *IJCAI'17*, pages 894–900, 2017.

[Caragiannis *et al.*, 2019] I. Caragiannis, N. Gravin, and X. Huang. Envy-freeness up to any item with high Nash welfare: The virtue of donating items. In *ACM-EC'19*, pages 527–545, 2019.

[Castiglioni *et al.*, 2021] M. Castiglioni, D. Ferraioli, N. Gatti, and G. Landriani. Election manipulation on social networks: Seeding, edge removal, edge addition. *JAIR*, 71:1049–1090, 2021.

[Chaudhury *et al.*, 2021] B. Chaudhury, T. Kavitha, K. Mehlhorn, and A. Sgouritsa. A little charity guarantees almost envy-freeness. *SICOMP*, 50(4):1336–1358, 2021.

[Chen and Csáji, 2023] J. Chen and G. Csáji. Optimal capacity modification for many-to-one matching problems. In *AAMAS'23*, pages 2880–2882, 2023.

[Chen and Schlotter, 2025] J. Chen and I. Schlotter. Control in stable marriage and stable roommates: Complexity and algorithms. Technical Report arXiv:2502.01215, 2025.

[Chen *et al.*, 2017] J. Chen, P. Faliszewski, R. Niedermeier, and N. Talmon. Elections with few voters: Candidate control can be easy. *JAIR*, 60:937–1002, 2017.

[Chen *et al.*, 2018] J. Chen, D. Hermelin, M. Sorge, and H. Yedidsion. How hard is it to satisfy (almost) all roommates? In *ICALP'18*, pages 35:1–35:15, 2018.

[de Haan, 2017] R. de Haan. Complexity results for manipulation, bribery and control of the Kemeny judgment aggregation procedure. In *AAMAS'17*, pages 1151–1159, 2017.

[Dimitrov *et al.*, 2007] D. Dimitrov, S. Sung, and Y. Xu. Procedural group identification. *MSS*, 54(2):137–146, 2007.

[Dorn *et al.*, 2021] B. Dorn, R. de Haan, and I. Schlotter. Obtaining a proportional allocation by deleting items. *Algorithmica*, 83(5):1559–1603, 2021.

[Dubey and Shapley, 1979] P. Dubey and L. Shapley. Mathematical properties of the Banzhaf power index. *MOR*, 4(2):99–131, 1979.

[Elkind *et al.*, 2011] E. Elkind, P. Faliszewski, and A. Slinko. Cloning in elections: Finding the possible winners. *JAIR*, 42:529–573, 2011.

[Erdélyi *et al.*, 2015] G. Erdélyi, M. Fellows, J. Rothe, and L. Schend. Control complexity in Bucklin and fallback voting: A theoretical analysis. *JCSS*, 81:632–660, 2015.

[Erdélyi *et al.*, 2020] G. Erdélyi, C. Reger, and Y. Yang. The complexity of bribery and control in group identification. *JAAMAS*, 34(1):8, 2020.

[Faliszewski *et al.*, 2009] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Llull and Copeland voting computationally resist bribery and constructive control. *JAIR*, 35:275–341, 2009.

[Faliszewski *et al.*, 2011] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. Multimode control attacks on elections. *JAIR*, 40:305–351, 2011.

[Faliszewski *et al.*, 2025] P. Faliszewski, L. Janeczko, D. Knop, J. Pokorný, Š. Schierreich, M. Słuszniak, and K. Sornat. Participatory budgeting project strength via candidate control. In *AAMAS'25*, 2025.

[Gale and Shapley, 1962] D. Gale and L. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.

[Gokhale *et al.*, 2024] S. Gokhale, S. Singla, S. Narang, and R. Vaish. Capacity modification in the stable matching problem. In *AAMAS'24*, pages 697–705, 2024.

[Gupta and Jain, 2025] S. Gupta and P. Jain. Manipulation with(out) money in matching market. In *ADT'24*, pages 273–287, 2025.

[Hemaspaandra and Schnoor, 2016] E. Hemaspaandra and H. Schnoor. Dichotomy for pure scoring rules under manipulative electoral actions. In *ECAI'16*, pages 1071–1079, 2016.

[Hemaspaandra *et al.*, 2007] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Anyone but him: The complexity of precluding an alternative. *AIJ*, 171(5–6):255–285, 2007.

[Hosseini *et al.*, 2020] H. Hosseini, S. Sikdar, R. Vaish, H. Wang, and L. Xia. Fair division through information withholding. In *AAAI'20*, pages 2014–2021, 2020.

[Kaczmarek and Rothe, 2024a] J. Kaczmarek and J. Rothe. Control by adding players to change or maintain the Shapley–Shubik or the Penrose–Banzhaf power index in weighted voting games is complete for $\mathrm{NP}^{\mathrm{PP}}$. In *ECAI'24*, pages 3525–3532, 2024.

[Kaczmarek and Rothe, 2024b] J. Kaczmarek and J. Rothe. Controlling weighted voting games by deleting or adding players with or without changing the quota. *AMAI*, 92(3):631–669, 2024.

[Kaczmarek *et al.*, 2025] J. Kaczmarek, J. Rothe, and N. Talmon. Control by adding or deleting edges in graph-restricted weighted voting games. *JAIR*, 82:1077–1145, 2025.

[Kamiyama, 2025] N. Kamiyama. Modifying an instance of the super-stable matching problem. *IPL*, 189:106549, 2025.

[Karh Bet *et al.*, 2024] G. Karh Bet, J. Rothe, and R. Zorn. Complexity of candidate control for single nontransferable vote and bloc voting. In *ADT'24*, pages 3–17, 2024.

[Lin, 2012] A. Lin. *Solving Hard Problems in Election Systems*. PhD thesis, RIT, Rochester, NY, USA, 2012.

[Loreggia *et al.*, 2015] A. Loreggia, N. Narodytska, F. Rossi, B. Venable, and T. Walsh. Controlling elections by replacing candidates or votes (extended abstract). In *AAMAS'15*, pages 1737–1738, 2015.

[Maushagen and Rothe, 2018] C. Maushagen and J. Rothe. Complexity of control by partitioning veto elections and of control by adding candidates to plurality elections. *AMAI*, 82(4):219–244, 2018.

[Maushagen and Rothe, 2020] C. Maushagen and J. Rothe. The last voting rule is home: Complexity of control by partition of candidates or voters in maximin elections. In *ECAI'20*, pages 163–170, 2020.

[Maushagen *et al.*, 2024] C. Maushagen, D. Niclaus, P. Nüsken, J. Rothe, and T. Seeger. Toward completing the picture of control in Schulze and ranked pairs elections. In *IJCAI'24*, pages 2940–2948, 2024.

[Meir *et al.*, 2008] R. Meir, A. Procaccia, J. Rosenschein, and A. Zohar. Complexity of strategic behavior in multi-winner elections. *JAIR*, 33:149–178, 2008.

[Menton and Singh, 2012] C. Menton and P. Singh. Manipulation and control complexity of Schulze voting. Technical Report arXiv:1206.2111v1, 2012.

[Menton and Singh, 2013] C. Menton and P. Singh. Control complexity of Schulze voting. In *IJCAI'13*, pages 286–292, 2013.

[Menton, 2013] C. Menton. Normalized range voting broadly resists control. *Theory of Computing Systems*, 53(4):507–531, 2013.

[Mnich and Schlotter, 2020] M. Mnich and I. Schlotter. Stable marriage with covering constraints: A complete computational trichotomy. *Algorithmica*, 82:1136–1188, 2020.

[Neveling and Rothe, 2021] M. Neveling and J. Rothe. Control complexity in Borda elections: Solving all open cases of offline control and some cases of online control. *AIJ*, 298:103508, 2021.

[Nguyen and Vohra, 2018] T. Nguyen and R. Vohra. Near-feasible stable matchings with couples. *American Economic Review*, 108(11):3154–3169, 2018.

[Parkes and Xia, 2012] D. Parkes and L. Xia. A complexity-of-strategic-behavior comparison between Schulze's rule and ranked pairs. In *AAAI'12*, pages 1429–1435, 2012.

[Rey and Rothe, 2018] A. Rey and J. Rothe. Structural control in weighted voting games. *BEJTE*, 18(2):1–15, 2018.

[Russel, 2007] N. Russel. Complexity of control of Borda count elections. Master's thesis, RIT, Rochester, NY, USA, 2007.

[Samet and Schmeidler, 2003] D. Samet and D. Schmeidler. Between liberalism and democracy. *JET*, 110(2):213–233, 2003.

[Segal-Halevi, 2022] E. Segal-Halevi. Redividing the cake. *JAAMAS*, 36(1):14:1–14:36, 2022.

[Shapley and Shubik, 1954] L. Shapley and M. Shubik. A method of evaluating the distribution of power in a committee system. *Amer. Pol. Sci. Rev.*, 48(3):787–792, 1954.

[Tan, 1990] J. Tan. A maximum stable matching for the roommates problem. *BIT Num. Math.*, 30:631–640, 1990.

[Tan, 1991] J. Tan. Stable matchings and stable partitions. *IJCM*, 39:11–20, 1991.

[Yang and Dimitrov, 2018] Y. Yang and D. Dimitrov. How hard is it to control a group? *JAAMAS*, 32(5):672–692, 2018.

[Yang and Dimitrov, 2023] Y. Yang and D. Dimitrov. Group control for consent rules with consecutive qualifications. *MSS*, 121:1–7, 2023.

[Yang, 2023] Y. Yang. Complexity of manipulating and controlling approval-based multiwinner voting. Technical Report arXiv:2302.11291v2, 2023.