

Preventing Manipulation by Restricting Information

It was shown in the last chapter that computational complexity does not seem to be a very strong barrier against manipulation. Consequently, we need to look for new barriers. In this chapter we examine some preliminary ideas to prevent manipulation for the cases where there is one manipulator, by restricting the manipulator's information about the other voters' votes. We recall that in all previously studied manipulation problems, it is normally assumed that the manipulator has full information about the votes of the non-manipulators. The argument often given is that if it is NP-hard with full information, then it only can be at least as computationally difficult with partial information. However, when there is only one manipulator, computing a manipulation is in \mathbf{P} for most common voting rules, including all positional scoring rules, Copeland, maximin, and voting trees (see Table 3.1). The only known exceptions are STV (Bartholdi and Orlin, 1991), ranked pairs (Xia et al., 2009), and Nanson's and Baldwin's rules (Narodytska et al., 2011). It is not clear whether it is computationally easy for a single manipulator to find a manipulation when she only has partial information for other rules.

In this chapter, we first model how one manipulator computes a manipulation based on partial information about the other votes. For example, the manipulator may know that some voters prefer one alternative to another, but might not be able to know all pairwise comparisons for all voters. We suppose the knowledge of the manipulator is described by an *information set* E . This is some subset of possible profiles of the non-manipulators which is known to contain the true profile. Given an information set and a pair of votes U and V , if for every profile in E , the manipulator is not worse off voting U than voting V , and there exists a profile in E such that the manipulator is strictly better off voting U , then we say that U *dominates* V . If there exists a vote U that dominates the true preferences of the manipulator then the manipulator has an incentive to vote untruthfully. We call this a *dominating manipulation*. If there is no such vote, then a cautious manipulator might have little incentive to vote strategically.

We are interested in whether a voting rule r is *immune* to dominating manipulations, meaning that a voter's true preferences are never dominated by another vote. If r is not immune to dominating manipulations, we are interested in whether r is *resistant*, meaning that computing whether a voter's true preferences are dominated by another vote U is NP-hard, or *vulnerable*, meaning that this problem is in P. These properties depend on both the voting rule and the form of the partial information. Interestingly, it is not hard to see that most voting rules are immune to manipulation when the partial information is just the current winner. For instance, with any majority consistent rule (for example, plurality), a risk averse manipulator will still want to vote for her most preferred alternative. This means that the chairman does not need to keep the current winner secret to prevent such manipulations. On the other hand, if the chairman lets slip more information, many rules stop being immune. With most scoring rules, if the manipulator knows the current scores, then the rule is no longer immune to such manipulation. For instance, when her most

preferred alternative is too far behind to win, the manipulator might vote instead for a less preferred candidate who can win.

In this chapter, we focus on the case where the partial information is represented by a profile P_{po} of partial orders, and the information set E consists of all linear orders that extend P_{po} . The dominating manipulation problem is related to the *possible/necessary winner* problems, which I have briefly talked about in Section 1.6 and Section 2.3. We recall that in possible/necessary winner problems, we are given an alternative c and a profile of partial orders P_{po} that represents the partial information of the voters' preferences. We are asked whether c is the winner for *some* extension of P_{po} (that is, c is a *possible winner*), or whether c is the winner for *every* extension of P_{po} (that is, c is a *necessary winner*). We note that in the possible/necessary winner problems, there is no manipulator and P_{po} represents the chair's partial information about the votes. In dominating manipulation problems, P_{po} represents the partial information of the manipulator about the non-manipulators.

In the following sections, we start with the special case where the manipulator has complete information. In this setting the dominating manipulation problem reduces to the standard manipulation problem, and many common voting rules are vulnerable to dominating manipulation (from known results). When the manipulator has no information, we show that a wide range of common voting rules are immune to dominating manipulation. When the manipulator's partial information is represented by partial orders, our results are summarized in Table 6.1.

Our results are encouraging. For most voting rules r we study in this paper (except plurality and veto), hiding even a little information makes r resistant to dominating manipulation. If we hide all information, then r is immune to dominating manipulation. Therefore, limiting the information available to the manipulator

¹ All hardness results hold even when the number of undetermined pairs in each partial order is no more than a constant.

Table 6.1: Computational complexity of the dominating manipulation problems with partial orders, for common voting rules.

	DOMINATING MANIPULATION
STV	Resistant (Proposition 6.2.2)
Ranked pairs	Resistant (Proposition 6.2.2)
Borda	Resistant (Theorem 6.3.1)
Copeland	Resistant (Corollary 6.3.7)
Voting trees	Resistant (Corollary 6.3.7)
Maximin	Resistant (Theorem 6.3.8)
Plurality	Vulnerable (Algorithm 5)
Veto	Vulnerable (Similar to plurality)

appears to be a promising way to prevent strategic voting.

6.1 Framework for Manipulation with Partial Information

We now introduce the framework of this paper. In this chapter, we suppose there are $n - 1 \geq 1$ non-manipulators and one manipulator to make notion easier. The information the manipulator has about the votes of the non-manipulators is represented by an *information set* E . The manipulator knows for sure that the profile of the non-manipulators is in E . However, the manipulator does not know exactly which profile in E it is. Usually E is represented in a compact way. Let \mathcal{I} denote the set of all possible information sets in which the manipulator may find herself.

Example 6.1.1. *Suppose the voting rule is r .*

- *If the manipulator has no information, then the only information set is $E = \mathcal{F}_{n-1}$. Therefore $\mathcal{I} = \{\mathcal{F}_{n-1}\}$. Here we recall that \mathcal{F}_{n-1} is the set of all $(n-1)$ -profiles*
- *If the manipulator has complete information, then $\mathcal{I} = \{\{P\} : P \in \mathcal{F}_{n-1}\}$.*
- *If the manipulator knows the current winner (before the manipulator votes), then the set of all information sets the manipulator might know is $\mathcal{I} = \{E_1, E_2, \dots, E_m\}$, where for any $i \leq m$, $E_i = \{P \in \mathcal{F}_{n-1} : r(P) = c_i\}$.*

Let V_M denote the true preferences of the manipulator. Given a voting rule r

and an information set E , we say that a vote U *dominates* another vote V , if for every profile $P \in E$, we have $r(P \cup \{U\}) \succeq_{V_M} r(P \cup \{V\})$, and there exists $P' \in E$ such that $r(P' \cup \{U\}) \succ_{V_M} r(P' \cup \{V\})$. In other words, when the manipulator only knows the voting rule r and the fact that the profile of the non-manipulators is in E (and no other information), voting U is a strategy that dominates voting V . We define the following two decision problems.

Definition 6.1.2. *Given a voting rule r , an information set E , the true preferences V_M of the manipulator, and two votes V and U , we are asked the following two questions.*

- *Does U dominate V ? This is the DOMINATION problem.*
- *Does there exist a vote V' that dominates V_M ? This is the DOMINATING MANIPULATION problem.*

We stress that usually E is represented in a compact way, otherwise the input size would already be exponentially large, which would trivialize the computational problems. Given a set \mathcal{I} of information sets, we say a voting rule r is *immune* to dominating manipulation, if for every $E \in \mathcal{I}$ and every V_M that represents the manipulator's preferences, V_M is not dominated; r is *resistant* to dominating manipulation, if DOMINATING MANIPULATION is NP-hard (which means that r is not immune to dominating manipulation, assuming $P \neq NP$); and r is *vulnerable* to dominating manipulation, if r is not immune to dominating manipulation, and DOMINATING MANIPULATION is in P.

6.2 Manipulation with Complete/No Information

In this section we focus on the following two special cases: (1) the manipulator has complete information, and (2) the manipulator has no information. It is not hard to see that when the manipulator has complete information, DOMINATING MANIPULA-

TION coincides with the standard manipulation problem. Therefore, our framework of dominating manipulation is an extension of the traditional manipulation problem, and we immediately obtain the following proposition from the Gibbard-Satterthwaite theorem (Gibbard, 1973; Satterthwaite, 1975).

Proposition 6.2.1. *When $m \geq 3$ and the manipulator has full information, a voting rule satisfies non-imposition and is immune to dominating manipulation if and only if it is a dictatorship.*

The following proposition directly follows from the computational complexity of the manipulation problems for some common voting rules (Bartholdi et al., 1989a; Bartholdi and Orlin, 1991; Conitzer et al., 2007; Zuckerman et al., 2009; Xia et al., 2009).

Proposition 6.2.2. *When the manipulator has complete information, STV, ranked pairs, Nanson’s and Baldwin’s rules are resistant to DOMINATING MANIPULATION; all positional scoring rules, Copeland, voting trees, and maximin are vulnerable to dominating manipulation.*

Next, we investigate the case where the manipulator has no information. We obtain the following positive results.

Theorem 6.2.3. *When the manipulator has no information, any Condorcet consistent voting rule r is immune to dominating manipulation.*

Proof. For the sake of contradiction, let U dominates V_M . Because $U \neq V_M$, there exist two alternatives a and b such that $a \succ_{V_M} b$ and $b \succ_U a$. We prove the theorem in the following two cases.

Case 1: $n - 1$ is even. For any j such that $1 \leq j \leq (n - 1)/2$, we let $V_{2j-1} = [a \succ b \succ (\mathcal{C} \setminus \{a, b\})]$, where the alternatives in $\mathcal{C} \setminus \{a, b\}$ are ranked according to

the ascending order of their subscripts; let $V_{2j} = [b \succ a \succ \text{Rev}(\mathcal{C} \setminus \{a, b\})]$. Here $\text{Rev}(\mathcal{C} \setminus \{a, b\})$ is the reverse of $\mathcal{C} \setminus \{a, b\}$. Let $P = (V_1, \dots, V_{n-1})$. It follows that a is the Condorcet winner for $P \cup \{V_M\}$ and b is the Condorcet winner for $P \cup \{U\}$. Because $a \succ_{V_M} b$, V_M is not dominated by U , which contradicts the assumption.

Case 2: $n - 1$ is odd. For any j such that $1 \leq j \leq (n - 2)/2$, we let $V_{2j-1} = [a \succ b \succ (\mathcal{C} \setminus \{a, b\})]$ and $V_{2j} = [b \succ a \succ \text{Rev}(\mathcal{C} \setminus \{a, b\})]$. Suppose $a = c_{i_1}$ and $b = c_{i_2}$. Let $V_{n-1} = \begin{cases} V_1 & \text{if } i_1 > i_2 \\ V_2 & \text{if } i_1 < i_2 \end{cases}$. Let $P = (V_1, \dots, V_{n-1})$. It follows that a is the Condorcet winner for $P \cup \{V_M\}$ and b is the Condorcet winner for $P \cup \{U\}$, which contradicts the assumption. \square

Theorem 6.2.4. *When the manipulator has no information, Borda is immune to dominating manipulation.*

Proof. For the sake of contradiction, let U dominates V_M . Because $U \neq V_M$, there exists $i^* \leq m$ such that $\text{Alt}(V_M, i^*) \neq \text{Alt}(U, i^*)$ and for every $i < i^*$, $\text{Alt}(V_M, i) = \text{Alt}(U, i)$. That is, i^* is the first position from the top where the alternatives in V_M and U are different. Let $c_{i_1} = \text{Alt}(V_M, i^*)$ and $c_{i_2} = \text{Alt}(U, i^*)$. We prove the theorem in the following three cases.

Case 1: $n - 1$ is even. For any $i < i' \leq m$, let $V_M^{[i, i']}$ denote the sub-linear-order of V_M that starts at the i th position of V_M and ends at the i' th position of V_M . For any j such that $1 \leq j \leq (n - 1)/2$, we let $V_{2j-1} = [V_M^{[i^*, m]} \succ \text{Rev}(V_M^{[1, i^*-1]})]$ and $V_{2j} = [\text{Rev}(V_M^{[i^*, m]}) \succ \text{Rev}(V_M^{[1, i^*-1]})]$. Let $P = (V_1, \dots, V_{n-1})$. It follows that $\text{Borda}(P \cup \{V_M\}) = c_{i_1}$ and $\text{Borda}(P \cup \{U\}) = c_{i_2}$. We note that $c_{i_1} \succ_{V_M} c_{i_2}$, which contradicts the assumption.

Case 2: $n - 1$ is odd and c_1 is ranked within top i^* positions in V_M . For any j such that $1 \leq j \leq (n - 2)/2$, we let $V_{2j-1} = [c_1 \succ c_2 \succ \dots \succ c_m]$ and $V_{2j} = [c_m \succ c_{m-1} \succ \dots \succ c_1]$. Let $V_{n-1} = \text{Rev}(V)$ and $P = (V_1, \dots, V_{n-1})$. It

follows that $\text{Borda}(P \cup \{V\}) = c_1$ and $\text{Borda}(P \cup \{U\}) \neq c_1$, which contradicts the assumption.

Case 3: $n - 1$ is odd and c_1 is not ranked within top i^* positions in V_M . Let V_1, \dots, V_{n-2} be defined the same as in Case 2. Let $V' = [V_M^{[i^*, m]} \succ \text{Rev}(V_M^{[1, i^* - 1]})]$. Let $U' = [U^{[i^*, m]} \succ \text{Rev}(U^{[1, i^* - 1]})]$. It follows that $\text{Borda}(V', V_M) = c_{i_1}$. Let $a = \text{Borda}(V', U)$. If $a \neq c_{i_1}$, then $c_{i_1} \succ_{V_M} a$. This is because the alternative ranked within top $i^* - 1$ positions in V_M gets exactly the average score in $\{V', U\}$, which means that in order for any of them to win, the scores of all alternative in $\{V', U\}$ must be the same. However, due to the tie-breaking mechanism, the winner is c_1 , which contradicts the assumption that c_1 is not ranked within top i^* positions in V_M . Let $P' = (V_1, \dots, V_{n-2}, V')$, we have that $\text{Borda}(P' \cup \{V_M\}) = c_{i_1} \succ_{V_M} a = \text{Borda}(P' \cup \{U\})$, which contradicts the assumption. If $a = c_{i_1}$, then $\text{Borda}(U', V_M) = \text{Borda}(V', U) = a = c_{i_1}$. Let $P^* = (V_1, \dots, V_{n-2}, U')$. We have $\text{Borda}(P^* \cup \{V_M\}) = c_{i_1} \succ_{V_M} c_{i_2} = \text{Borda}(P^* \cup \{U\})$, which is a contradiction.

Therefore, the theorem is proved. \square

Theorem 6.2.5. *When the manipulator has no information and $n \geq 6(m - 2) + 1$, any positional scoring rule is immune to dominating manipulation.*

Proof. For the sake of contradiction, let U dominates V_M . Let $c = \arg \max_{c^*} \{\vec{s}_m(V_M, c^*) : \vec{s}_m(V_M, c^*) > \vec{s}_m(U, c^*)\}$. It follows that there exists an alternative c' such that $\vec{s}_m(V_M, c') < \vec{s}_m(V_M, c)$ and $\vec{s}_m(U, c') = \vec{s}_m(V_M, c)$. It follows that $s_m(V_M, c) > \vec{s}_m(V_M, c')$ and $\vec{s}_m(U, c') = \vec{s}_m(V_M, c) > \vec{s}_m(U, c)$.

We prove the theorem for the case where $c = c_1$ and $c' = c_2$. The other cases can be proved similarly. Let M_{m-2} denote the cyclic permutation such that $c_3 \rightarrow c_4 \rightarrow \dots \rightarrow c_m \rightarrow c_3$. For any $k \in \mathbb{N}$ and any $c \in \mathcal{C} \setminus \{c_1, c_2\}$, we let $M_{m-2}^0(c) = c$ and $M_{m-2}^k(c) = M(M_{m-2}^{k-1}(c))$. Let $W = [c_1 \succ c_2 \succ c_3 \succ \dots \succ c_m]$ and $W' = [c_2 \succ c_1 \succ c_3 \succ \dots \succ c_m]$. Let P_1 denote the $6(m - 2)$ -profile that is composed of three copies

of $\{W, W', M_{m-2}(W), M_{m-2}(W'), \dots, M_{m-2}^{m-3}(W), M_{m-2}^{m-3}(W')\}$.

If $n - 1$ is even, then let P be composed of P_1 plus $(n - 1)/2 - 3(m - 2)$ copies of $\{W, W'\}$. If $n - 1$ is odd, then let W^* denote the a vote obtained from V_M by exchanging the positions of c and c' and let P be composed of $P_1 \cup \{W^*\}$ plus $\lfloor (n - 1)/2 \rfloor - 3(m - 2)$ copies of $\{W, W'\}$. Because $\vec{s}_m(1) > \vec{s}_m(m)$, we have that $r(P \cup \{V_M\}) = c_1$ and $r(P \cup \{U\}) = c_2$. We note that $c_1 \succ_{V_M} c_2$. Therefore, we obtain a contradiction, which means that V_M is not dominated. \square

These results demonstrate that the information that the manipulator has about the votes of the non-manipulators plays an important role in determining strategic behavior. When the manipulator has complete information, many common voting rules are vulnerable to dominating manipulation, but if the manipulator has no information, then many common voting rules become immune to dominating manipulation.

6.3 Manipulation with Partial Orders

In this section, we study the case where the manipulator has partial information about the votes of the non-manipulators. We suppose the information is represented by a profile P_{po} composed of partial orders. That is, the information set is $E = \{P \in \mathcal{F}_n : P \text{ extends } P_{po}\}$. We note that the two cases discussed in the previous section (complete information and no information) are special cases of manipulation with partial orders. Consequently, by Proposition 6.2.1, when the manipulator's information is represented by partial orders and $m \geq 3$, no voting rule that satisfies non-imposition and non-dictatorship is immune to dominating manipulation. It also follows from Theorem 6.2.4 that STV and ranked pairs are resistant to dominating manipulation. The next theorem states that even when the manipulator only misses a tiny portion of the information, Borda becomes resistant to dominating

manipulation.

Theorem 6.3.1. DOMINATION and DOMINATING MANIPULATION with partial orders are NP-hard for Borda, even when the number of unknown pairs in each vote is no more than 4.

Proof. We only prove that DOMINATION is NP-hard, via a reduction from EXACT COVER BY 3-SETS (x3C). The proof for DOMINATING MANIPULATION is similar to the proof of the NP-hardness of the possible winner problems under positional scoring rules in Xia and Conitzer (2011a).

In an x3C instance, we are given two sets $\mathcal{V} = \{v_1, \dots, v_q\}$, $\mathcal{S} = \{S_1, \dots, S_t\}$, where for any $j \leq t$, $S_j \subseteq \mathcal{V}$ and $|S_j| = 3$. We are asked whether there exists a subset \mathcal{S}' of \mathcal{S} such that each element in \mathcal{V} is in exactly one of the 3-sets in \mathcal{S}' . We construct a DOMINATION instance as follows.

Alternatives: $\mathcal{C} = \{c, w, d\} \cup \mathcal{V}$, where d is an auxiliary alternative. Therefore, $m = |\mathcal{C}| = q + 3$. Ties are broken in the following order: $c > w > \mathcal{V} > d$.

Manipulator's preferences and possible manipulation: $V_M = [w > c > d > \mathcal{V}]$. We are asked whether $V = V_M$ is dominated by $U = [w > d > c > \mathcal{V}]$.

The profile of partial orders: Let $P_{po} = P_1 \cup P_2$, defined as follows.

First part (P_1) of the profile: For each $j \leq t$, We define a partial order O_j as follows.

$$O_j = [w > S_j > d > \text{Others}] \setminus [\{w\} \times (S_j \cup \{d\})]$$

That is, O_j is a partial order that agrees with $w > S_j > d > \text{Others}$, except that the pairwise relations between (w, S_j) and (w, d) are not determined (and these are the only 4 unknown relations). Let $P_1 = \{O_1, \dots, O_t\}$.

Second part (P_2) of the profile: We first give the properties that we need P_2 to satisfy, then show how to construct P_2 in polynomial time. All votes in P_2 are linear orders that are used to adjust the score differences between alternatives. Let

$P'_1 = \{w \succ S_i \succ d \succ \text{Others} : i \leq t\}$. That is, P'_1 ($|P'_1| = t$) is an extension of P_1 (in fact, P'_1 is the set of linear orders that we started with to obtain P_1 , before removing some of the pairwise relations). Let $\vec{s}_m = (m-1, \dots, 0)$. P_2 is a set of linear orders such that the following holds for $Q = P'_1 \cup P_2 \cup \{V\}$:

- (1) For any $i \leq q$, $\vec{s}_m(Q, c) - \vec{s}_m(Q, v_i) = 1$, $\vec{s}_m(Q, w) - \vec{s}_m(Q, c) = 4q/3$.
- (2) For any $i \leq q$, the scores of v_i and w, c are higher than the score of d in any extension of $P_1 \cup P_2 \cup \{V\}$ and in any extension of $P_1 \cup P_2 \cup \{U\}$.
- (3) The size of P_2 is polynomial in $t + q$.

We now show how to construct P_2 in polynomial time. For any alternative $a \neq d$, we define the following two votes: $W_a = \{[a \succ d \succ \text{Others}], [\text{Rev}(\text{Others}) \succ a \succ d]\}$, where $\text{Rev}(\text{Others})$ is the reversed order of the alternatives in $\mathcal{C} \setminus \{a, d\}$. We note that for any alternative $a' \in \mathcal{C} \setminus \{a, d\}$, $\vec{s}_m(W, a) - \vec{s}_m(W, a') = 1$ and $\vec{s}_m(W, a') - \vec{s}_m(W, d) = 1$. Let $Q_1 = P'_1 \cup \{V\}$. P_2 is composed of the following parts:

- (1) $tm - \vec{s}_m(Q_1, c)$ copies of W_c .
- (2) $tm + 4q/3 - \vec{s}_m(Q_1, w)$ copies of W_w .
- (2) For each $i \leq q$, there are $tm - 1 - \vec{s}_m(Q_1, v_i)$ copies of W_{v_i} .

We next prove that V is dominated by U if and only if c is the winner in at least one extension of $P_{po} \cup \{V\}$. We note that for any $v \in \mathcal{V} \cup \{w\}$, the score of v in V is the same as the score of v in U . The score of c in U is lower than the score of c in V . Therefore, for any extension P^* of P_{po} , if $r(P^* \cup \{V\}) \in (\{w\} \cup \mathcal{V})$, then $r(P^* \cup \{V\}) = r(P^* \cup \{U\})$ (because d cannot win). Hence, for any extension P^* of P_{po} , voting U can result in a different outcome than voting V only if $r(P^* \cup V) = c$. If there exists an extension P^* of P_{po} such that $r(P^* \cup \{V\}) = c$, then we claim that the manipulator is strictly better off voting U than voting V . Let P_1^* denote the extension of P_1 in P^* . Then, because the total score of w is no more than the total score of c , w is ranked lower than d at least $\frac{q}{3}$ times in P_1^* . Meanwhile, for each $i \leq q$, v_i is not ranked higher than w more than one time in P_1^* , because otherwise

the total score of v_i will be strictly higher than the total score of c . That is, the votes in P_1^* where $d \succ w$ make up a solution to the x3C instance. Therefore, the only possibility for c to win is for the scores of c, w , and all alternatives in \mathcal{V} to be the same (so that c wins according to the tie-breaking mechanism). Now, we have $w = r(P^* \cup \{U\})$. Because $w \succ_{V_M} c$, the manipulator is better off voting U . It follows that V is dominated by U if and only if there exists an extension of $P_{po} \cup \{V\}$ where c is the winner.

The above reasoning also shows that V is dominated by U if and only if the x3C instance has a solution. Therefore, DOMINATION is NP-hard. For the DOMINATING MANIPULATION problem, we add to P_{po} a profile P_E defined as follows. For each $e \in \mathcal{V} \cup \{w\}$ and each $i \leq l - 1$, we obtain a vote $V_{e,i}$ from V_M by exchanging the alternative ranked in the $(i+1)$ th position and e , and then exchanging the alternative ranked in the i th position and d ; let $O_{e,i}$ denote the partial order obtained from $V_{e,i}$ by removing $d \succ e$. Let M denote the following cyclic permutation $c \rightarrow w \rightarrow d \rightarrow \mathcal{V} \rightarrow A \rightarrow c$. Let P_E denote q copies of $\{O_{e,i}, M(V_{e,i}), M(V_{e,i})^2, \dots, M^{l-1}(V_{e,i}) : e \in \mathcal{V} \cup \{w\}, i \leq l - 1\}$. We note that in an extension P_E^* of P_E where the extension of $O_{e,i}$ is $V_{e,i}$, then the scores of the alternatives in P_E^* are the same.

For any vote W where there exists $v \in \mathcal{V}$ such that the score difference between w and v is different from the score difference between w and v in V_M , there must exist $v' \in \mathcal{V}$ such that the score difference between w and v' in W is strictly smaller than their score difference in V_M . Then, it is not hard to find an extension of P_{po} such that if the manipulator votes V_M , then w wins, and if the manipulator votes W , then v' wins, which means that V_M is not dominated by W . Therefore, if V_M is dominated by another W , then the score differences between w and the alternatives in \mathcal{V} are the same across V_M and W . Following the same reasoning as for the DOMINATION problem, we conclude that DOMINATING MANIPULATION is NP-hard. \square

Theorem 6.3.1 can be generalized to a class of scoring rules similar to the class of rules in Theorem 1 in Xia and Conitzer (2011a), which does not include plurality or veto. In fact, as we will show later, plurality and veto are vulnerable to dominating manipulation.

We now investigate the relationship to the possible winner problem in more depth. In a possible winner problem (r, P_{po}, c) , we are given a voting rule r , a profile P_{po} composed of n partial orders, and an alternative c . We are asked whether there exists an extension P of P_{po} such that $c = r(P)$. Intuitively, both DOMINATION and DOMINATING MANIPULATION seem to be harder than the possible winner problem under the same rule. Next, we present two theorems, which show that for any WMG-based rule, DOMINATION and DOMINATING MANIPULATION are harder than two special possible winner problems, respectively.

We first define a notion that will be used in defining the two special possible winner problems. For any instance of the possible winner problem (r, P_{po}, c) , we define its *WMG partition* $\mathbb{R} = \{R_{c'} : c' \in \mathcal{C}\}$ as follows. For any $c' \in \mathcal{C}$, let $R_{c'} = \{\text{WMG}(P) : P \text{ extends } P_{po} \text{ and } r(P) = c'\}$. That is, $R_{c'}$ is composed of all WMGs of the extensions of P_{po} , where the winner is c' . It is possible that for some $c' \in \mathcal{C}$, $R_{c'}$ is empty. For any subset $\mathcal{C}' \subseteq \mathcal{C} \setminus \{c\}$, we let $G_{\mathcal{C}'}$ denote the weighted majority graph where for each $c' \in \mathcal{C}'$, there is an edge $c' \rightarrow c$ with weight 2, and these are the only edges in $G_{\mathcal{C}'}$. We are ready to define the two special possible winner problems for WMG-based voting rules.

Definition 6.3.2. *Let d^* be an alternative and let \mathcal{C}' be a nonempty subset of $\mathcal{C} \setminus \{c, d^*\}$. For any WMG-based voting rule r , we let $PW_1(d^*, \mathcal{C}')$ denote the set of possible winner problems (r, P_{po}, c) satisfying the following conditions:*

1. *For any $G \in R_c$, $r(G + G_{\mathcal{C}'}) = d^*$.*
2. *For any $c' \neq c$ and any $G \in R_{c'}$, $r(G + G_{\mathcal{C}'}) = r(G)$.*

3. For any $c' \in \mathcal{C}'$, $R_{c'} = \emptyset$.

We recall that R_c and $R_{c'}$ are elements in the WMG partition of the possible winner problem.

Definition 6.3.3. Let d^* be an alternative and let \mathcal{C}' be a nonempty subset of $\mathcal{C} \setminus \{c, d^*\}$. For any WMG-based voting rule r , we let $PW_2(d^*, \mathcal{C}')$ denote the problem instances (r, P_{po}, c) of $PW_1(d^*, \mathcal{C}')$, where for any $c' \in \mathcal{C} \setminus \{c, d^*\}$, $R_{c'} = \emptyset$.

Theorem 6.3.4. Let r be a WMG-based voting rule. There is a polynomial time reduction from $PW_1(d^*, \mathcal{C}')$ to DOMINATION with partial orders, both under r .

Proof. Let (r, P_{po}, c) be a $PW_1(d^*, \mathcal{C}')$ instance. We construct the following DOMINATION instance. Let the profile of partial orders be $Q_{po} = P_{po} \cup \{\text{Rev}(d^* > c > \mathcal{C}' > \text{Others})\}$, $V = V_M = [d^* > c > \mathcal{C}' > \text{Others}]$, and $U = [d^* > \mathcal{C}' > c > \text{Others}]$. Let P be an extension of P_{po} . It follows that $\text{WMG}(P \cup \{\text{Rev}(d^* > c > \mathcal{C}' > \text{Others}), V\}) = \text{WMG}(P)$, and $\text{WMG}(P \cup \{\text{Rev}(d^* > c > \mathcal{C}' > \text{Others}), U\}) = \text{WMG}(P) + G_{c'}$. Therefore, the manipulator can change the winner if and only if $\text{WMG}(P) \in R_c$, which is equivalent to c being a possible winner. We recall that by the definition of $PW_1(d^*, \mathcal{C}')$, for any $G \in R_c$, $r(G + G_{c'}) = d^*$; for any $c' \neq c$ and any $G \in R_{c'}$, $r(G + G_{c'}) = c'$; and $d^* >_V c$. It follows that $V (=V_M)$ is dominated by U if and only if the $PW_1(d^*, \mathcal{C}')$ instance has a solution. \square

Theorem 6.3.4 can be used to prove that DOMINATION is NP-hard for Copeland, maximin, and voting trees, even when the number of undetermined pairs in each partial order is bounded above by a constant. It suffices to show that for each of these rules, there exist d^* and \mathcal{C}' such that $PW_1(d^*, \mathcal{C}')$ is NP-hard. To prove this, we can modify the NP-completeness proofs of the possible winner problems for Copeland, maximin, and voting trees by Xia and Conitzer Xia and Conitzer (2011a).

Corollary 6.3.5. DOMINATION *with partial orders is NP-hard for Copeland, maximin, and voting trees, even when the number of unknown pairs in each vote is bounded above by a constant.*

Proof. Copeland: We tweak the reduction in the NP-completeness proof of PW w.r.t. Copeland (Xia and Conitzer, 2011a, Theorem 3) by letting $D(c, v) = 1$ for any alternative $v \in \mathcal{V}$ and use the tie-breaking mechanism where $w > c > \text{Others}$. Let $d^* = w$, $\mathcal{C}' = B$, $V = U = [w > c > \mathcal{C}' > \text{Others}]$ and $W = [w > \mathcal{C}' > c > \text{Others}]$. It follows that the alternatives in B never wins the elections, and if c wins the election in an extension P of P_{po} , then the Copeland score of c is $8t + 1$ and the Copeland score of $w = 8t$. However, in the weighted majority graph $\text{WMG}(P) + G_{\mathcal{C}'}$, c loses to all alternatives in \mathcal{C}' in their pairwise elections, which means that the Copeland score of c is $t + 1$. Consequently w is the winner. On the other hand, for any extension P where c is not the winner, w is the winner, and w is also the winner in the weighted majority graph $\text{WMG}(P) + G_{\mathcal{C}'}$. Therefore, the PW instance is a $\text{PW}_1(d^*, \mathcal{C}')$ instance.

Maximin: We tweak the reduction in the NP-completeness proof of PW w.r.t. maximin (Xia and Conitzer, 2011a, Theorem 5) by letting $D(w', w) = t$. Let $d^* = w$, $\mathcal{C}' = \{w'\}$, $V = U = [w > c > w' > \mathcal{V}]$ and $W = [w > w' > c > \mathcal{V}]$. We adopt the tie-breaking mechanism where $w > c > \mathcal{V} > w'$. It is easy to check that w' never wins the elections. If c wins the election in an extension P of P_{po} , then the minimum pairwise score of c is $-t + 2$, and the minimum pairwise score of w and the alternatives in \mathcal{V} are $-t$. We note that in the majority graph $\text{WMG}(P) + G_{\mathcal{C}'}$, the minimum pairwise score of c is $-t$ (against w'), which means that $r(\text{WMG}(P) + G_{\mathcal{C}'}) = w$. For any extension P of P_{po} such that $r(P) \neq c$, it easy to check that the winner is in $\{w\} \cup \mathcal{V}$, and the minimum pairwise scores of them are the same as in the weighted majority graph $\text{WMG}(P) + G_{\mathcal{C}'}$. Therefore, the PW instance is a $\text{PW}_1(d^*, \mathcal{C}')$ instance.

Voting trees: We tweak the reduction in the NP-completeness proof of PW w.r.t. voting trees (Xia and Conitzer, 2011a, Theorem 7) by letting $D(c, d) = 1$. Let $d^* = w$, $\mathcal{C}' = \{d\}$, $V = U = [w > c > d > \text{Others}]$ and $W = [w > d > c > \text{Others}]$. For any extension P of P_{po} where c wins, the winner for the weighted majority graph $\text{WMG}(P) + G_{\mathcal{C}'}$ is w , because c loses to d in the first round, and w beats any other alternatives (except c) in their pairwise elections. For any extension P of P_{po} where c does not win, the winner is w . Therefore, the PW instance is a $\text{PW}_1(d^*, \mathcal{C}')$ instance. \square

Theorem 6.3.6. *Let r be a WMG-based voting rule. There is a polynomial-time reduction from $\text{PW}_2(d^*, \mathcal{C}')$ to DOMINATING MANIPULATION with partial orders, both under r .*

Proof. The proof is similar to the proof for Theorem 6.3.4. We note that d^* is the manipulator's top-ranked alternative. Therefore, if c is not a possible winner, then V ($= V_M$) is not dominated by any other vote; if c is a possible winner, then V is dominated by $U = [w > \mathcal{C}' > c > \text{Others}]$. \square

Similarly, we have the following corollary.

Corollary 6.3.7. *DOMINATING MANIPULATION with partial orders is NP-hard for Copeland and voting trees, even when the number of unknown pairs in each vote is bounded above by a constant.*

It is an open question if $\text{PW}_2(d^*, \mathcal{C}')$ with partial orders is NP-hard for maximin. However, we can directly prove that DOMINATING MANIPULATION is NP-hard for maximin by a reduction from x3C.

Theorem 6.3.8. *DOMINATING MANIPULATION with partial orders is NP-hard for maximin, even when the number of unknown pairs in each vote is no more than 4.*

Proof. We prove the hardness result by a reduction from x3C. Given an x3C instance $\mathcal{V} = \{v_1, \dots, v_q\}$, $\mathcal{S} = \{S_1, \dots, S_t\}$, where $q = t > 3$, we construct a DOMINATING MANIPULATION instance as follows.

Alternatives: $\mathcal{V} \cup \{c, w, w'\}$. Ties are broken in the order $w > \mathcal{V} > c > w'$.

First part P_1 of the profile: for each $i \leq t$, we start with the linear order $V_i = [w > S_i > c > (\mathcal{V} \setminus S_i) > w']$, and subsequently obtain a partial order O_i by removing the relations in $\{w\} \times (S_i \cup \{c\})$. For each $i \leq t$, we let O'_i be a partial order obtained from $V'_i = [w > v_i > \text{Others}]$ by removing $w > v_i$. We let O' be a partial order obtained from $V' = [w' > w > \text{Others}]$ by removing $w' > w$. Let P_1 be the profile composed of $\{O_1, \dots, O_t\}$, 2 copies of $\{O'_1, \dots, O'_t\}$, and 3 copies of O' . Let P'_1 denote the extension of P_1 that consists of V_1, \dots, V_t , 2 copies of $\{V'_1, \dots, V'_t\}$, and 3 copies of V' .

Second part P_2 of the profile: P_2 is defined to be a set of linear orders such that the pairwise score differences of $P'_1 \cup P_2 \cup \{V\}$ satisfy:

$$(1) \quad D(w, c) = 2t + \frac{2q}{3}, \quad D(w', w) = 2t + 6, \quad D(w', c) = 2t, \quad \text{and for all } i \leq q, \\ D(w, v_i) = 2t + 4 \quad \text{and} \quad D(v_i, w') = 4(t + q).$$

$$(2) \quad D(l, r) \leq 1 \quad \text{for all other pairwise scores not defined in (1).}$$

Manipulator's preferences: $V_M = [w > \mathcal{V} > c > w']$.

We note that in any extension of $P_1 \cup P_2$, after the manipulator changes her vote from V_M to $[w > \mathcal{V} > w' > c]$, the only change made to the weighted majority graph is that the weight on $w \rightarrow c$ increases by 2. Since w' never wins in any extension, if c does not win when the manipulator votes for V_M , then the winner does not change after the manipulator changes her vote to $[w > \mathcal{V} > w' > c]$. It follows from the proof of Theorem 6.3.4, Corollary 6.3.5, and Theorem 5 in Xia and Conitzer (2011a) that if the x3C instance has a solution, then V_M is dominated by $U = [w > \mathcal{V} > w' > c]$.

Suppose that the x3C instance does not have a solution, we next show that V_M is not dominated by any vote.

For the sake of contradiction, suppose the x3C instance does not have a solution and V_M is dominated by a vote U . There are following cases.

Case 1: There exist $v_i \in \mathcal{V}$ such that $w \succ_V v_i$ and $v_i \succ_U w$. We let P^* be the extension of $P_1 \cup P_2$ obtained from $P'_1 \cup P_2$ as follows. (1) Let $w \succ w'$ in 3 extensions of O' (we recall that there are $q > 3$ copies of O' in P_1). (2) Let $v_i \succ w$ in 2 extensions of O'_i . It is easy to check that in P^* , the minimum pairwise score of w is $-2t$ (via w') and the minimum pairwise score of v_i is $-2t$ (via w). Therefore, due to the tie-breaking mechanism, w wins. However, if the manipulator changes her vote from V_M to U , then the minimum pairwise score of w at most $-2t$ and the minimum pairwise score of v_i is at least $-2t + 2$, which means that v_i wins. We note that $w \succ_V v_i$. This contradicts the assumption that U dominates V_M .

Case 2: $w \succ_W v_i$ for each $v_i \in \mathcal{V}$. By changing her vote from V_M to U , the manipulator might reduce the minimum score of U by 2, increase the minimum score of c by 2, or increase the minimum score of w' by 2. Therefore, by changing her vote to U , the manipulator would either make no changes, make w lose, or make c win (we note that w' is not winning anyway). In each of these three cases the manipulator is not better off, which means that U does not dominate V_M . This contradicts the assumption. \square

For plurality and veto, there exist polynomial-time algorithms for both DOMINATION and DOMINATING MANIPULATION. Given an instance of DOMINATION, denoted by (r, P_{po}, V_M, V, U) , we say that U is a *possible improvement* of V , if there exists an extension P of P_{po} such that $r(P \cup \{U\}) \succ_{V_M} r(P \cup \{V\})$. It follows that U dominates V if and only if U is a possible improvement of V , and V is not a possible improvement of U . We first introduce an algorithm (Algorithm 4) that checks

whether U is a possible improvement of V for plurality.

Let c_{i^*} (resp., c_{j^*}) denote the top-ranked alternative in V (resp., U). We will check whether there exists $0 \leq l \leq n$, $d, d' \in \mathcal{C}$ with $d' \succ_{VM} d$, and an extension P^* of P_{po} , such that if the manipulator votes for V , then the winner is d , whose plurality score in P^* is l , and if the manipulator votes for U , then the winner is d' . We note that if such d, d' exist, then either $d = c_{i^*}$ or $d' = c_{j^*}$ (or both hold). To this end, we solve multiple maximum-flow problems defined as follows.

Let $\mathcal{C}' \subset \mathcal{C}$ denote a set of alternatives. Let $\vec{e} = (e_1, \dots, e_m) \in \mathbb{N}^m$ be an arbitrary vector composed of m natural numbers such that $\sum_{i=1}^m e_i \geq n$. We define a maximum-flow problem $F_{\mathcal{C}'}^{\vec{e}}$ as follows.

Vertices: $\{s, O_1, \dots, O_n, c_1, \dots, c_m, y, t\}$.

Edges:

- For any O_i , there is an edge from s to O_i with capacity 1.
- For any O_i and c_j , there is an edge $O_i \rightarrow c_j$ with capacity 1 if and only if c_j can be ranked in the top position in at least one extension of O_i .
- For any $c_i \in \mathcal{C}'$, there is an edge $c_i \rightarrow t$ with capacity e_i .
- For any $c_i \in \mathcal{C} \setminus \mathcal{C}'$, there is an edge $c_i \rightarrow y$ with capacity e_i .
- There is an edge $y \rightarrow t$ with capacity $n - \sum_{c_i \in \mathcal{C}'} e_i$.

For example, $F_{\{c_1, c_2\}}^{\vec{e}}$ is illustrated in Figure 6.1.

It is not hard to see that $F_{\mathcal{C}'}^{\vec{e}}$ has a solution whose value is n if and only if there exists an extension P^* of P_{po} , such that (1) for each $c_i \in \mathcal{C}'$, the plurality of c_i is exactly e_i , and (2) for each $c_{i'} \notin \mathcal{C}'$, the plurality of $c_{i'}$ is no more than $e_{i'}$. Now, for any pair of alternatives $d = c_i, d' = c_j$ such that $d' \succ_{VM} d$ and either $d = c_{i^*}$ or $d' = c_{j^*}$, we define the set of *admissible maximum-flow problems* A_{Plu}^l to be the set




FIGURE 6.1: $F_{\{c_1, c_2\}}^{\vec{e}}$.

of maximum flow problems $F_{c_i, c_j}^{\vec{e}}$ where $e_i = l$, and if $F_{c_i, c_j}^{\vec{e}}$ has a solution, then the manipulator can improve the winner by voting for U . More precisely, we define A_{Plu}^l as follows.

- If $i = i^*$ and $j \neq j^*$, then let $e_i = l$, $e_j = l + 1 - \delta(c_j, c_i)$, and $e_{j^*} = \min(l + 1 - \delta(j^*, i), e_j - 1 - \delta(j^*, j))$. For any $c_{i'} \in \mathcal{C} \setminus \{c_i, c_j, c_{j^*}\}$, we let $e_{i'} = \min(l + 1 - \delta(i', i), e_j - \delta(i', j))$. Let $A_{\text{Plu}}^l = \{F_{\{c_i, c_j\}}^{\vec{e}}\}$.
- If $i \neq i^*$ and $j = j^*$, then let $e_i = l$, $e_j = l - \delta(c_j, c_i)$, and $e_{i^*} = \min(l - 1 - \delta(i^*, i), e_j + 1 - \delta(i^*, j))$. For any $c_{i'} \in \mathcal{C} \setminus \{c_i, c_j, c_{i^*}\}$, we let $e_{i'} = \min(l - \delta(i', i), e_j + 1 - \delta(i', j))$. Let $A_{\text{Plu}}^l = \{F_{\{c_i, c_j\}}^{\vec{e}}\}$.
- If $i = i^*$ and $j = j^*$, then we define A_{Plu}^l as follows.
 - Let $e_i = l$, $e_j = l + 1 - 2\delta(c_j, c_i)$. For any $c_{i'} \in \mathcal{C} \setminus \{c_i, c_j\}$, we let $e_{i'} = \min(l + 1 - \delta(i', i), e_j + 1 - \delta(i', j))$.
 - Let $e'_i = e'_j = l$. For any $c_{i'} \in \mathcal{C} \setminus \{c_i, c_j\}$, we let $e'_{i'} = \min(l + 1 - \delta(i', i), e_j + 1 - \delta(i', j))$. Let $\vec{e}' = (e'_1, \dots, e'_m)$.
 - Let $A_{\text{Plu}}^l = \{F_{\{c_i, c_j\}}^{\vec{e}}, F_{\{c_i, c_j\}}^{\vec{e}'}\}$.

Algorithm 4 solves all maximum-flow problems in A_{Plu}^l to check whether U is a possible improvement of V .

Algorithm 4: PossibleImprovement(V, U)

Let $c_{i^*} = \text{Alt}(V, 1)$ and $c_{j^*} = \text{Alt}(U, 1)$.
for any $0 \leq l \leq n$ and any pair of alternatives $d = c_i, d' = c_j$ such that $d' \succ_{V_M} d$ and either $d = c_{i^*}$ or $d' = c_{j^*}$ **do**
 | Compute A_{Plu}^l .
 | **for** each maximum-flow problem $F_{\mathcal{C}'}^{\vec{e}}$ in A_{Plu}^l **do**
 | **if** $\sum_{c_i \in \mathcal{C}'} e_i \leq n$ and the value of maximum flow in $F_{\mathcal{C}'}^{\vec{e}}$ is n **then**
 | Output that the U is a possible improvement of V , terminate the algorithm.
 | **end**
 | **end**
end
Output that U is not a possible improvement of V .

A_{Plu}^l The algorithm for DOMINATION (Algorithm 5) runs Algorithm 4 twice to check whether U is a possible improvement of V , and whether V is a possible improvement of U .

Algorithm 5: Domination

if PossibleImprovement(V, U) = “yes” and PossibleImprovement(U, V) = “no”
then
 | Output that V is dominated by U .
end
else
 | Output that V is not dominated by U .
end

The algorithm for DOMINATING MANIPULATION for plurality simply runs Algorithm 5 $m - 1$ times. In the input we always have that $V = V_M$, and for each alternative in $\mathcal{C} \setminus \{\text{Alt}(V, 1)\}$, we solve an instance where that alternative is ranked first in U . If in any step V is dominated by U , then there is a dominating manipulation; otherwise V is not dominated by any other vote. The algorithms for DOMINATION and DOMINATING MANIPULATION for veto are similar.

6.4 Summary

We have shown in this chapter that for many common voting rules, restricting the manipulator's information about the other voters' votes is an effective way to make dominating manipulation computationally hard, or even impossible. Analysis of manipulation with partial information provides insight into what needs to be kept confidential in an election. For instance, in a plurality or veto election, revealing (perhaps unintentionally) part of the preferences of non-manipulators may open the door to strategic voting.

In Chapter 4, 5, and 6 we have seen some recent work and discussions on using computational complexity as a barrier against manipulation. However, a more important question that should be asked is: Why should we even try to prevent manipulation and other types of strategic behavior? In the next chapter, we will show that indeed, the strategic behavior of the voters can lead to extremely undesirable outcomes, in a type of voting games which we call *Stackelberg voting games*.