Homework #1

Deadline: Tuesday, 3 November 2026, 10:00

Exercise 1 (10 points)

Provide a one-paragraph definition of the voting rule you presented in class, as well as one sentence each on one positive and one negative property of that rule. Remember to cite the resources you used to find the information you presented. This exercise is intended to serve as a reminder of what you presented in class; what counts is the presentation itself.

If you did not present, instead pick a paper published in 2026 on any topic in computational social choice you like and submit a one-page summary written in your own words.

For this exercise we will use a simplified grading scale. It will be graded with a 10 for any serious attempt, with a 6 for any barely acceptable attempt, and with a 1 in all other cases.

Exercise 2 (10 points)

For the sake of simplicity, for your solution restrict yourself to the case of $m = 3$ alternatives.

In analogy to the definition of a Condorcet winner, a *Condorcet loser* is an alternative that would lose against any other alternative in a pairwise majority contest. It seems uncontroversial to say that a Condorcet loser, when it exists, is a particularly unattractive alternative. So we should hope that the voting rule we use will not elect it.

(a) Give an example that shows that the plurality rule *can* elect a Condorcet loser.

(b) Prove that the Borda rule *never* elects a Condorcet loser.

(c) Is the Borda rule the only positional scoring rule that never elects a Condorcet loser? Provide either a proof of the uniqueness of Borda or an example for another such rule.

Hint: For part (c), do not confuse different representations of a rule with different rules. For instance, the positional scoring rule induced by the scoring vector $(7, 5, 3)$ is the same rule as the positional scoring rule induced by $(2, 1, 0)$. Both are representations of the Borda rule.

Exercise 3 (10 points)

Recall that most natural voting rules are not *resolute* and will sometimes return tied winners. Therefore, for this exercise, assume that every irresolute voting rule you might consider is paired with a *lexicographic tie-breaking rule*. So there always is a single election winner.

Consider an election with n voters and m alternatives. Suppose alternative x wins. Let us call the *regret* of voter i the number of alternatives that dominate x in her preference order (i.e., this is a number between 0 and $m-1$). We might compare voting rules by the *average regret* they generate, across all n voters, for a given profile. Now suppose every voter draws her true preference order from the uniform probability distribution over all $m!$ logically

possible orders (this is known as the *impartial culture assumption*). Under this assumption, we can calculate the *expected regret* of voter i under voting rule F. If we average over all voters, we obtain the *average expected regret* of F. Answer the following questions:

(a) Suppose you want to use a voting rule that minimises average regret. Which of the rules discussed in class is the best choice according to this design objective? Prove that, relative to this objective, the rule you have identified really is better than any of the other voting rules we have seen in class.

(b) Show that the average *expected* regret is not a useful design objective by proving that, as n goes to infinity, all voting rules have the same average expected regret.