Homework #1

Deadline: Tuesday, 7 April 2020, 18:00

Question 1 (10 marks)

In analogy to the definition of a Condorcet winner, a *Condorcet loser* is an alternative that would lose against every other alternative in a pairwise majority contest.

(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.

Remark: It is in fact possible to show that the Borda rule is the *only* positional scoring rule (with a strictly descending scoring vector) that satisfies this property. This is not part of the exercise, but you may still want to think about it.

Question 2 (10 marks)

I ask you to not discuss this particular exercise with anybody before the submission deadline, as that would render it pretty much pointless almost immediately.

In the late 1980's, Hervé Moulin published a paper in which he shows that every voting rule that always elects the Condorcet winner when it exists must allow for situations where a voter has an incentive to not vote at all rather than to vote in line with her truthful preferences. Track down this paper and then answer the following questions:

(a) Provide the full bibliographic information for the paper.

(b) Briefly describe how you tracked down the paper.

(c) Give a compact and precise statement of the aforementioned result, using the notation and terminology of the course.

(d) Find one political scientist, one philosopher, and one computer scientist who have cited Moulin's paper in the past 10 years. For each of them, provide full bibliographic information of the contribution in question (including a link), provide evidence that the person in question indeed is a political scientist/philosopher/computer scientist, and describe the reason for the citation in a few words.

Question 3 (10 marks)

Recall that most voting rules are not *resolute*. Therefore, for this exercise, assume that every irresolute voting rule is paired with a *lexicographic tie-breaking rule*. So there always is a single 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 be interested in the *average regret* a given voting rule F generates 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.