

Assignment mechanisms with public preferences and independent types*

Francisco Silva[†]

September 24, 2021

Abstract

The literature on delegation considers the problem of an uninformed decision maker and an informed but biased agent. I extend that analysis to the case of multiple agents under two assumptions: independent private information and public preferences. In the optimal mechanism, agents assign points to the various alternatives, which then get mapped into scores, so that the alternative with the largest score wins. Each alternative's score is the sum of points received plus an extra term that is larger when the agents who have a strong preference for that alternative assign points to the alternatives they like less.

JEL classification: D82,

Keywords: delegation, mechanism design, contests.

*I would like to thank Mustafa Dogan, Ju Hu, Samir Mamadehussene and Juan Pereyra for their useful comments.

[†]Department of Economics, Pontificia Universidad Catolica de Chile, Vicuña Mackenna 4860, Piso 3. Macul, Santiago, Chile. (email: franciscosilva@uc.cl).

1 Introduction

In this paper, I consider the classical problem of an uninformed decision maker (DM) who designs a mechanism that incentivizes informed agents to share their private information with her. The agents and the DM have different preferences over the various alternatives available to the DM, so the DM cannot simply ask for the agents' private information. When there is a single agent, the problem of finding the DM's optimal mechanism has been solved. Holmstrom (1984) points out that, by the revelation principle, it trivially follows that the DM should *delegate* her decision to the agent; specifically, the DM first pre-selects a subset of (lotteries over) alternatives and then the agent selects one of them. Delegation mechanisms allow the DM to use some of the private information of the agent while also mitigating the potential danger of having the agent choose an alternative that the DM does not like (through the initial choice of admissible alternatives). The vast *delegation* literature that followed Holmstrom (1984) has focused almost entirely on the one-agent case and has determined many of the properties of the optimal delegation mechanism (Melumad and Shibano, 1991; Armstrong, 1995; Martimort and Semenov, 2006; Alonso and Matouschek, 2008; Kovac and Mylovanov, 2009; Koessler and Martimort, 2012; Amador and Bagwell, 2013).

Much less is known of the general problem with multiple agents. If the setting has transfers, i.e., if the alternatives the DM chooses from are multidimensional and one of the dimensions enters linearly in the utility functions of the agents, then optimal mechanisms have been found in a variety of settings (Myerson, 1981, Myerson and Satterthwaite, 1983, Maskin and Riley, 1984, among many others). However, settings without transfers have proven to be harder to analyze. Holmstrom (1984) mentions that one possibly interesting mechanism would be the delegation mechanism: the DM pre-selects a set of alternatives and then delegates the final decision to one of the agents. Indeed, Alonso, Brocas and Carrillo (2014) and Gan, Hu and Weng (2021), who study settings with only two agents, find that a form of delegation is optimal among mechanisms that are incentive compatible in dominant strategies. There is, however, no mention of whether that mechanism is still optimal for the DM among mechanisms that are Bayes-Nash incentive compatible, a concept that is more standard in mechanism design; in fact, I show it is not. Even the social choice literature, which studies a variation of this problem by assuming that the DM has some specific utility function like the (weighted) sum of the agents' utilities, has only been able to fully characterize optimal incentive compatible mechanisms where there are only two

alternatives available to the DM.¹

Part of the reason why the problem with multiple agents has been so hard to solve is that the differences in the players' preferences vary depending on the agents' private information, i.e., in addition to not knowing the agents' private information, the DM does not know their preferences. Indeed, in most of literature, there is a one-to-one mapping between the agents' private information and their (private) preferences. However, there are a series of applications where agents have public preferences *and* private information that the DM would like to have. For example, suppose that the State assembles a group of experts that provide advice on how to legislate over gun ownership. Each expert has private information over the potential danger (or lack thereof) of legalizing gun ownership and is funded by an interest group that seeks to either further or restrict gun ownership. It is publicly known which groups support which experts. Another example is the problem of an arbitrator who must decide how to split the assets of some bankrupted firm among the many self-interested parties who have information on how those assets should be distributed. The example I refer to most often is that of a contest between multiple agents, where each agent wants to be chosen but has information on which of the agents would be a better winner. For example, the members of the board of some organization must decide a leader among them. Every member wants to be the one selected and has information over which member would be the best leader.

I study these types of applications by considering the general problem with multiple agents but assuming that agents have public preferences, i.e., I break the direct link between preferences and private information. In this paper, I focus on the case of independent private information.

The optimal mechanism is not obvious at first glance. Clearly any form of (weighted) majority rule works poorly because it gives incentives for agents to vote for their favorite alternatives rather than for those which provide more value to the DM. It works especially poorly in contests, where each of the contestants has a vote, because every

¹If there are only two alternatives and the agents' private information is independent, Azrieli and Kim (2014) have shown that the optimal incentive compatible mechanism is the weighted majority rule. In symmetric settings, the optimal mechanism becomes the simple majority rule (Schmitz and Troger, 2012). When there are more than two alternatives, little is known about optimal mechanisms in general, even though there is some promising work done on various specific settings (Borgers and Postl, 2009, Miralles, 2012, Goldlucke and Troger, 2018, and Bhaskar and Sadler, 2020). There is some characterization of optimal ordinal mechanisms, which only use the ordinal preferences of the agents (Majumbar and Sen, 2004). However, it is also known that ordinal mechanisms are not optimal (Kim, 2017). There is also a large literature which studies mechanisms that are incentive compatible in dominant strategies (e.g. Gibbard, 1973; Satterthwaite, 1975; etc.).

agent would vote for himself, so that the winner would end up being chosen randomly. More broadly, score mechanisms, as introduced by Myerson (2002), also have the same problem that agents simply vote for their preferred alternatives.²

One alternative that is very common in contests is the *adjusted* majority rule, which works as the simple majority rule except that each agent cannot vote for himself.³ The appeal of the adjustment is that it seemingly negates the primary desire of each contestant to vote for himself. However, this mechanism also has its share of problems. First, in asymmetric settings where, a priori, there are some contestants that are better than others, the adjusted majority rule does not induce agents to vote for the candidates they think would provide more value to the DM. Second, even in symmetric settings, where contestants do have an incentive to vote for who they think is the best candidate, the mechanism still excludes a lot of information because, not only is it an ordinal mechanism, it only uses the very top of the order of each candidate, i.e., it only asks for each candidate's top option.

Delegation is certainly a plausible mechanism for the DM. Indeed, going beyond Holmstrom (1984), the DM could select a set of agents and a subset of alternatives such that every agent selected is indifferent among all alternatives and then ask that group of agents to collectively determine the winning alternative. However, I show in the paper that, in general, none of these mechanisms is optimal.

I find that the (Bayes-Nash incentive compatible) optimal mechanism is as follows. First, each agent assigns points to each alternative. Those points then determine each alternative's score and then, the alternative with the largest score is chosen. The difference to score mechanisms is that each alternative's score does not depend only on the sum of points received. Each alternative j 's score is the sum of two parts: α_j and β_j . The first part α_j represents the sum of points alternative j receives like in a score mechanism. The second part β_j has two key properties. On the one hand, each β_j is larger when agents who have a strong preference for alternative j assign a lot of points to their least preferred alternatives. On the other hand, when some agent i does assign points to his least favorite alternatives, the increase in each β_j is weighted by agent i 's preferences. So, for example, say that there are three alternatives - a , b and c - and that agent 1 prefers alternative a , then b and then c . Suppose agent 1 observes an increase in the value for the DM of alternative c and reacts by increasing the amount of points

²In a score mechanism, each agent assigns points to each alternative and then the alternative with the most amount of points wins. Majority rule is a score mechanism.

³High-profile examples of the adjusted majority rule include the election of the pope, which did not allow cardinals to vote for themselves up until 1945, and the election of the winner in the Eurovision song contest among many other examples that can easily be found in any search engine.

assigned to alternative c . On the one hand, α_c increases, which increases alternative c 's score. On the other hand, because agent 1 is giving additional points to his least favorite alternative, β_a, β_b and β_c all increase, but β_a increases more than β_b , which increases more than β_c (for each set of reports of the other agents). Overall, this leads to alternative c replacing alternatives a and b as the winning alternative for some reports of the other agents (which harms agent 1), while for some other reports, alternative b is replaced by alternative a (which benefits agent 1). These two effects taken together leave agent 1 indifferent but improve the DM's expected payoff on account of assigning a larger probability to alternative c .

This explanation also provides intuition for the second result of the paper, which is about the optimal mechanism when there are only two alternatives (labeled a and b) available to the DM. Suppose that agent 1 strictly prefers a over b . It is not possible that agent 1 has enough incentives to give points to alternative b if those points increase alternative b 's probability of winning, because that would lower agent 1's expected utility. So, the DM can only benefit from the agents' private information when there are at least three alternatives. As a result, I find that, provided agents are not indifferent between the two alternatives, it is not possible for the DM to gain from interacting with the agents; she should just select the winning alternative with whatever public information there is. While this result is also true in the standard delegation literature with a single agent, it contrasts with some of the prominent literature on arbitration, which considers a very similar setting - two privately informed agents dispute the ownership of an asset - but assumes that agents have perfectly correlated signals (Gibbons, 1998; Mylovanov and Zapechelnuk, 2013). Indeed, in Gibbons (1998), the DM can achieve her favorite outcome if she is able to observe her own exogenous and imperfect signal.⁴

The optimal mechanism takes on a simpler form in contests, where each agent is an alternative and is only interested in being selected: Agents assign points to every other agent; i.e., they do not assign points to themselves. Then, each agent's score is the sum of points received plus an increasing function of the points assigned by the agent. The agent wins the contest if his score is the highest. Once again, when some agent i assigns a lot of points to some agent j , he does not alter his own probability of being chosen (which is all he cares about) but does increase the probability of agent j

⁴Pereyra and Silva (2021) and Bloch et al. (2021), who study a similar model but with multiple agents, also show how the DM can incentivize agents to communicate when she has access to exogenous signals.

being chosen at the expense of the other agents.⁵

If we further assume symmetry and restrict attention only to ordinal mechanisms, the optimal mechanism becomes extremely simple: agents rank every agent, including themselves. Then, each agent's score is an increasing function of how others have ranked him and a decreasing function of how the agent has ranked himself. Once again, the winner is the agent with the largest score. Unlike cardinal mechanisms, which ask each agent to be able to quantify by how much some alternatives are better than others, ordinal mechanisms only ask agents to be able to rank the alternatives, which makes them more appealing (Bogomolnaia and Moulin, 2001; Carroll, 2018). Assuming ex-ante symmetry between the alternatives is particularly compelling in contests where the DM is neutral; in asymmetric settings, the DM would prefer to design mechanisms that favor some alternatives over others.

The paper continues as follows. In section 2, I present the model; in section 3; I provide a simple example that illustrates some of the results; in section 4, I discuss optimal mechanisms in general; in section 5, I discuss contests; in section 6, I conclude.

2 Model

There is a decision maker (DM) and I agents. The DM must decide between J alternatives. The payoff of each agent i depends only on the alternative j that is chosen and is denoted by $u_{ij} \geq 0$. For each i , vector $u_i = (u_{i1}, \dots, u_{iJ})$ is assumed to be public. Each agent i observes a private signal $s_{ij} \in S_{ij} \subset \mathbb{R}$ relative to each alternative j . For each i , vector $s_i = (s_{i1}, \dots, s_{iJ}) \in S_i \equiv S_{i1} \times \dots \times S_{iJ}$ represents agent i 's private type, which is assumed to be independent across agents. Each set S_{ij} is assumed to be finite and the probability distribution of s_i is denoted by p_i .

The DM's payoff depends on the alternative chosen and on the private information held by the agents. I denote the DM's payoff when choosing alternative j by $v^j(s)$, where $s = (s_1, \dots, s_I) \in S$, and assume that

$$v^j(s) = \sum_{i=1}^I v_i^j(s_{ij}),$$

⁵Goldlucke and Troger (2018) study a specific contest where the "winner" is chosen to perform some public service and assume the agents' preferences are type-dependent. De Clippel et al. (2021) study a dynamic contest between two agents.

where each function $v_i^j : S_{ij} \rightarrow \mathbb{R}$ is strictly increasing for each i . Notice that, without loss of generality, one can assume that $v_i^j(s_{ij}) = s_{ij}$ for all i and j .⁶

Note: Before proceeding, it is important to remind the reader of the usual caveat when dealing with independent types (e.g., Branco, 1996). Even though agents have independent private information, that does not mean that their opinions over which alternatives provide more value for the DM are also independent. Take as an example the case where a hiring committee decides over which candidate to hire. Three experts are a part of the hiring committee - one in micro theory, one in macroeconomics and one in empirical microeconomics - and there are three candidates; one from each of the experts' field. The experts' preferences are public: each expert prefers to hire the candidate of his field and is indifferent between the other two. However, the Dean, who is the decision maker of the scenario, is not indifferent and would like to select the most able candidate.

When forming an opinion over each candidate, each member of the recruiting committee relies on public information and on private information. The public information is the information that is also available to the Dean. It includes information on the candidates' CV, their publication and teaching records, how complimentary are the recommendation letters, etc.. The private information each expert has is more taste-based and relies heavily on each agent's introspection: is the candidate's work important?; is it novel?; is the candidate a good fit with the department? More formally, it is likely that agent i 's opinion over candidate j denoted by $o_{ij} = c_j + \varepsilon_{ij}$, where c_j is common to all agents and represents the public information available about candidate j and ε_{ij} represents the private information held by agent i about candidate j . The assumption of this paper is that ε_{ij} is independent across i for all j ; it is not that o_{ij} is independent across i . Indeed, if c_j is random, then the agents' opinions over the different candidates would be positively correlated even when the agents' private information is independent.

An allocation is a function $x : S \rightarrow [0, 1]^J$ such that

$$\sum_{j=1}^J x_j(s) = 1$$

⁶One can simply assume that the signals each agent observes are $v_i^j(s_{ij})$ instead of s_{ij} .

for all $s \in S$. An allocation x is incentive compatible (IC) if

$$E_{s_{-i}} \left(\sum_{j=1}^J x(s_i, s_{-i}) u_{ij} \right) = E_{s_{-i}} \left(\sum_{j=1}^J x(s'_i, s_{-i}) u_{ij} \right)$$

for all $s_i, s'_i \in S_i$ and for all i . Notice that this incentive compatibility condition is a direct consequence of assuming that the agents' private types are independent. In general, incentive compatibility forces the allocation to be such that each agent always prefers to report his type truthfully given that other agents also report truthfully. Because types are independent, the beliefs that each type has over the reports of the other agents will be the same. This means that different types face the same set of lotteries over alternatives to choose from. Therefore, all types must be indifferent.

By the revelation principle (Myerson, 1979), the problem of finding an optimal mechanism reduces to that of finding an optimal allocation. An optimal allocation maximizes the DM's expected payoff among all IC allocations. An optimal allocation trivially exists, because there is always a solution to any linear program where the (finite) choice variables are probabilities. An optimal mechanism is a mechanism for which there is a Bayes-Nash equilibrium which induces an optimal allocation.

In part of the paper, I specifically study contests. A contest is such that $I = J$ and

$$u_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}$$

for all $i, j \in I \times J$.

3 Example

In order to boost morale, the manager of a firm decides to organize a costume contest during Halloween in which the firm's employees participate. Each employee is supposed to wear a costume during the Halloween party and, at the end of the party, the best costume is selected and the employee is rewarded. For transparency, the manager does not participate in determining the winning costume; instead, she determines the mechanism by which a winner is chosen.

Each one of the $I \geq 3$ contestants has very simple and known preferences: they want to win and, if they cannot win, they are indifferent as to whom should win. Formally,

$u_{ii} = 1$ and $u_{ij} = 0$ for all i and for all $j \neq i$. During the party, each contestant observes everyone else's costume. To keep matters simple (I study the general case in the text), let us assume that, after the party, each contestant i has one costume that he finds the best one and all other costumes are equally good.⁷ For example, if $I = 3$, then

$$s_i \in \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\},$$

where $s_i = (1, 0, 0)$, for example, is interpreted as contestant i thinking that contestant 1's costume is the best one. Furthermore, assume a symmetric setting so that each vector s_i has the same probability of $\frac{1}{I}$ for all i . The manager, who strives for fairness, is assumed to have a payoff of $\sum_{i=1}^I s_{ij}$ when contestant j is chosen to win the contest.

The manager knows of a few popular mechanisms. One option would be to elect the winner through majority rule. Naturally, that would not work well as all contestants would vote for themselves, so there would be a massive one-vote tie at the end of the party.

Another option would be to divide the employees into two groups - jurors and contestants - before the party starts. Jurors would be told they could not win and would be instructed to deliberate together and determine a winner among the contestants. This more general form of delegation does make sense in that the mechanism provides enough incentives for jurors to choose the best costume among the group of contestants. It turns out, however, that delegation is not optimal. To see why that is, let us consider the simpler case when $I = 3$. In this case, the only delegation that makes sense for the manager is to delegate on one single agent who then picks among the other two. Seeing as the setting is completely symmetric, the manager is indifferent as to whom to delegate, so she selects random delegation, where each contestant is made the (single) juror with probability $\frac{1}{3}$. Consider the following vector of signals:

$$s_1 = (0, 1, 0), s_2 = (1, 0, 0), s_3 = (1, 0, 0).$$

Under random delegation, contestant 1 wins with probability $\frac{2}{3}$, contestant 2 wins with probability $\frac{1}{3}$ and contestant 3 does not win. Seeing as more employees prefer contestant 1's costume over contestant 2's costume, the manager would be better off if contestant 1 won with probability 1. Indeed, the same happens for all of the following five equally

⁷Formally, assume that $s_{ij} \in \{0, 1\}$ for all i, j and that $\sum_{j=1}^J s_{ij} = 1$.

likely vectors of signals:

$$s_1 = (0, 0, 1), s_2 = (1, 0, 0), s_3 = (1, 0, 0),$$

$$s_1 = (0, 1, 0), s_2 = (1, 0, 0), s_3 = (0, 1, 0),$$

$$s_1 = (0, 1, 0), s_2 = (0, 0, 1), s_3 = (0, 1, 0),$$

$$s_1 = (0, 0, 1), s_2 = (0, 0, 1), s_3 = (1, 0, 0),$$

$$s_1 = (0, 0, 1), s_2 = (0, 0, 1), s_3 = (0, 1, 0).$$

In all of these, the most preferred costume only wins with probability $\frac{2}{3}$. Suppose we change the random delegation mechanism as follows: whenever one of these six vectors occurs, the contestant with the most preferred costume wins with probability 1; for all other vectors, the winner is chosen according to the random delegation mechanism.

It is clear that the new mechanism makes the manager better off, as the most liked costume is chosen more often. Moreover, it also follows that each contestant's expected payoff of choosing each signal remains unchanged. For example, consider how contestant 1's expected payoff when his type is $s_1 = (0, 1, 0)$ changes from random delegation to the new mechanism. Agent 1 increases his chances of winning by $\frac{1}{3}$ when $s_2 = (1, 0, 0)$ and $s_3 = (1, 0, 0)$, and loses $\frac{1}{3}$ when $s_2 = (1, 0, 0)$ and $s_3 = (0, 1, 0)$. Seeing as those two events are equally likely, agent 1's expected payoff given type $s_1 = (0, 1, 0)$ remains the same. Therefore, the new mechanism is also incentive compatible and is strictly preferred by the manager to the random delegation mechanism.

After realizing that the delegation mechanism is not optimal, the manager again considers the majority rule but with a twist; contestants are not allowed to vote for themselves. In a symmetric setting such as this, contestants have an incentive to vote "truthfully": whenever his preferred costume is not his own, each contestant votes for his preferred costume; if his own costume is his preferred costume, he randomizes over whom to vote for. When $I = 3$, this mechanism generates an expected payoff for the manager of $\frac{3}{2}$. The reason why this mechanism is not optimal is that it does not distinguish between contestant i voting for some contestant j who contestant i thinks has the best costume, and contestant i , convinced that he has the best costume, but being forced to vote for somebody else, voting for contestant j . Ideally for the manager, in the former case, contestant j 's probability of winning would be larger than in the latter.

Alerted to this issue, the manager quickly suggests that, rather than being forced

to vote for others, each contestant should be allowed to abstain. Such a mechanism works very poorly, however, because each contestant, interested in maximizing his own chances of winning, would prefer to abstain over voting for his favorite costume. Indeed, the optimal mechanism, which I characterize below, is similar to the majority rule with abstention but contestants who abstain get an extra penalty.

In the text, I characterize the optimal mechanism in general (Propositions 1 and 2) and then provide a further characterization for the case of contests (Propositions 3 and 4). Applying these results to the example returns the following optimal mechanism. Each contestant has a vote. Let $z_i \in \{1, \dots, I\}$ denote the vote of contestant i and notice that each contestant can vote for himself. One can interpret a vote on one's self as an abstention. Once everyone has voted, each contestant receives a score as follows:

$$score_i = \sum_{j \neq i} \mathbf{1}\{z_j = i\} - k \mathbf{1}\{z_i = i\}$$

for $k \geq 0$. In words, candidate i 's score is the sum of votes received by others minus some penalty k if candidate i decides to abstain. The winner of the contest is the contestant with the largest score. The number of contestants I determines k and the tie breaking rules (the probability that each contestant wins when multiple contestants have the highest score). Specifically, both k and the tie breaking rules are chosen such that each contestant is indifferent between abstaining and voting for others.

One can show that, if I is small, then $k = 1$ and, in the event of a tie, the tie-breaking rules favor contestants who abstain.⁸ For example, if $I = 3$, then $k = 1$, and, whenever there is a tie between a contestant who abstained and a contestant who voted for somebody else, the contestant who abstained wins with probability $\frac{14}{18}$.⁹ In that case, the manager's expected payoff is $\frac{127}{81} > \frac{3}{2}$. When I is sufficiently large, $k = 0$ and the tie breaking rules favor those who vote for others.

Even though the tie breaking rules are important in the mechanism, as they ensure that agents are indifferent, if the information structure of the agents is sufficiently rich (i.e., if the types of the agents are approximately continuous), then ties happen with arbitrarily small probability (I show this in the appendix).

⁸If all contestants with the highest score have abstained or have not abstained, the tie breaking rules are even.

⁹For example, if $s_1 = (1, 0, 0)$, $s_2 = (1, 0, 0)$ and $s_3 = (0, 0, 1)$, contestants 1 and 3 have the highest score of 0. Contestant 1 then wins with probability $\frac{14}{18}$ while contestant 3 wins with probability $\frac{4}{18}$.

4 Optimal mechanisms

If the DM could trust the agents to report truthfully, she should simply implement the following *naive* mechanism: agents assign points to each alternative; the alternative with the most points received is chosen. Formally, each agent i would report s_{ij} for each alternative j (so that s_{ij} would be the amount of points assigned by agent i to alternative j) and the winning alternative would be

$$j^w \in \arg \max \alpha_j(s),$$

where $\alpha_j : S \rightarrow \mathbb{R}$ is such that

$$\alpha_j(s) = \sum_{i=1}^I s_{ij}.$$

Naturally, were the DM to commit to implementing this mechanism, the agents would simply assign as many points as possible to their preferred alternatives and as few points as possible to their least favorite alternatives. Therefore, as discussed in the example, the DM must somehow penalize agents who assign a lot of points to their favorite alternatives enough to keep them indifferent between all of their reports. As I discuss in more detail in the next section, in contests, this is easier to do as it is enough to punish agents who assign a lot of points to themselves and very few points to others. In general, though, it is not clear a priori, what it means for an agent to assign a lot of points to his favorite alternatives. The following definition formalizes this idea.

Definition 1 For each agent i and any pair $s_i, s'_i \in S_i$, $s_i \succ_i s'_i$ if and only if, for all pairs of alternatives j, j' ,

$$u_{ij} > u_{ij'} \Rightarrow s_{ij} - s_{ij'} \geq s'_{ij} - s'_{ij'}.$$

When $s_i \succ_i s'_i$ agent i assigns more points to his favorite alternatives with vector s_i than with vector s'_i in the sense that, should the DM implement the naive mechanism, agent i would always (weakly) benefit from reporting s_i over s'_i . For example, suppose that $J = 3$ and that $u_{i1} > u_{i2} > u_{i3}$. Consider signals $s_i = (2, 1, 0)$ and $s'_i = (0, 1, 2)$. Were the DM to implement the naive mechanism, agent i would prefer to report s_i over s'_i , because α_1 would increase more than α_2 and α_2 would increase more than α_3 . Therefore, $s_i \succ_i s'_i$. Notice that \succ_i is a partial order because, for example, if

$s_i'' = (0, 2, 0)$, then signals s_i and s_i'' would not be comparable: when going from s_i'' to s_i , α_1 increases more than α_2 but α_2 increases less than α_3 .

Any vector of functions $\mu = (\mu_1, \dots, \mu_I)$ such that $\mu_i : S_i \rightarrow \mathbb{R}$ for all i is called *regular* if, for each i and for any pair $s_i, s_i' \in S_i$ for which $s_i \succ_i s_i'$, $\mu_i(s_i) \leq \mu_i(s_i')$.

The first result of the paper is that the following class of direct mechanisms are optimal. Each agent i simultaneously assigns points $s_{ij} \in S_{ij}$ to each alternative j . Vector s determines each alternative j 's score as follows:

$$\text{score}_j(s) = \alpha_j(s) + \beta_j(s),$$

where

$$\beta_j(s) \equiv \sum_{i=1}^I \mu_i(s_i) u_{ij}$$

for some $\mu = (\mu_1, \dots, \mu_I)$. The alternative with the largest score is chosen. Notice that each such mechanism is completely described by a vector of functions μ and a tie-breaking rule T , which determine which alternative is chosen with what probability, should there be multiple alternatives with the highest score. If one defines $\wp(J)$ as the set of all subsets of $\{1, \dots, J\}$ excluding the empty set, one can formally define T as follows: $T : S \times \wp(J) \rightarrow [0, 1]^J$ such that

$$\sum_{j=1}^J T_j(s, \Omega) = 1$$

and

$$j \notin \Omega \Rightarrow T_j(s, \Omega) = 0$$

for all j , for all $s \in S$ and for all $\Omega \in \wp(J)$, where each $T_j(s, \Omega)$ is interpreted as the probability that alternative j is chosen given $s \in S$ and given that all alternatives in set $\Omega \in \wp(J)$ have the highest score. Each direct mechanism with function μ and tie breaking rule T is denoted by $Mech(\mu, T)$. If $Mech(\mu, T)$ has a Bayes-Nash equilibrium where agents report truthfully, I say that the allocation that is generated by that equilibrium is truthfully induced by $Mech(\mu, T)$. Formally, if $Mech(\mu, T)$ truthfully induces allocation x^* , then

$$x_j^*(s) = T_j(s, \Omega^\mu(s))$$

for all j and $s \in S$, where $\Omega^\mu(s)$ represents the set of alternatives with the highest score given μ and s .

Proposition 1 *i) For any optimal allocation x^* , there is a regular $\mu : S \rightarrow \mathbb{R}^I$ and a tie breaking rule $T : S \times \wp(J) \rightarrow [0, 1]^J$ such that $Mech(\mu, T)$ truthfully induces x^* .*

ii) If there is some $\mu : S \rightarrow \mathbb{R}^I$ such that

$$\sum_{s_i \in S_i} p_i(s_i) \mu_i(s_i) = 0 \quad (1)$$

for all i and a tie-breaking rule $T : S \times \wp(J) \rightarrow [0, 1]^J$ for which $Mech(\mu, T)$ truthfully induces some allocation x^ , then allocation x^* is optimal.*

Part i) of proposition 1 states that one can restrict attention to the type of direct mechanisms described above with regular functions μ . Even though these mechanisms resemble score mechanisms, in that the alternative with the largest score is the one that is chosen, they are fundamentally different because each alternative's score is not just the sum of points received; it has a second component β , which aligns the incentives of the agents with the DM's. Whenever agent i reports $s_i \succ_i s'_i$ he assigns more points to his favorite alternatives than with s'_i (in the sense described above). To keep agent i indifferent between reporting s_i and s'_i , the mechanism rewards agent i when reporting s'_i as follows. Because $s_i \succ_i s'_i$, then $\mu_i(s'_i) \geq \mu_i(s_i)$, which, in turn, implies that $\beta_j(s'_i, s_{-i}) \geq \beta_j(s_i, s_{-i})$ for all j and for all s_{-i} , i.e., every alternative's β is larger when agent i gives more points to the alternatives he likes the least (i.e., when he reports s'_i). However, this increase in the alternatives' scores is weighted by agent i 's preferences; the increase is larger for the alternatives that agent i prefers.

Part ii) states that, in order to find an optimal mechanism, it is enough to find a μ such that (1) holds and a tie-breaking rule T such that agents prefer to report truthfully given mechanism $Mech(\mu, T)$. Condition (1) states that the average $\mu_i(s_i)$ is zero, thereby preventing $\mu_i(s_i)$ to be consistently high, which rules out sub-optimal allocations where agent i is a dictator.

4.1 Example revisited

Let us consider again the example of section 3 when $I = J = 3$. If one realizes that, for any symmetric setting, there is an optimal allocation that is symmetric, the

problem becomes much simpler and can be solved using only part i) of proposition 1. Specifically, by symmetry, it follows that

$$\mu_1(1, 0, 0) = \mu_2(0, 1, 0) = \mu_3(0, 0, 1) \equiv a$$

and

$$\mu_1(0, 1, 0) = \mu_1(0, 0, 1) = \mu_2(1, 0, 0) = \mu_2(0, 0, 1) = \mu_3(1, 0, 0) = \mu_3(0, 1, 0) \equiv b.$$

As a result, finding the optimal mechanism boils down to finding a , b and a tie-breaking rule T for when agents who assign points to themselves have the same highest score as agents who do not assign points to themselves. One can then verify that the set of (a, b, T) which works is any a and b such that $a = b - 2$ with a tie breaking rule of $\frac{14}{18}$ in favor of the agents who assign more points to themselves. Seeing as any such (a, b, T) leads to the same allocation, that allocation is optimal.

Had one not realized that there was a symmetric solution, part ii) of proposition 1 would have been required as it is not the case that any (μ, T) for which agents report truthfully given $Mech(\mu, T)$ is optimal. For example, suppose $\mu_1(s_1) = k$ for all $s_1 \in S_1$ and $\mu_2(s_2) = \mu_3(s_3) = 0$ for all $s_2 \in S_2$ and $s_3 \in S_3$. If k is sufficiently large, the allocation that would be truthfully implemented would be a constant allocation where alternative $j = 1$ would always be chosen, which is clearly not optimal. Part ii) states that one can only be certain of having found an optimal mechanism if μ satisfies condition (1).

4.2 How does the optimal mechanism help the DM?

To get a better sense of what the mechanism does to help the DM, let us go through the following exercise: say that agent 1 goes from observing (and truthfully reporting) s_1 to observing (and truthfully reporting) s'_1 , where $s_{1j} = s'_{1j}$ for all alternatives j except for one $j = j^*$ for which $s_{1j^*} < s'_{1j^*}$. The only difference between the two signals is that, in the latter, agent 1 believes there is more value for the DM in picking j^* . What changes in terms of which alternatives get chosen?

One possibility is that nothing changes, so that $x^*(s'_1, s_{-1}) = x^*(s_1, s_{-1})$ for all $s_{-1} \in S_{-1}$, where x^* represents the optimal allocation induced by the mechanism. Instead, let us consider the case where there is some $s_{-1} \in S_{-1}$ for which $x^*(s'_1, s_{-1}) \neq x^*(s_1, s_{-1})$. In this scenario, as I detail below, the probability that alternative j^* is

chosen increases. For each alternative j and each vector $s_{-1} \in S_{-1}$, let

$$\Delta_j(s_{-1}) \equiv \text{score}_j(s'_1, s_{-1}) - \text{score}_j(s_1, s_{-1}).$$

Using proposition 1, it follows that

$$\Delta_j(s_{-1}) = s'_{1j} - s_{1j} + (\mu_1(s'_1) - \mu_1(s_1)) u_{1j}$$

so that $\Delta_j(s_{-1})$ is actually independent of s_{-1} . There are different cases to consider depending on agent 1's preferences: $\mu_1(s'_1) = \mu_1(s_1)$, $\mu_1(s'_1) > \mu_1(s_1)$ and $\mu_1(s'_1) < \mu_1(s_1)$. Let us consider the first two, as the third case is analogous to the second one.

Case 1: $\mu_1(s'_1) = \mu_1(s_1)$

In this case, it follows that $\Delta_{j'} = \Delta_{j''}$ for any pair j', j'' such that $j', j'' \neq j^*$, i.e., the order of scores between alternatives other than j^* stay the same. However, $\Delta_{j^*} > \Delta_j$ for all $j \neq j^*$. As a result, giving points to alternative j^* has the simple effect of replacing some of the other alternatives by alternative j^* for some vectors $s_{-1} \in S_{-1}$. For agent 1 to be indifferent, it would have to be that some of the alternatives that are replaced are preferred by agent 1 to alternative j^* and some are not.

Case 2: $\mu_1(s'_1) > \mu_1(s_1)$

In this case, for any pair j', j'' such that $j', j'' \neq j^*$ it follows that $\Delta_{j'} > \Delta_{j''}$ if and only if $u_{1j'} > u_{1j''}$, i.e., the score of agent 1's preferred alternative increases more than his second most preferred alternative, which increases more than his third preferred alternative and so on. However, alternative j^* , which would be low in agent 1's preference order in order to generate $\mu_1(s'_1) > \mu_1(s_1)$, receives a boost of $s'_{1j} - s_{1j}$ that makes it go up that ranking. As a result, for some vectors $s_{-1} \in S_{-1}$, there will be some alternative j'' that gets replaced by some other alternative $j' \neq j^*$ such that $u_{1j''} > u_{1j'}$ (which is favorable to agent 1), while for some other vectors $s_{-1} \in S_{-1}$, some alternative j will be replaced by alternative j^* . In the latter case, there will certainly be some $s_{-1} \in S_{-1}$ such that the alternative j that is replaced by j^* is such that $u_{1j} > u_{1j^*}$ to preserve incentive compatibility. Overall, agent 1 is left indifferent by reporting s'_1 but increases the likelihood of alternative j^* being selected and that is what helps the DM; she ends up increasing the odds of selecting alternatives which receive more points.

For clarity, let us say that there are four alternatives and that $u_{i1} > u_{i2} > u_{i3} > u_{i4}$ with $j^* = 4$. The following table displays how each score changes when agent 1 goes

from reporting s_1 to reporting s'_1 :

j	α_j	β_j
1	\leftrightarrow	$\uparrow\uparrow\uparrow\uparrow$
2	\leftrightarrow	$\uparrow\uparrow\uparrow$
3	\leftrightarrow	$\uparrow\uparrow$
4	\uparrow	\uparrow

The only α_j that increases is α_4 , due to the additional points received from agent 1. However, because agent 1 assigns additional points to an alternative he dislikes, the β_j that increase more are those of the alternatives that agent 1 prefers. Overall, for some reports of the other agents, alternative 4 replaces alternatives 1, 2 and 3, which harms agent 1, but, for some other reports, alternative 1 replaces alternatives 2 and 3, and alternative 2 replaces alternative 3, which benefits agent 1. The combination of these two effects leave him indifferent.

4.3 Proof of proposition 1

For each agent i , let $s_i^* \in S_i$ denote some arbitrary type of agent i and let $\widehat{S}_i \equiv S_i \setminus \{s_i^*\}$. It follows that any allocation x^* is an optimal allocation if and only if it maximizes the DM's expected payoff

$$V(x) = \sum_{s \in S} P(s) \sum_{j=1}^J x_j(s) \sum_{i=1}^I s_{ij},$$

where

$$P(s) = \prod_{i=1}^I p_i(s_i),$$

subject to the incentive constraints,

$$\sum_{s_{-i} \in S_{-i}} P(s_i, s_{-i}) \sum_j (x_j(s_i, s_{-i}) - x_j(s_i^*, s_{-i})) u_{ij} \equiv IC_i(s_i) = 0$$

for all i and $s_i \in \widehat{S}_i$, the feasibility constraints,

$$P(s) \sum_j (x_j(s) - 1) = 0$$

for all $s \in S$, and the non-negativity constraints,

$$P(s) x_j(s) \geq 0$$

for all $s \in S$ and j .

By the Karush-Kuhn-Tucker theorem, an allocation x^* is optimal if and only if it satisfies all constraints and there is $(\lambda, \gamma, \delta)$, where $\lambda = (\lambda_1, \dots, \lambda_I)$, $\lambda_i : \widehat{S}_i \rightarrow \mathbb{R}$ for all i , $\gamma : S \rightarrow \mathbb{R}$, $\delta = (\delta_1, \dots, \delta_J)$ and $\delta_j : S \rightarrow \mathbb{R}_+$ for all j , such that

$$\frac{\partial \mathcal{L}(x^*, \lambda, \gamma, \delta)}{\partial x_j(s)} = 0 \text{ for all } j \text{ and } s \in S \quad (2)$$

and

$$\delta_j(s) x_j^*(s) = 0 \text{ for all } j \text{ and } s \in S, \quad (3)$$

where

$$\mathcal{L}(x, \lambda, \gamma, \delta) \equiv V(x) + \sum_{i=1}^I \sum_{s_i \in \widehat{S}_i} \lambda_i(s_i) IC_i(s_i) + \sum_{s \in S} P(s) \sum_{j=1}^J (\gamma(s) + \delta_j(s)) x_j(s).$$

4.3.1 Proof of part i)

Let allocation x^* be optimal and let $(\lambda^*, \gamma^*, \delta^*)$ be the corresponding Lagrange multipliers such that (2) and (3) hold. For each i , let $\widehat{\mu}_i : S_i \rightarrow \mathbb{R}$ be such that

$$\widehat{\mu}_i(s_i) = \begin{cases} \lambda_i^*(s_i) & \text{if } s_i \in \widehat{S}_i \\ - \sum_{\widehat{s}_i \in \widehat{S}_i} \frac{p_i(\widehat{s}_i)}{p_i(s_i^*)} \lambda_i^*(\widehat{s}_i) & \text{if } s_i = s_i^* \end{cases}$$

and let

$$score_j^*(s) = \alpha_j(s) + \sum_{i=1}^I \widehat{\mu}_i(s_i) u_{ij}$$

for any $s \in S$ and j .

Claim 1 For any $s \in S$ and j ,

$$score_j^*(s) < \max_{j' \neq j} score_{j'}^*(s) \Rightarrow x_j^*(s) = 0.$$

Proof. Notice that (2) can be written as

$$score_j^*(s) + \delta_j^*(s) + \gamma^*(s) = 0$$

for each $s \in S$ and j . I proceed by contradiction. Take any (j', j'') and any $s \in S$ such that $score_{j'}^*(s) > score_{j''}^*(s)$ but assume that $x_{j''}^*(s) > 0$. In that case, it follows that $\delta_{j''}^*(s) = 0$ by (3), which implies that

$$score_{j'}^*(s) + \delta_{j'}^*(s) = score_{j''}^*(s),$$

which is a contradiction, because $\delta_{j'}^*(s) \geq 0$. ■

If $\hat{\mu} = (\hat{\mu}_1, \dots, \hat{\mu}_I)$ is regular, then the statement follows trivially by setting $\mu = \hat{\mu}$ and by defining the following tie-breaking rule T :

$$T_j(s, \Omega) = \begin{cases} x_j^*(s) & \text{if } \Omega = \Omega^\mu(s) \\ \frac{\mathbf{1}\{j \in \Omega\}}{|\Omega|} & \text{if } \Omega \neq \Omega^\mu(s) \end{cases}$$

for all $s \in S$ and $\Omega \in \wp(J)$.¹⁰ In general, however, there could be multiple Lagrange multipliers associated with allocation x^* , so it is possible that not all of them lead to a regular $\hat{\mu}$. To overcome this issue, in the appendix, I describe an algorithm Γ that transforms $\hat{\mu}$ into a regular μ (i.e., $\Gamma(\hat{\mu}) = \mu$) such that, when one defines each alternative j 's score as

$$score_j(s) = \alpha_j(s) + \sum_{i=1}^I \mu_i(s_i) u_{ij},$$

it follows that, for any $s \in S$ and j ,

$$score_j(s) < \max_{j' \neq j} score_{j'}(s) \Rightarrow x_j^*(s) = 0.$$

4.3.2 Proof of part ii)

Take any (μ, T) such $Mech(\mu, T)$ truthfully induces allocation x^* . In order to show that x^* is optimal, it is enough to find Lagrange multipliers $(\lambda^*, \gamma^*, \delta^*)$ such that (2)

¹⁰Notice that any $\Omega \neq \Omega^\mu(s)$ is "off-the-path" in that set Ω contains alternatives that do not have the highest score. As a result, for all such Ω , $T(s, \Omega)$ can be defined arbitrarily.

and (3) hold. It can be verified that the following multipliers satisfy both conditions, provided condition (1) holds:

$$\lambda_i^*(s_i) = \mu_i(s_i)$$

for all $s_i \in \widehat{S}_i$ and for all i ,

$$\gamma^*(s) = -\max_j \text{score}_j(s)$$

for all $s \in S$ and

$$\delta_j^*(s) = -\gamma^*(s) - \text{score}_j(s)$$

for all $s \in S$ and j .

4.4 Two alternatives

In the optimal mechanism, the DM is able to incentivize each agent to share his private information by promising to leave him indifferent regardless of what he reports. When an agent reports that a certain alternative he does not like has more value for the DM, the DM is able to increase the likelihood that said alternative is chosen but compensates the agent by also increasing the likelihood that some of the alternatives the agent does like are chosen. When there are only two alternatives, this logic fails, because providing information in support of one's least favorite alternative directly implies that one's preferred alternative is selected less often. As a result, the DM does not benefit from interacting with the agents.

Proposition 2 *If $J = 2$ and agents have strict preferences, there is an optimal allocation x^* that is independent of $s \in S$.*

Proof. For each agent i , let

$$k_i \equiv \sum_{s_i \in S_i} p_i(s_i) s_{i1} - \sum_{s_i \in S_i} p_i(s_i) s_{i2}$$

and let μ be as follows: for each agent i ,

$$\mu_i(s_i) = \frac{k_i}{u_{i1} - u_{i2}} - \frac{s_{i1} - s_{i2}}{u_{i1} - u_{i2}}$$

for each $s_i \in S_i$. Notice that each $\mu_i(s_i)$ is well defined because $u_{i1} \neq u_{i2}$ for all i (agents have strict preferences). By construction, μ satisfies (1) and is ordered. Let

the tie-breaking rule T be such that

$$T_1(s, \{1, 2\}) = T_2(s, \{1, 2\}) = \frac{1}{2}$$

for all $s \in S$. Finally, notice that

$$\begin{aligned} score_1(s) - score_2(s) &= \sum_{i=1}^I (s_{i1} - s_{i2} + \mu_i(s_i)(u_{i1} - u_{i2})) = \sum_{i=1}^I k_i \\ &= \sum_{\hat{s} \in S} P(\hat{s}) \alpha_1(\hat{s}) - \sum_{\hat{s} \in S} P(\hat{s}) \alpha_2(\hat{s}), \end{aligned}$$

which is independent of s . As a result, allocation x^* is truthfully implemented by mechanism $Mech(\mu, T)$, where

$$x_1^*(s) = \begin{cases} 1 & \text{if } \sum_{\hat{s} \in S} P(\hat{s}) \alpha_1(\hat{s}) > \sum_{\hat{s} \in S} P(\hat{s}) \alpha_2(\hat{s}) \\ \frac{1}{2} & \text{if } \sum_{\hat{s} \in S} P(\hat{s}) \alpha_1(\hat{s}) = \sum_{\hat{s} \in S} P(\hat{s}) \alpha_2(\hat{s}) \\ 0 & \text{if } \sum_{\hat{s} \in S} P(\hat{s}) \alpha_1(\hat{s}) < \sum_{\hat{s} \in S} P(\hat{s}) \alpha_2(\hat{s}) \end{cases}$$

for all $s \in S$. Therefore, by proposition 1, part ii), allocation x^* is optimal, despite being independent of s . ■

5 Contests

In many applications, the DM (real or abstract) is asked to choose one of I self-interested contestants, who are indifferent as to whom should be chosen should it not be them. A lot of times, those contestants have information over which contestants provide more value to the DM. Contests are then just a special case of the general model described in section 3, so that the optimal mechanism can be found using proposition 1. Given the special characteristics of contests, it is possible to further refine the optimal mechanism.

For each $\tau = (\tau_1, \dots, \tau_I)$ such that $\tau_i : S_i \rightarrow \mathbb{R}$ and tie-breaking rule $T : S \times \wp(J) \rightarrow [0, 1]^J$, define mechanism $Mech^c(\tau, T)$ as follows: Each agent i simultaneously assigns points $s_{ij} \in S_{ij}$ to each alternative $j \neq i$. Vector s determines each alternative j 's score

as follows:

$$score_j^c(s) = \sum_{i \neq j} s_{ij} + \tau_j(s_j).$$

As before, the alternative with the largest score is chosen and ties are broken according to T . Vector τ is called *external* if each τ_i is independent of s_{ii} .

This new mechanism is very similar to the general mechanism of the previous section except that each agent's score does not depend on the points he assigns himself, provided that τ is external.

Proposition 3 *i) For any optimal allocation x^* , there is an external $\tau : S \rightarrow \mathbb{R}^I$ such that each τ_i is weakly increasing with s_{ij} and a tie breaking rule $T : S \times \varphi(J) \rightarrow [0, 1]^J$ such that $Mech^c(\mu, T)$ truthfully induces x^* .*

ii) If there is some external $\tau : S \rightarrow \mathbb{R}^I$ such that

$$\sum_{s_i \in S_i} p_i(s_i) \tau_i(s_i) = \sum_{s_i \in S_i} p_i(s_i) s_{ii} \quad (4)$$

for all i and a tie-breaking rule $T : S \times \varphi(J) \rightarrow [0, 1]^J$ for which $Mech^c(\tau, T)$ truthfully induces some allocation x^ , then allocation x^* is optimal.*

Proof. For part i), it follows by part i) of proposition 1 that, for any optimal allocation x^* , there is some (μ, T) such that x^* is truthfully induced by $Mech(\mu, T)$ and μ is ordered. In mechanism $Mech(\mu, T)$, agent i 's report of s_{ii} only affects $score_i$. As a result, if, for each agent i , one arbitrarily defines some $s_{ii}^* \in S_{ii}$, it follows that, for any $s \in S$,

$$score_i(s) < \max_{j \neq i} score_j(s) \Rightarrow score_i(s_{ii}^*, s_{i,-i}, s_{-i}) \leq \max_{j \neq i} score_j(s_{ii}, s_{i,-i}, s_{-i}).$$

In words, it must be that agent i 's report of s_{ii} does not matter in determining which scores are the highest; if it did, agent i would always report whatever s_{ii} increased his score the most. Define

$$\tau_i(s_i) \equiv s_{ii}^* + \mu_i(s_{ii}^*, s_{i,-i})$$

for all $s_i \in S_i$ and notice that, by construction, τ is external, for each i , τ_i is (weakly) increasing with s_{ij} for all $j \neq i$ because μ is ordered and allocation x^* is truthfully induced by $Mech^c(\tau, T)$.

As for part ii), suppose that $Mech^c(\tau, T)$ truthfully induces some allocation x^* and

let μ be such that

$$\mu_i(s_i) \equiv \tau_i(s_i) - s_{ii}$$

for all i and for all $s_i \in S_i$. Because, by construction,

$$\sum_{i \neq j} s_{ij} + \tau_j(s_j) = \sum_{i=1}^I s_{ij} + \mu_j(s_j)$$

and

$$\sum_{s_i \in S_i} p_i(s_i) \mu_i(s_i) = 0,$$

the result follows by proposition 1, part b). ■

Part i) states that one can always implement an optimal allocation by assigning a score to each agent which only depends on the amount of points other agents have assigned him. In this way, each agent's assessment of himself - the s_{ii} of each agent i - may only enter the mechanism through its impact on the tie-breaking rule. Nevertheless, as I show in the appendix, the probability that there is a tie can be made arbitrarily small as the agents' types become quasi-continuous. Part ii) states the converse: if there is a mechanism where agents do not assign points to themselves that truthfully induces some allocation, that allocation is optimal, provided condition (4) holds.

5.1 Ordinal mechanisms under symmetry

Ordinal mechanisms are mechanisms that only use the information over how agents are ordered in terms of their value for the DM. They are appealing because they require less of each agent (Bogomolnaia and Moulin, 2001; Carroll, 2018). Each agent is not required to know precisely how much better to the DM some alternatives are compared to others; he is simply required to be able to rank them according to how much value for the DM they provide.

Specifically, for this part of the paper, I make two additional assumptions.

Assumption A: For all i , $p_i(s_i) = 0$ for all $s_i \in S_i$ such that there is j, j' for which $s_{ij} = s_{ij'}$. In words, I assume that each agent i always believes that either agent j or agent j' provide more value to the DM (there are no ties).

Assumption B: I assume symmetry: S_{ij} is the same for all i, j and each s_i has the same symmetric distribution.

Assumption B is convenient because it implies that there are optimal allocations (ordinal and otherwise) that are symmetric (one can always generate a symmetric optimal allocation by finding a non-symmetric optimal allocation and randomizing over each agent's identity). That leads to the simpler mechanism described below. Nevertheless, for non-symmetric settings, the optimal ordinal allocation can be found using proposition 3.

Additional notation: For each agent i and each $s_i \in S_i$, let vector $r_i(s_i) = (r_{i1}(s_i), r_{i2}(s_i), \dots, r_{iI}(s_i))$, where each $r_{ij}(s_i) \in \{1, \dots, I\}$ represents agent j 's ranking according to agent i . So, for example, $r_{ij}(s_i) = 1$ if s_{ij} is the largest element of vector s_i , $r_{ij}(s_i) = 2$ if s_{ij} is the second largest and so on. Let also $r(s) = (r_1(s_1), r_2(s_2), \dots, r_I(s_I))$. Let $\kappa : \{1, \dots, I\} \rightarrow \mathbb{R}$ be such that

$$\kappa(a) = E(s_{ij} | r_{ij}(s_i) = a)$$

for all $a \in \{1, \dots, I\}$; $\kappa(a)$ represents the expected value added for the DM of selecting agent j , given that agent i ranks agent j in the a^{th} position. It then follows that κ is increasing.

Additional definitions: An allocation x is ordinal if and only if $x(s) = x(s')$ for all $s, s' \in S$ such that $r(s) = r(s')$. A tie-breaking rule $T : S \times \wp(J) \rightarrow [0, 1]^J$ is ordinal if and only if $T(s, \Omega) = T(s', \Omega)$ for all $\Omega \in \wp(J)$ and $s, s' \in S$ such that $r(s) = r(s')$.

For each $\eta : \{1, \dots, I\} \rightarrow \mathbb{R}$ and ordinal tie-breaking rule $T : S \times \wp(J) \rightarrow [0, 1]^J$, define mechanism $Mech^o(\eta, T)$ as follows: Each agent i simultaneously ranks each alternative j (including himself) by reporting r_i . Vector r determines each alternative j 's score as follows:

$$score_j^o(r) = \sum_{i \neq j} \kappa(r_{ij}) - \eta(r_{jj}).$$

The alternative with the largest score is chosen with ties being resolved using T . If $Mech^o(\eta, T)$ has a Bayes-Nash equilibrium where agents report truthfully ($r_{ij} = r_{ij}(s_i)$ for all i and j), the allocation that is generated by that equilibrium is truthfully induced by $Mech^o(\eta, T)$.

Proposition 4 *i) If allocation x^* is an optimal symmetric ordinal allocation, there is a weakly decreasing function $\eta : \{1, \dots, I\} \rightarrow \mathbb{R}$ and an ordinal tie breaking rule $T : S \times \wp(J) \rightarrow [0, 1]^J$ such that $\text{Mech}^o(\eta, T)$ truthfully induces it.*

ii) If there is some function $\eta : \{1, \dots, I\} \rightarrow \mathbb{R}$ such that

$$\frac{1}{I} \sum_{r_{ii} \in \{1, \dots, I\}} \eta(r_{ii}) = \frac{1}{I} \sum_{r_{ii} \in \{1, \dots, I\}} \kappa(r_{ii}). \quad (5)$$

and some ordinal tie breaking rule $T : S \times \wp(J) \rightarrow [0, 1]^J$ such that $\text{Mech}^o(\eta, T)$ truthfully induces some allocation x^ , then allocation x^* is an optimal ordinal allocation.*

Proof. Consider the following information structure:

$$\widehat{S}_{ij} = \{\kappa(1), \kappa(2), \dots, \kappa(I)\}$$

for any i, j with \widehat{p}_i being such that i) $\widehat{p}_i(\widehat{s}_i) = 0$ for any s_i for which $\widehat{s}_{ij} = \widehat{s}_{ij'}$ for some pair (j, j') and ii) $\widehat{p}_i(\widehat{s}_i) = \frac{1}{I}$ for any other $\widehat{s}_i \in \widehat{S}_i$. By construction, a symmetric ordinal allocation is optimal if and only if it is an optimal allocation given the new information structure. As a result, by proposition 3, part i), it follows that, for each optimal symmetric ordinal allocation x^* , there is an external $\tau : \widehat{S} \rightarrow \mathbb{R}^I$ and a tie breaking rule $T : \widehat{S} \times \wp(J) \rightarrow [0, 1]^J$ such that $\text{Mech}^c(\tau, T)$ truthfully induces x^* . For each agent i , let

$$\widetilde{S}_i \equiv \left\{ \widehat{s}_i \in \widehat{S}_i : \widehat{p}_i(\widehat{s}_i) > 0 \right\}$$

and notice that the symmetry of x^* implies that there is some (τ, T) for which $\text{Mech}^c(\tau, T)$ truthfully induces x^* such that τ_i is the same for all i and is such that $\tau_i(\widehat{s}_i) = \tau_i(\widehat{s}'_i)$ for all $\widehat{s}_i, \widehat{s}'_i \in \widetilde{S}_i$ for which $\widehat{s}_{ii} = \widehat{s}'_{ii}$. Moreover, because τ_i is weakly increasing with each \widehat{s}_{ij} , it follows that $\tau_i(\widehat{s}_i) \geq \tau_i(\widehat{s}'_i)$ for all $\widehat{s}_i, \widehat{s}'_i \in \widetilde{S}_i$ such that $\widehat{s}_{ii} \leq \widehat{s}'_{ii}$. As result, part i) follows by defining $\eta(r_{ii}) = -\tau_i(\widehat{s}_i)$ for any $\widehat{s}_i \in \widetilde{S}_i$ such that $\widehat{s}_{ii} = r_{ii}$.

As for part ii), suppose there is some $\eta : \{1, \dots, I\} \rightarrow \mathbb{R}$ such that (5) holds and some tie breaking rule $T : S \times \wp(J) \rightarrow [0, 1]^J$ such that $\text{Mech}^o(\eta, T)$ truthfully induces some allocation x^* . Let $\tau : \widehat{S} \rightarrow \mathbb{R}^I$ be such that $\tau_i(\widehat{s}_i) = -\eta(\widehat{s}_{ii})$ for any $\widehat{s}_i \in \widetilde{S}_i$ and $\tau_i(\widehat{s}_i) = M$ for any $\widehat{s}_i \notin \widetilde{S}_i$ where $M \in \mathbb{R}$ is sufficiently small to ensure that mechanism $\text{Mech}^c(\tau, T)$ truthfully induces allocation x^* given the new information structure. Notice that condition (5) implies that

$$\sum_{\widehat{s}_i \in \widetilde{S}_i} \widehat{p}_i(\widehat{s}_i) \tau_i(\widehat{s}_i) = \sum_{\widehat{s}_i \in \widetilde{S}_i} \widehat{p}_i(\widehat{s}_i) \widehat{s}_{ii}.$$

As a result, it follows by part ii) of proposition 3, that allocation x^* is optimal given the new information structure. ■

Notice that, in the optimal mechanism described, each agent is asked to also rank himself and not just others as it may have been expected given proposition 3, where I show that agents do not report their own value. The way to reconcile the two results is by noticing that, when an agent ranks himself, he is also providing information about the other agents' values. For example, if an agent ranks himself first, the information that is passed on to the DM is that the other agents' values are not that high. Therefore, an agent ranking himself is useful not because of the information he provides about his own actual value but about the value of others. Naturally, to make sure agents have an incentive to report truthfully, ranking oneself high lowers one's score because it also lowers everybody else's score (compared to ranking oneself lower).

By nature of its information structure, the optimal mechanism of the example is an ordinal mechanism like the one described. Recall that in the example, each agent simply ranks one of the agents as the best costume and all other costumes are tied in second place. Therefore, each agent's score is the sum of an increasing function of how others rank him (the sum of votes received by others), plus a decreasing function of how the agent ranks himself (whether the agent votes for somebody else or not).

6 Conclusion

6.1 On the complexity of the optimal mechanism

Part of the motivation of discussing ordinal mechanisms in section 5.1. is that the optimal ordinal mechanism (for contests and under symmetry) are easier to understand by the agents than the general optimal mechanism. In ordinal mechanisms, agents simply have to rank alternatives and trust that a truthful report does not harm them. Nevertheless, one could argue that the optimal ordinal mechanism is still more complicated than the adjusted majority rule, where agents simply vote for their favorite alternative and are not allowed to vote for themselves. A middle ground solution would be the following mechanism.

An alternative simple mechanism: Like the adjusted majority rule, have each agent vote for one of the other agents, but allow each agent to choose to make their vote a "strong" vote. So, each agent can choose whom to vote for and whether their vote

is strong or weak. Compared to a weak vote, a strong vote by some agent i on some agent j increases both agent j and agent i 's score (so that agent i is indifferent between weak and strong voting). The reason why this alternative mechanism improves on the adjusted majority rule is as follows. In the adjusted majority rule, when agent i votes for agent j , he communicates to the DM that agent j provides more value than any other agent different than i . However, in this new proposed mechanism, in addition to reporting that agent j is better than any other agent different than i , agent i can also say whether agent j is better than agent i or not; he gives a strong vote in the former case and a weak vote in the latter case. In that way, the new mechanism takes into account that if an agent receives a weak vote, their value is not as high as if they were to receive a strong vote, so their score does not increase as much. Naturally, in order for agents to have incentives to give strong votes, their own score must also increase.

Lastly, it is important to note that in none of these mechanisms - the optimal mechanisms, the adjusted majority rule and the alternative simple mechanism - have a dominant strategy; indeed, reporting truthfully is optimal for each agent provided others do so as well. For that reason, for each of these mechanisms to work, each agent must i) trust that it is best for him to report truthfully when others report truthfully and ii) trust that everybody else will report truthfully. This is in contrast with the delegation mechanism, where reporting truthfully is a dominant strategy for the agents.

6.2 On the requirement that agents be indifferent

One feature of the optimal mechanism is that agents are always indifferent as to what to do. The challenge of the DM is precisely to design a mechanism that elicits as much information as possible from the agents while keeping them indifferent. There are various models and papers where the optimal mechanism is such that some agents for some types are indifferent. However, one can usually build an alternative mechanism where agents are not indifferent that works almost as well. That is not possible in the model I consider; in fact, it is not possible in any delegation setting with a single agent either. Indeed, as I discuss in section 2, because the agents' preferences are public and types are independent, *any* incentive compatible mechanism must leave every agent indifferent between what to report; not just optimal mechanisms. In that sense, while the reader might have concerns over whether agents will play as specified in the optimal mechanisms when indifferent, those concerns extend to any mechanism. What the optimal mechanisms presented in the paper have in their favor

is that they induce truthful reporting. That is, when indifferent, agents are supposed to assign points and/or rank alternatives according to their perceived value of each alternative. Reporting truthfully seems like a good default as most people experience some discomfort when lying or deceiving (Gneezy, 2005).

6.3 On the likelihood that there are multiple highest scores

In general, in the optimal mechanism, there might be multiple alternatives with the highest score. When that happens, it might be that the tie-breaking rule is not even; a feature that might be undesirable or unrealistic in symmetric contests. Nevertheless, that feature is largely non-existent in richer information structures. Specifically, I show in the appendix that if each agent's type is quasi-continuous (basically, if the set of types is large and the probability of each type is small), the probability that there is a tie is negligible. This does not mean that ties can then be assumed to be broken evenly in and of itself, because ties would still have a positive, albeit (arbitrarily) small, probability. However, relying again on the idea that agents have an intrinsic preference for reporting truthfully, one can introduce the notion of an ε -*equilibrium* (Fudenberg and Levine, 1988). In a truthful ε -*equilibrium*, agents report truthfully provided that no deviation returns an expected payoff larger than $\varepsilon > 0$. Equipped with this new equilibrium concept, it then follows that, provided types are quasi-continuous, one can take any optimal mechanism and then impose that ties are broken evenly. This only slightly alters the agents' incentives but not enough to make them change their (truthful) report.

7 Appendix

7.1 Proof of proposition 1 (remaining steps)

I complete the proof of part i) by describing an algorithm Γ such that $\Gamma(\hat{\mu}) \equiv \mu$ is regular and such that

$$score_j(s) < \max_{j' \neq j} score_{j'}(s) \Rightarrow x_j^*(s) = 0,$$

for all $s \in S$ and j .

For each agent i and type $s_i \in S_i$, let

$$\Sigma_i(s_i) \equiv \{s'_i \in S_i : s'_i \succ s_i\}$$

and notice that, if $s'_i \in \Sigma_i(s''_i)$, then $s''_i \notin \Sigma_i(s'_i)$. Let $\bar{c}_i \equiv |S_i|$ and recursively define set S_i^c as follows:

$$S_i^0 = \{s_i \in S_i : \Sigma_i(s_i) = \emptyset\}$$

and

$$S_i^c = \left\{ s_i \in S_i \setminus \bigcup_{c'=0}^{c-1} S_i^{c'} : \Sigma_i(s_i) \subset \bigcup_{c'=0}^{c-1} S_i^{c'} \right\}$$

for $c = 1, \dots, \bar{c}_i$, where $\bar{c}_i \equiv |S_i|$. Finally, define μ_i recursively as follows:

$$\forall s_i \in S_i^0, \mu_i(s_i) = \widehat{\mu}_i(s_i)$$

and

$$\forall s_i \in S_i^c, \mu_i(s_i) = \max \left\{ \widehat{\mu}_i(s_i), \max_{s'_i \in \Sigma_i(s_i)} \mu_i(s'_i) \right\}$$

for $c = 1, \dots, \bar{c}_i$. Notice that, by construction, $\mu = (\mu_1, \dots, \mu_I)$ is regular.

Claim 2 For any $s \in S$, for all j ,

$$\text{score}_j(s) < \max_{j' \neq j} \text{score}_{j'}(s) \Rightarrow x_j^*(s) = 0. \quad (6)$$

Proof. The proof is by induction. For all $k = 0, 1, \dots, I$, and for all $s \in S$ and j , let

$$\text{score}_j^k(s) = \sum_{i=1}^I s_{ij} + \sum_{i=1}^I \varpi_i^k(s_i) u_{ij},$$

where

$$\varpi_i^k(s_i) = \begin{cases} \mu_i(s_i) & \text{if } i \leq k \\ \widehat{\mu}_i(s_i) & \text{if } i > k \end{cases}.$$

Essentially, each step k transforms λ into $\widehat{\mu}$ one agent at a time; in particular,

$$\text{score}_j^*(s) = \text{score}_j^0(s) \text{ and } \text{score}_j(s) = \text{score}_j^I(s).$$

For each step k and each $s \in S$, let

$$\Omega^k(s) \equiv \left\{ j : \text{score}_j^k(s) \geq \max_{j' \neq j} \text{score}_{j'}^k(s) \right\},$$

and, for each agent i ,

$$\bar{u}_i^k(s) = \max_{j \in \Omega^k(s)} u_{ij} \text{ and } \underline{u}_i^k(s) = \min_{j \in \Omega^k(s)} u_{ij}.$$

The claim can then be restated as follows: for any $s \in S$, for all j such that $x_j^*(s) > 0$, then $j \in \Omega^k(s)$. Notice that I have already established this for $k = 0$. Therefore, it is enough to prove for all $k = 1, \dots, I$ that if $x_j^*(s) > 0 \Rightarrow j \in \Omega^{k-1}(s)$ then $x_j^*(s) > 0 \Rightarrow j \in \Omega^k(s)$ for all $s \in S$ and j .

Suppose the statement is false, so that there is some step k , $s \in S$ and some alternative j such that $x_j^*(s) > 0$ but where $j \notin \Omega^k(s)$. By the induction hypothesis, it follows that $j \in \Omega^{k-1}(s)$. This implies that $\underline{u}_k^k(s) > u_{ij}$, while, by construction, $\underline{u}_k^k(s) \geq \bar{u}_k^{k-1}(s)$. On the other hand, the fact that $\Omega^k(s) \neq \Omega^{k-1}(s)$ implies that there is some $s'_k \in S_k$ such that $s'_k \succ_k s_k$ and $\mu_k(s'_k) = \mu_k(s_k) = \hat{\mu}_k(s'_k) > \hat{\mu}_k(s_k)$. Notice that, by construction, $\underline{u}_k^{k-1}(s'_k, s'_{-k}) \geq \bar{u}_k^{k-1}(s_k, s'_{-k})$ for all $s'_{-k} \in S_{-k}$. Furthermore, $\underline{u}_k^{k-1}(s'_k, s_{-k}) = \underline{u}_k^k(s'_k, s_{-k}) \geq \bar{u}_k^k(s) \geq \underline{u}_k^k(s) > u_{ij}$. As a result, given the induction hypothesis, it follows that

$$E_{s'_{-k}}(x^*(s'_k, s'_{-k})) > E_{s'_{-k}}(x^*(s_k, s'_{-k})),$$

which is a contradiction, because x^* is incentive compatible. ■

7.2 On the likelihood that there are multiple highest scores

In the text, it is mentioned that the likelihood that there are multiple alternative with the same highest score can be made arbitrarily small if types becomes quasi-continuous. In this section, I make that statement precise. Consider a tournament setting and assume that $I \geq 3$.

Proposition 5 *Consider any optimal allocation x^* and any mechanism $\text{Mech}^c(\tau, T)$ that truthfully implements it. Let $\delta > 0$ be such that, for all $k \in \mathbb{R}$,*

$$\Pr \{s_{ij} - s_{ij'} = k\} < \delta$$

for all i, j and $j' \neq j$. It follows that

$$\Pr \{score_j^c(s) = score_{j'}^c(s)\} < \delta$$

for all j and $j' \neq j$.

Proof. Fix any pair j, j' . I show that

$$\Pr \{score_j^c(s) = score_{j'}^c(s)\} < \delta.$$

Let i^* be some agent $i \neq j, j'$. Notice that

$$\begin{aligned} & \Pr \{score_j^c(s) = score_{j'}^c(s)\} \\ = & \sum_{s_{-i^*}} \Pr \{s_{-i^*}\} \Pr \{score_j^c(s_{i^*}, s_{-i^*}) = score_{j'}^c(s_{i^*}, s_{-i^*}) | s_{-i^*}\}. \end{aligned}$$

Furthermore, for any s such that $score_j(s) = score_{j'}(s)$ there is some $a : S_{-i^*} \rightarrow \mathbb{R}$ such that

$$s_{i^*j} - s_{i^*j'} = a(s_{-i^*}).$$

As a result,

$$\Pr \{score_j^c(s_{i^*}, s_{-i^*}) = score_{j'}^c(s_{i^*}, s_{-i^*}) | s_{-i^*}\} < \delta$$

for all $s_{-i^*} \in S_{-i^*}$. As a result,

$$\Pr \{score_j^c(s) = score_{j'}^c(s)\} < \sum_{s_{-i^*}} \Pr \{s_{-i^*}\} \delta < \delta.$$

■

References

- [1] Alonso, R., Brocas, I., & Carrillo, J. D. (2014). Resource allocation in the brain. *The Review of Economic Studies*, 81(2), 501-534.
- [2] Alonso, R., & Matouschek, N. (2008). Optimal delegation. *The Review of Economic Studies*, 75(1), 259-293.
- [3] Amador, M., & Bagwell, K. (2013). The theory of optimal delegation with an application to tariff caps. *Econometrica*, 81(4), 1541-1599.
- [4] Armstrong, M. (1995). Delegation and discretion. *Unpublished manuscript*.
- [5] Azrieli, Y., & Kim, S. (2014). Pareto efficiency and weighted majority rules. *International Economic Review*, 55(4), 1067-1088.
- [6] Bloch, F., Dutta, B. & Dziubinski, M. (2021). Selecting a winner with external referees. Working paper.
- [7] Bogomolnaia, A., & Moulin, H. (2001). A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2), 295-328.
- [8] Börgers, T., & Postl, P. (2009). Efficient compromising. *Journal of Economic Theory*, 144(5), 2057-2076.
- [9] Branco, F. (1996). Common value auctions with independent types. *Economic design*, 2(1), 283-309.
- [10] Carroll, G. (2018). On mechanisms eliciting ordinal preferences. *Theoretical Economics*, 13(3), 1275-1318.
- [11] de Clippel, G., Eliaz, K., Fershtman, D., & Rozen, K. (2021). On selecting the right agent. *Theoretical Economics*, 16(2), 381-402.
- [12] Fudenberg, D., & Levine, D. K. (1988). Open-loop and closed-loop equilibria in dynamic games with many players. *Journal of Economic Theory*, 44(1), 1-18.
- [13] Gan, T., Hu, J., & Weng, X. (2021). Optimal contingent delegation. Working paper.
- [14] Gibbard, A. (1973). Manipulation of voting schemes: a general result. *Econometrica*, 587-601.

- [15] Gibbons, R. (1988). Learning in equilibrium models of arbitration. *American Economic Review*, 78 (5), 896-912.
- [16] Gneezy, U. (2005). Deception: The role of consequences. *American Economic Review*, 95(1), 384-394.
- [17] Goldlücke, S., & Tröger, T. (2018). Assigning an unpleasant task without payment. Working paper.
- [18] Holmstrom, B. (1984). On the theory of delegation. Boyer, and R. Kihlstrom (eds.) *Bayesian Models in Economic Theory* (New York: North-Holland) 115-141.
- [19] Kim, S. (2017). Ordinal versus cardinal voting rules: A mechanism design approach. *Games and Economic Behavior*, 104, 350-371.
- [20] Koessler, F., & Martimort, D. (2012). Optimal delegation with multi-dimensional decisions. *Journal of Economic Theory*, 147(5), 1850-1881.
- [21] Kováč, E., & Mylovanov, T. (2009). Stochastic mechanisms in settings without monetary transfers: The regular case. *Journal of Economic Theory*, 144(4), 1373-1395.
- [22] Majumdar, D., & Sen, A. (2004). Ordinarily Bayesian incentive compatible voting rules. *Econometrica*, 72(2), 523-540.
- [23] Martimort, D., & Semenov, A. (2006). Continuity in mechanism design without transfers. *Economics Letters*, 93(2), 182-189.
- [24] Maskin, E., & Riley, J. (1984). Monopoly with incomplete information. *The RAND Journal of Economics*, 15(2), 171-196.
- [25] Melumad, N. D., & Shibano, T. (1991). Communication in settings with no transfers. *The RAND Journal of Economics*, 173-198.
- [26] Miralles, A. (2012). Cardinal Bayesian allocation mechanisms without transfers. *Journal of Economic Theory*, 147(1), 179-206.
- [27] Myerson, R. B. (1979). Incentive compatibility and the bargaining problem. *Econometrica*, 61-73.
- [28] Myerson, R. B. (1981). Optimal auction design. *Mathematics of operations research*, 6(1), 58-73.

- [29] Myerson, R. B. (2002). Comparison of scoring rules in Poisson voting games. *Journal of Economic Theory*, 103(1), 219-251.
- [30] Myerson, R. B., & Satterthwaite, M. A. (1983). Efficient mechanisms for bilateral trading. *Journal of Economic theory*, 29(2), 265-281.
- [31] Mylovanov, T., & Zapechelnuk, A. (2013). Optimal arbitration. *International Economic Review*, 54(3), 769-785.
- [32] Pereyra, J. & Silva, F. (2021). Optimal assignment mechanisms with imperfect verification. Working paper.
- [33] Satterthwaite, M. A. (1975). Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of economic theory*, 10(2), 187-217.
- [34] Schmitz, P. W., & Tröger, T. (2012). The (sub-) optimality of the majority rule. *Games and Economic Behavior*, 74(2), 651-665.