An Overview of Social Choice and Mechanism Design

Tao Wang

April 16, 2020

1 Arrow Impossibility Theorem

Arrow Impossibility THM tells a sad story. Whenever there are more than three outcomes to be compared, a social aggregation of individual preferences over these outcomes necessarily replicate single member's preference. The lucky guy is called a dictator. But his fortune is the society's misfortune.

A few conditions are imposed over the preferences. First, preference is rational in that it is complete and transitive. Second, it is Paretian. If everyone likes apple over the pear, society as a whole should like apples over pears. Third, independent irrelevant alternatives(IIA) states that if the individual preferences over two outcomes stay the same, the social preference over the two should stay the same no matter what changes happen to preferences on other outcomes. Third, the space of outcomes is unrestricted.

2 A Paretian Social Choice Function

Thus far, the process through which individuals' preference aggregate takes an abstract form of a social choice function. It is a mapping from a profile of each individuals' type in society to a certain social outcome. One of the desirable features of a social choice function is ex-post efficient or Paretian. It implies that given the profile types, no other outcome is Pareto improvement compared to the outcome assigned by the social choice function. Or put it another way, for all type profiles, there is to an alternative to the social choice function that makes some individual strictly better off without hurting the rest.

3 Implementation of a Social Choice Function

3.1 Revealation Principle

However, things change when the individual type is private information. There are cases in which individuals find it optimal to misrepresent their own type. Therefore, the social choice function cannot be directly implemented based on

true types. Instead, one has to rely on indirect mechanisms that map the message sent by individuals to social outcomes through certain rules. The word implementation means the mechanisms induce equilibrium outcome the same to that directly implemented through the social choice function. In general, we don't know if a certain social choice function can be implemented. But the revelation principle assures us that if a social choice function can be implemented through a mechanism, then it can be directly implemented. So we need to only focus on the direct implementation of a social choice function.

3.2 Dominant Strategy Implementation

If we as a planner or a mechanism designer want to implement a social choice function in dominant strategy, that is if we make the collective allocation based on this social choice function, everyone finds it a dominant strategy to be him/herself instead of pretending someone else, should we feel confident about achieving the goal? Unfortunately, not so. Given an unrestricted outcome space and more than three outcomes, a social choice function is implemented in a dominant strategy if and only if it is dictatorial. This echoes the spirit of Impossibility THM. (Muller-Satterthwaite)

What does this exactly mean? A dictatorial social choice function simply maximizes the utility of an individual member in any case. No matter what others' type and this individual's type are, the social choice function always gives the dictator the best result he desires.

Some corollaries. Dominant strategy implementation implies that the social choice function is Paretian, which satisfies IIA and monotonic. The first two conditions together with Arrow Impossibility THM lead to dictatorial. What does monotonic mean for a social choice function? A monotonic social choice function is one such that, given a type profile, to each individual, if the social outcome is no better compared to her being a different type, then the outcome obtained based on the alternative profile is the same to the current one.

It turns out that there are very restrictive conditions required for social choice functions that can be implemented in a dominant strategy. For instance, in the second pricing auction, the equilibrium bid being the private value is a weakly dominant strategy. What general properties does this social choice function that is implicitly implemented by second pricing auction have?

3.3 2nd-Price Auction as an Example of Dominant Implementation

Consider a specific context with an indicator function that maps profile of types into a yes or no decision to undertake a social project. Together with this decision rule and a function of transfer constitute the social choice function in this context. The utility takes a quasi-linear form that consists of a benefit term proportional to the private valuation conditional on the project begin undertaken and a transfer term. No external financing. It is ex-post efficient. And if the transfer by an individual is equal to the total benefit of the project

to rest of the society plus a function independent from the individual's type. Or put it in another way, if the transfer by each one is equal to the benefit to the rest of the society solely through the project choice. Then this mechanism can be implemented in a dominant strategy.

Let's check the social choice function implicitly implemented by the second pricing auction meet these conditions. First, it is ex-post efficient. Aggregate welfare cannot be Pareto improved as the agent with the highest valuation end up getting the project, while transfers across agents have no aggregate welfare implications. Second, the decision rule always assigns the object the highest value agent. Third, for the non-pivotal bidder, the transfer is zero while the benefit to the rest of the society with or without her is equal to the highest valuation, thus the difference being zero. Forth, for the pivotal bidder, the payment is the second high valuation. The rest of the social benefit is zero when she is bidding while the second-highest valuation when she is not bidding. The tax is exactly equal to the externality caused by her reporting through the mechanism.

3.4 Bayes-Nash Implementation

It is too stringent to expect a social choice to be implemented in such a way that everyone is always happy to take the desired action no matter what the rest of society does. Instead, we may count on them to take the desired action as an average best response to the rest of the players. This is Bayesian Nash equilibrium implementation. That is to say, the mechanism induces equilibrium actions that each individual can only do their best in expectation and the equilibrium leads to the social choice. Our main question, again, is what kind of social choice functions could be implemented in this manner? According to the revelation principle, we can focus on direct implementation in B-N equilibrium. So we ask, what kind of conditions are necessary for a social choice function to be directly or truthfully implemented?

3.5 1-st Price Auction as an Example of B-N Implementation

Here is a specific example. First pricing sealed bid auction is a mechanism that induces the B-N equilibrium result. We know that the symmetric equilibrium of agents consists of bidding a fraction of their private valuation, specifically (n-1)/n for n bidders. The social choice function implicitly implemented by this mechanism is the following: first, anyone with the highest valuation gets the object, the same as in the first bid pricing. Second, in the case of wine, the transfer is equal to one half of the winner's valuation in the two-bidder case. This is different from the second pricing auction. (The social choice function is again ex-post efficient as the mechanism implies the highest valuation get the object.) According to the revelation principle, social choice function can be also directly implemented. That is the seller asks each one's private value

and assigns the object to the highest bidder and asks for half of their private valuation.

To generalize the conditions required for a social choice function to be truthfully implemented, we again consider the setting with quasi-linear utility and a decision function to undertake the public project subject to group resource constraints. The benefit term is proportional to the private valuation and a transformation from the yes/no project choice. By reporting different types, the agent affects the project choice. The second term is again a transfer depending upon the announced type profile.

Three conditions are important. First, the types are independently distributed across agents. Second, the expected benefit should be a non-decreasing function of the type. In the case of first pricing auction, the "benefit" is simply the yes/no answer if she wins the object. The conditions mean the expected chance of winning given she reports her true valuation (and so do others) is non-decreasing with her own private valuation. Loosely speaking, everyone reports their true valuation and the higher the valuation is, the more likely she is to win. True. Third, the expected utility of any type given everyone reports true types is equal to the expected utility of the lowest type plus the cumulative expected "benefit" from all types in between. In the auction example, the expected utility of the lowest type is zero. No chance to win anything or needs to pay anything. The expected "benefit" for any given valuation is the probability that such a valuation is above all other bidders. The expected utility for the same type is the valuation times the expected probability of winning, subtracting the payment times the probability of winning, which is half of the valuation conditional on wining. It turns out the conditions hold for the first pricing auction, which confirms that it could be implemented truthfully.