Strategic communication and manipulation

Summary

In the first part of the thesis, we study the sender-receiver games. In this part we study the behaviour of the partially informed receiver and fully informed sender. Receiver can use the information provided by the sender to make his decisions. However while making decisions based on this information, the receiver must take into consideration that the sender is strategic and self-interested.

In Chapter 2, we define the sender-receiver stopping games. This was the first model introduced which combines the features from dynamic sender-receiver games (where sender and receiver interact repeatedly for finitely or infinitely many periods) and the stopping games (where the duration of the game is determined by the actions of the players involved). In this setting the sender and the receiver interact repeatedly until the receiver stops the game. Payoff functions of the players are monotonically increasing in the state of the world. So both players prefer that the game ends when the state is high enough, although due to different payoff functions, they have different interpretation of the state being high enough.

Since the sender can only send one of the two possible messages to the sender, the receiver obtains very little information about the state of the world. If the state is high enough for the sender, he has incentive to send a message which is interpreted as the suggestion for the receiver to stop the game. And if the state is low, he has incentive to send the other message which is interpreted as the suggestion for the receiver to continue the game. Although the receiver may have different cardinal preferences, obeying the sender is a best response for the receiver.

In the finite horizon, the regular strategy profile, that is the strategy profile in which the sender sends the message according to his optimal threshold values (sincere strategy) and the receiver obeys the sender's message (obeying strategy) is a Perfect Bayesian Equilibrium (PBE),

provided the payoffs are undiscounted or discounted with sufficiently large discount factor. Moreover, this is the unique responsive PBE, that is, PBE in which the receiver plays non-babbling strategy.

In the infinite horizon, uniqueness of responsive PBE also holds for discounted payoff functions with sufficiently large discount factor, among the essentially Markovian strategy profiles. However, when the payoffs are undiscounted, there is no responsive PBE. This is the consequence of the fact that, in the infinite horizon game when the payoffs are undiscounted, the sender always expects to obtain higher state in the future with probability 1, so he always sends the suggestion to continue, and the receiver obeys. So the game never stops and the both players receive payoff zero.

In Chapter 3, we consider one of the natural extension of the sender-receiver games in which there are multiple senders. In the case of only one sender, under responsive PBE, the receiver must obey the sender's suggestions. In the setting with many senders, the receiver obtains more information about the state of the world, hence the receiver can make better informed decisions. Due to the complexity of the interaction between the senders and receiver, this setting is much more difficult to analyze. We only consider the finite horizon setting.

In this setting the obeying strategy is not relevant since the receiver have information about the state from multiple senders. The reactive strategy of the receiver is a strategy with threshold on the number messages, such that the receiver stops the game if he receives more number of 'quit' messages than the threshold and continues if he receives less number of 'quit' messages than the threshold. This strategy is the most natural extension of obeying strategy from the previous setting. The regular profile in this setting corresponds to the strategy profile in which each sender plays the sincere strategy and the receiver plays the reactive strategy. Regular strategy profile is a PBE. Moreover, no set of senders can collude together to obtain better payoff. However, the receiver may be able to collude with one or more senders to obtain strictly better payoff for all the players involved in the coalition. In this setting, the uniqueness of PBE does not hold.

We study the externality effect of an additional sender over the receiver and the remaining senders. With more senders, the receiver obtains more information about the state, however it does not necessarily mean that the receiver gets higher expected payoff. We construct specific example of games such that the receiver obtains strictly higher payoff in a game with one sender than another game with two senders, where one of the senders in two sender game has same payoff as in the first game. Having an additional sender may be strictly better, strictly worse or indifferent for the remaining senders when there are at least 2 remaining senders. But having an additional sender could only be strictly worse or indifferent for the remaining senders when he is the lone sender.

In the second part of thesis which consists Chapter 4, we consider a model of multidimensional binary domain. In this domain, agents who have a yes or no opinion on different issues called components, vote for alternatives which are binary vectors. The preferences of agents over the alternatives are determined by the number of disagreements on components with their top alternative. The objective of this chapter is to study if agents can manipulate the social choice functions by misreporting their true preferences.

We study rules, which are social choice functions that satisfy unanimity, anonymity and neutrality. It turns out the preference domain in this setting is highly restricted. So, the standard Gibbard–Satterthwaite theorem does not hold. Particularly, we show that the component-wise majority rules are strategy-proof. Strategy-proofness is a fairly weak concept in this setting, so the characterization of all strategy-proof rules is not feasible. Due to this, we study stronger notion of the weak group strategy-proofness in which the group of agents can collectively try to manipulate the outcome of the social choice function which is strictly better for every agent in the coalition. Surprisingly, the component-wise majority rules are not weakly group strategy-proof.

We show the existence of rules which are weakly group strategy-proof for certain cases.

Even stronger notion we study is strong group strategy-proofness, in which group of agents can collectively try to manipulate the outcome of the social choice function which is weakly better for every agent in the coalition and strictly better for at least one agent. We show that if there are at least 3 components and 4 agents, there does not exist any strong group strategy-proofness rule. In the case of 3 components and 3 agents, we show that component-wise majority is the unique strong group strategy-proof rule.