Monotonicity and Bayes-Nash implementation

Every day we are faced with strategic situations the outcome of which depends on our interaction with other people. Think for example of everyday situations like voting with colleagues on where to go for lunch, bidding on a rare Ramones record on Ebay or playing a round of Risk with friends. Game theory attempts to model people's behavior in such strategic situations mathematically.

The topic of this thesis falls into a sub-field of game theory called mechanism design. Mechanism design is the branch of game theory that concerns itself with designing the rules of games. A game can yield several possible outcomes. Each participating agent has a set of possible actions at his disposal that he can choose from. His choice of action influences the outcome of the game. How exactly is specified by a mechanism which defines the rules of the game. Such a mechanism consists of two rules, an allocation rule and a payment rule. Based on the actions taken by all the agents, the allocation rule determines the outcome of the game. Similarly, the payment rule determines for each agent a payment to make or to receive.

In order to illustrate the above let us consider the auction for a painting as an example. The agents in this auction game are the participating. One possible outcome of the auction is for example that bidder A wins the painting. Another one would be for bidder B to win. The action a bidder can take to influence the auction outcome is to make a money bid for the painting. Based on the bids from all the bidders an allocation rule determines who gets the painting and a payment rule determines how much everyone has to pay.

An agent’s outcome preferences are expressed by his valuation function. His type represents information available to him that is influential to his valuation of the different possible outcomes. Coming back the painting auction example, an agent's type can for example be a representation of how much he likes the painting. In this thesis we focus on direct revelation mechanisms, a special class of mechanisms. In a direct revelation mechanism the only action an agent can take is to announce a type. An agent's strategy is called dominant if it maximizes his utility for every possible combination of actions taken by the other agents. One goal of the mechanism designer could be to design a direct revelation mechanism which is dominant strategy incentive compatible. This means creating a mechanism such that truthfully announcing his own type is a dominant strategy for every agent. In this thesis we are concerned with Bayes-Nash settings. In these settings agents do not know which types the others have but the distribution of potential types for each agent is publicly known. More specifically, we are investigating Bayes-Nash incentive compatible direct revelation mechanisms. This means that truthfully reporting his own type maximizes each agent's expected utility, given that all the other agents also report truthfully.

In a recent stream of literature researchers are concerned with the characterization of dominant strategy incentive compatible direct revelation mechanisms in terms of a monotonicity condition on the allocation alone. That is, they want to characterize precisely those allocation rules for which they can guarantee the existence of a payment rule that makes the resulting mechanism dominant strategy incentive compatible.

In Chapters 2 and 3 we adapt this idea to Bayes-Nash incentive compatible mechanisms. In Chapter 2 we consider direct revelation mechanisms in settings where agents have independently distributed, one-dimensional types and quasi-linear utility functions. Agents are allowed to have interdependent valuations. We show that monotonicity, a condition on the allocation rule, is a necessary condition for Bayes-Nash incentive compatibility. This condition expresses that an agent's expected gain in valuation for truthfully reporting a type t

instead of misreporting a type *s* should be at least as big as his expected gain in valuation for misreporting *t* instead of truthfully reporting *s*.

In order to establish the sufficiency of monotonicity for Bayes-Nash incentive compatibility we take a network approach similar to the one taken by Gui, Müller and Vohra (2004) and Saks and Yu (2005) in their work with regard to dominant strategy incentive compatible mechanisms. Recognizing that the constraints inherent in the definition of Bayes-Nash incentive compatibility have a natural network interpretation we build complete directed graphs for agents' type spaces. To do so we associate a node with each type and put a directed edge between each ordered pair of nodes. The length of the edge going from the node associated with type *s* to the node associated with the type *t* is defined as the cost of manipulation. That is, the expected difference in an agent's valuation for truthfully reporting *s* instead of misreporting *t*. Monotonicity corresponds to the absence of 2-cycles with negative length in these graphs.

We show that an allocation rule is Bayes-Nash incentive compatible if and only if the graphs described above contain no negative length cycles. This result is the Bayes-Nash equivalent to a finding by Rochet (1987) who shows that dominant strategy incentive compatibility can be characterized in terms of the absence of negative length cycles in similar graphs.

We demonstrate that in order for monotonicity to be a sufficient condition for Bayes-Nash incentive compatibility we have to put further restrictions on our setting. Such conditions are decomposition monotonicity of the costs of manipulation or agents' valuation functions satisfying non-decreasing expected differences. For these restrictions we can establish the sufficiency of monotonicity. Characterization results of Myerson (1981), Malakhov and Vohra (2004) and Feng (2008) for Bayes-Nash incentive compatible mechanisms follow as special cases from our characterization results. Their settings can be folded into our framework which allows for a broader class of type spaces and alternative forms interdependencies between agents' valuations.

Furthermore, we show how to construct corresponding payment rules for Bayes-Nash incentive compatible allocation rules by using shortest path length in the agent's graphs. An allocation rule satisfies revenue equivalence if all payment rules that make it Bayes-Nash incentive compatible yield the same expected payments up to an additive constant. We establish that revenue equivalence holds if agents' type spaces are convex and their expected valuations are convex in their own types. Thus, we establish revenue equivalence under the same conditions as Krishna and Maenner (2001).

In Chapter 3 we extend the analysis of Chapter 2 to settings with multi-dimensional type spaces. We follow the same network approach as before and keep the structure of the chapter as close as possible to the foregoing one. In this chapter we have to restrict agents' type spaces to be convex and their expected valuation functions to be convex in their own types. We show that if the costs of manipulation are decomposition monotone or agents' valuation functions satisfying non-decreasing expected differences then monotonicity in conjunction with an integrability condition is both necessary and sufficient for Bayes-Nash incentive compatibility. Although monotonicity alone is not sufficient anymore, we still achieve a characterization that is purely based on the valuations and the allocation rule. Characterization results of Jehiel, Moldovanu and Stacchetti (1999), Jehiel and Moldovanu (2001) and Müller, Perea and Wolf (2007) follow as special cases from our characterization results. Again we can establish revenue equivalence under the same conditions as Krishna and Maenner (2001).

In Chapter 4 we consider scoring auctions. In the painting auction example described before competition among the bidders works only via the price. Essentially every bidder announces a price he this willing to pay and whoever makes the highest offer gets the painting. However, sometimes the price of an item is not the only concern and other non-monetary attributes play a role as well. Consider for example several construction companies competing for a government contract to build a highway. One important feature is of course the price. Another one is for example the completion time. So one company might offer to build the highway for 20 million dollars within 6 months and another company might offer to build it for 15 million dollars but only within 9 months.

One common method employed to deal with the analysis of scoring auction settings is to skillfully transform the setting into that of a standard, price-only auction. Then known results about these price-only auctions can be applied. We do this for a scoring auction setting in which several items with multiple attributes need to be allocated. We show how to transform this setting into a combinatorial auction setting that looks like the linear valuations setting discussed in Section 3.4.1 of the foregoing chapter. In doing this we extend the work of Asker and Cantillon (2008) to a framework with multiple items. Asker and Cantillon (2008) consider a scoring auction setting in which a single item with multiple attributes has to be allocated. They transform this into a single item, price-only auction setting.