Foundations of Mechanism Design

Dinesh Garg and Y. Narahari
Electronic Commerce Laboratory
Department of Computer Science and Automation
Indian Institute of Science
Bangalore - India
{ dgarg,hari }@csa.iisc.ernet.in

Abstract

Mechanism design has emerged as an important tool to model, analyze, and solve decentralized design problems in engineering involving multiple agents that interact strategically in a rational and intelligent way. Mechanism design is concerned with settings where a social planner faces the problem of aggregating the announced preferences of multiple agents into a collective decision when the actual preferences are not publicly known. The objective of this paper is to provide a tutorial introduction to the foundations and key results in mechanism design theory.

Acronyms

SCF	Social Choice Function
IC	Incentive Compatibility (Compatible)
DSIC	Dominant Strategy Incentive Compatible
BIC	Bayesian Nash Incentive Compatible
AE	Allocative Efficiency (Allocatively Efficient)
BB	Budget Balance
IR	Individual Rationality
IIR	Interim Individually Rational
DSIC	Dominant Strategy Incentive Compatibility (Compatible)
VCG	Vickrey-Clarke-Groves Mechanisms
BIC	Bayesian Incentive Compatibility (Compatible)
dAGVA	d'Aspremont and Gérard-Varet mechanisms

Notation

n	Number of agents
N	Set of agents: $\{1, 2, \ldots, n\}$
Θ_i	Type set of Agent i
Θ	$= \Theta_1 \times \ldots \times \Theta_n$
Θ_{-i}	$= \Theta_1 \times \ldots \times \Theta_{i-1} \times \Theta_{i+1} \times \ldots \times \Theta_n$
θ_i	Actual type of agent i, $\theta_i \in \Theta_i$
θ	$= (\theta_1, \ldots, \theta_n)$
θ_{-i}	$= (\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_n)$
$\hat{\theta}_i$	Reported type of agent i, $\hat{\theta}_i \in \Theta_i$
$\hat{\theta}$	$= (\hat{\theta}_1, \ldots, \hat{\theta}_n)$
$\hat{\theta}_{-i}$	$= (\hat{\theta}_1, \ldots, \hat{\theta}_{i-1}, \hat{\theta}_{i+1}, \ldots, \hat{\theta}_n)$
$\Phi_i(.)$	A CDF on Θ_i
$\phi_i(.)$	A PDF on Θ_i
$\Phi(.)$	A CDF on Θ
$\phi(.)$	A PDF on Θ
X	Outcome Set
x	A particular outcome, $x \in X$
$u_i(.)$	Utility function of agent i
$f(.)$	A social choice function
$\mathscr{M}$	An indirect mechanism
$\mathscr{D}$	A direct revelation mechanism
$g(.)$	Outcome rule of an indirect mechanism
S_i	Set of actions available to agent i in an indirect mechanism
S	$= S_1 \times \ldots \times S_n$
b_i	Bid of agent i
b	$= (b_1, \ldots, b_n)$
b_{-i}	$= (b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n)$
$b^{(k)}$	k^{th} highest element in $(b_1, \ldots, b_n)$
$(b_{-i})^{(k)}$	k^{th} highest element in $(b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n)$
$s_i(.)$	Strategy of agent i
$s(.)$	$= (s_1(.), \ldots, s_n(.))$
K	Set of project choices
k	A particular project choice, $k \in K$
t_i	Monetary transfer to agent i
$v_i(.)$	Valuation function of agent i
$U_i(.)$	Expected utility function of agent i
$w(.)$	Social welfare function
F	Set of social choice functions
X_f	Set of feasible outcomes

1 Introduction

In the second half of the twentieth century, game theory and mechanism design have found widespread use in a gamut of applications in engineering. More recently, game theory and mechanism design have emerged as an important tool to model, analyze, and solve *decentralized design problems* in engineering involving multiple autonomous agents that interact strategically in a *rational* and *intelligent* way. The agents are *rational* in the game theoretic sense of making decisions consistently in pursuit of their own individual objectives. Each agent's objective is to maximize the expected value of his/her own payoff measured in some utility scale. *selfishness* or *self-interest* is an important implication of rationality. Each agent is *intelligent* in the game theoretic sense of knowing everything about the underlying game that a game theorist knows and each agent can make any inferences about the game that a game theorist can make. In particular, each agent is *strategic*, that is, takes into account his/her knowledge or expectation of behavior of other agents and is capable of doing the required computations.

The theory of *mechanism design* is concerned with settings where a policy maker (or social planner) faces the problem of aggregating the *announced preferences* of multiple agents into a collective (or social) decision when the *actual preferences* are not publicly known. Mechanism design theory uses the framework of *non-cooperative games with incomplete information* and seeks to study how the privately held preference information can be elicited and the extent to which the information elicitation problem constrains the way in which social decisions can respond to individual preferences. The main focus of mechanism design is the design of institutions or protocols that satisfy certain objectives, assuming that the individual agents, interacting through the institution, will act strategically and may hold private information that is relevant to the decision at hand.

1.1 Examples of Mechanism Design Problems

There are numerous examples of design problems where mechanism design has come to play a key role. These include: auctions and markets in electronic commerce; routing protocols in wired and wireless networks; resource allocation in computational grids; sponsored search auctions on the web; algorithms for selfish agents; cryptographic protocols for networks with selfish agents; etc. We describe some of these design problems in a succinct way below.

1.1.1 Mechanism Design and Algorithms for Selfish Agents

In distributed settings, there are certain algorithmic problems where the agents cannot be assumed to follow the algorithm but are driven by selfish goals [1]. In such situations, the agents are capable of manipulating the algorithm. For example, in a shortest path problem where edges are owned by individual agents and the costs of the edges are known only to the owning agents, the algorithm designer is faced with the challenge of eliciting the true costs of the edges from the owning agents before applying an algorithm for computing the shortest paths. The objective of the algorithm designer in such situations should be to come up with a scheme which ensures that the agents' interests are best served by behaving truthfully. Mechanism design theory provides a handle for studying and designing such algorithms. The mechanism is designed in a way that all agents are motivated to act truthfully and according to the wishes of the algorithm designer.

1.1.2 Mechanism Design and Selfish Routing

A major difficulty that is faced by networks such as road transportation networks, communication networks, and the Internet is that most of the time, the demand of the users for the network resources exceeds the available supply of the resources. This phenomenon causes congestion in the network. The traffic congestion can be avoided or mitigated if the arrival rates can be controlled and/or the traffic can be routed across the network in an effective manner. However, in most of these networks, the users of such networks are free to act according to their own interests and moreover, they are rational and intelligent in the sense that they care more for their own individual welfare and less for the health or performance of the overall network. Game theory and mechanism design can play an important role in the analysis and design of protocols in such networks. Indeed, in the past few years, there has been a spurt of research activities in this direction. See, for example, the work of Roughgarden [2, 3], Feigenbaum *et al* [4, 5], Hershberger and Suri [6], and Nisan *et al* [7, 8].

1.1.3 Mechanism Design and Ad-hoc Networks

Wireless ad-hoc networks also face the traffic congestion problem just like other wired networks such as communication networks and the Internet. However, due to the wireless nature of the communication medium, the physics of the problem is somewhat different. In the case of wireless ad-hoc networks, conservation of battery power by the individual nodes is of primary importance. Individual nodes are required to forward packets so as to ensure connectivity in an ad-hoc network and nodes therefore double up as routers. Forwarding of packets involves draining battery power and it may not always be in the self-interest of a node to forward packets. Design of protocols that stimulate cooperative actions from the nodes uses game theory and mechanism design in a significant way. There are several recent efforts in this direction. See, for example, the work of Anderegg and Eidenbenz [9] in unicast environment, where they use the VCG (Vickrey-Clarke-Groves) mechanism to compute a power efficient path using which the nodes transmit or forward the packets. Eidenbenz, Santi, and Resta [10] further generalize this model to achieve budget balance property. Wang and Li [11] propose strategy-proof pricing schemes for multicast environment. Suri [12] has generalized the model of Wang and Li and has proposed incentive compatible broadcast protocols in ad-hoc wireless networks.

1.1.4 Mechanism Design and Grid Computing

One of the most important issues concerned with grid computing is that of application scheduling. In a global grid setting, the individual users must be provided with an incentive to offer their resources. The situation becomes non-trivial because of the fact that these entities are rational and intelligent resource providers who, for strategic reasons, may not provide truthful information about their processing power and cost structure. Thus, resource allocation decisions in computational grids and other general grids have to take into account the rationality of the grid resource providers. In particular, there is a need for providing appropriate incentives to the nodes to stimulate their participation in grid computing. Mechanism design is extremely useful in designing incentives and also in coming up with auction based mechanisms for grid resource allocation. There are several recent efforts in this direction. See, for example, the work of Grosu and Chronopoulos [13, 14], Das and Grosu [15], Wellman, Walsh, Wurman, and Mackie-Mason [16], and Buyya [17]. Prakash [18] has recently proposed several innovative mechanisms for resource procurement in computational grids with rational resource providers.

1.1.5 Mechanism Design and Cryptography

Mechanism design can be viewed as a science of synthesis of protocols for selfish parties to achieve certain properties. In mechanism design, it is often assumed that the central authority (or social planner) can be trusted by the parties, but this might not always be true, especially in an Internet environment. If the social planner is corrupt, then it may misuse the information received from the agents. Privacy is therefore essential in order to ensure the social planner's credibility [19]. This problem was first stated by Varian [20]. Because cryptography deals with preserving privacy and integrity of data in computer and communication systems, it is obvious that techniques from cryptography may help implementing the mechanism design protocols in the real word. There has been some progress in the recent past in applying cryptographic tools and techniques to the problems of mechanism design. For example, see the work related to privacy preserving auctions by Noar *et al* [21] and Brandt [22].

1.1.6 Mechanism Design and the World Wide Web

We know that the world wide web has become an integral part of our day-to-day life. As engineers and scientists are making progress in developing innovative web based services, the hackers and spammers are bent on making such services collapse. Examples of such services include *search engines page ranking system, recommender systems,* and *reputation systems*. Note that it is not always the case that these hackers and spammers invest their energy and efforts just for nothing. More often than not, they do have a stake in not letting such systems run smoothly. The only solution for such a problem is to design the system in a way that there is no incentive for the hackers and spammers to put their energy and time in destroying it. Game theory and mechanism design theory play an important role in designing such systems. For example see the work of Altman and Tennenholtz [23, 24] for a mechanism design view of page ranking systems, and Gyongyi and Molina [25] for a kind of threat to the page ranking system, namely link spam.

1.1.7 Mechanism Design and Electronic Commerce

Electronic commerce is an area where the rationality in general and self-interest in particular of the participating agents is a significant factor to be taken into account. Game theory and mechanism design, therefore, have come to play a natural role here. For example, in a bargaining problem between a buyer and a seller, the seller would like to act as if the item is very expensive thus raising its price and the buyer would like to pretend to have a low value for the object to keep the price down. In this context, mechanism design helps to come up with a bargaining protocol that ensures an efficient trade of the good, so that successful trade occurs whenever the buyer's valuation exceeds that of the seller.

There are a variety of game theoretic design problems currently being explored in the area of electronic commerce. These design problems are well documented in the literature; the following is a representative listing:

- Auctions for selling spectrum [26]
- Mechanisms for selling advertising space through keyword auctions [27, 28, 29]
- Auctions for selling products/services as part of private marketplaces set up by e-business companies [30]
- Bandwidth exchanges [31]

- Procurement auctions and private marketplaces for e-procurement [32]
- Logistics and transportation marketplaces [33]
- Mechanisms for supply chain formation [34, 35, 36]

1.2 Outline of the Paper

- We start with definition of the mechanism design problem, the concept of a *social choice function*, the concept of a *mechanism*, and the notion of *implementation* of a social choice function by a mechanism. We also distinguish between direct revelation mechanisms and indirect mechanisms. These constitute the fundamental setup for mechanism design theory. We illustrate the concepts through several examples: bilateral trade, auctions for selling a single indivisible item, and a combinatorial auction. These examples are used throughout the rest of the paper for bringing out important insights.

- Next, we describe the desirable properties of a social choice function, which include *ex-post efficiency, non-dictatorial-ness, dominant strategy incentive compatibility* and *Bayesian Nash incentive compatibility.* We then state and prove a fundamental result in mechanism design theory, the *revelation theorem.*

- We then describe a landmark result - *Gibbard-Satterthwaite* impossibility theorem, which says that under fairly general conditions, no social choice function can satisfy the three properties - ex-post efficiency, non-dictatorial, and dominant strategy incentive compatibility simultaneously. This impossibility theorem, while ruling out implementability of certain desirable mechanisms, suggests two alternative routes to the design of useful mechanisms.

- The first of these two routes is to restrict the utility functions to what is known as a *quasi-linear environment*, and then investigating the properties of the social choice function. Here we show that ex-post efficiency is equivalent to a combination of two properties, namely, *allocative efficiency* and *budget balance.* The important result here is the existence of the classic *VCG (Vickrey-Clarke-Groves)* social choice functions (also known as VCG mechanisms), which are non-dictatorial, dominant strategy incentive compatible, and allocatively efficient.

- We next explore the second route suggested by the Gibbard-Satterthwaite impossibility theorem. Here we look for Bayesian incentive compatibility instead of dominant strategy incentive compatibility. We show that a class of social choice functions, known as *dAGVA* social choice functions are ex-post efficient, non-dictatorial, and Bayesian Nash incentive compatible. We then develop a characterization of Bayesian incentive compatible social choice functions in linear environment.

- The next topic we discuss is the celebrated *Revenue Equivalence Theorem* which states that the expected revenue produced by the English auction, Dutch auction, first price auction, and second price auction is the same, in the context of selling of a single indivisible item.

- This is followed by a discussion of the concept of individual rationality and the well known *Myerson-Satterthwaite* impossibility theorem.

- We next focus on design of *optimal mechanisms*. Here, we first define the notion of optimal mechanism design and then describe the seminal work of Myerson [37] on design of optimal auctions.

- Finally, we suggest several pointers to the literature to delve deeper into mechanism design theory.

2 The Mechanism Design Problem

The theory of mechanism design concerns with settings where a *social planner (or policy maker)* faces the problem of aggregating the individual preferences into a collective (or social) decision and the individuals' actual preferences are not publicly known. The *mechanism design problem* provides a good starting point for this theory. A formal description of the mechanism design problem is provided below.

1. There are n individuals (or agents), indexed by $i = 1, 2, \ldots, n$, who must make a collective choice from some set X, called the *outcome set* or *choice set*.

2. Prior to the choice, however, each agent i privately observes his preferences over X. Formally, this is modeled by supposing that agent i observes a parameter, or signal θ_i that determines his preferences. The parameter θ_i is referred to as agent i's type. The set of possible types of agent i is denoted by Θ_i.

3. The agents' types, denoted by $\theta = (\theta_1, \ldots, \theta_n)$ are drawn according to a probability distribution function $\Phi \in \Delta\Theta$, where $\Theta = \Theta_1 \times \ldots \times \Theta_n$, and $\Delta\Theta$ is the set of all the probability distribution functions over the set Θ. Let ϕ be the corresponding probability density function.

4. Each agent i is rational and intelligent and this fact is modeled by assuming that the agents always try to maximize a utility function $u_i : X \times \Theta_i \to \mathbb{R}$.

5. The probability density $\phi(\cdot)$, the type sets $\Theta_1, \ldots, \Theta_n$, and the utility functions $u_i(\cdot)$ are assumed to be *common knowledge* among the agents. Note that the utility function $u_i(\cdot)$ of agent i depends on both the outcome x and the type θ_i. Even though, the type θ_i is not common knowledge, by saying that $u_i(\cdot)$ is common knowledge we mean that for any given type θ_i, the social planner and every other agent can evaluate the utility function of agent i.

In the above situation, the social planner faces two problems:

1. **Preference Aggregation Problem**: The first problem that is faced by the social planner is the following: "For a given type profile $\theta = (\theta_1, \ldots, \theta_n)$ of the agents, which outcome $x \in X$ should be chosen?"

2. **Information Revelation (Elicitation) Problem**: Assuming that the preference aggregation problem has been solved, the next problem that is faced by the social planner is the following: "How do we extract the true type θ_i of each agent i, which is the private information of agent i?"

In what follows, we will address both these problems one by one. Let us first look at the preference aggregation problem. The *social planner or policy maker* can solve the preference aggregation problem by choosing an appropriate function known as *social choice function*, which is defined next.

3 Social Choice Function

Definition 3.1 *A social choice function (SCF) is a function $f : \Theta \to X$, which a social planner or policy maker uses to assign a collective choice $f(\theta_1, \ldots, \theta_n)$ to each possible profile of the agents' types $\theta = (\theta_1, \ldots, \theta_n) \in \Theta$.*

Figure 1 illustrates the idea behind social choice function. In what follows, we present a few examples

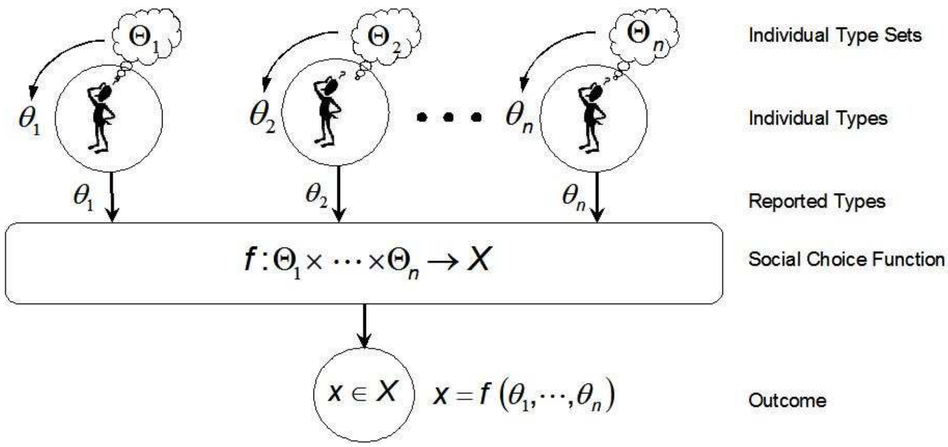

Figure 1: The idea behind social choice function

of the mechanism design problem and the social choice function that is being used by the social planner. Most of these examples are taken from [38].

3.1 Example: Allocation of a Single Unit of an Indivisible Private Good

Consider a situation where there is a set N of n agents and one of them is owning one unit of an indivisible good. The owner wants to trade this good with other agents by means of money. The other agents are interested in trading with him. Let us assume that there is an outside agent who plays the role of a broker and facilitates the trading activities among these agents. Each of the n agents interact directly with the broker agent in a personal manner but they do not interact among themselves. This broker agent can be viewed as a social planner. Here the problem of the broker agent is to decide whom to allocate the good and how much money to charge (or pay) from each agent. This can be viewed as a mechanism design problem and the various building blocks of the underlying mechanism design structure can be given in following manner.

1. **Outcome Set** X: An outcome in this case may be represented by a vector $x = (y_1, \ldots, y_n, t_1, \ldots, t_n)$, where $y_i = 1$ if the agent i receives the object, $y_i = 0$ otherwise, and t_i is the monetary transfer received by the agent i. The set of feasible alternatives is then

$$X = \left\{ (y_1, \ldots, y_n, t_1, \ldots, t_n) | y_i \in \{0, 1\}, t_i \in \mathbb{R} \ \forall\, i, \sum_{i=1}^{n} y_i = 1, \sum_{i=1}^{n} t_i \leq 0 \right\}$$

2. **Type Set** Θ_i: In this example, the type θ_i of an agent i can be viewed as his valuation of the good. We can take the set of possible valuations for agent i to be $\Theta_i = [\underline{\theta_i}, \overline{\theta_i}] \subset \mathbb{R}$, where $\underline{\theta_i}$ is the least valuation and $\overline{theta_i}$ is the highest valuation the agent may have for the item.

3. **Utility Function** $u_i(\cdot)$: The utility function of agent i can be given by

$$\begin{aligned} u_i(x, \theta_i) &= u_i(y_1, \ldots, y_n, t_1, \ldots, t_n, \theta_i) \\ &= \theta_i y_i + t_i \end{aligned}$$

4. **Social Choice Function** $f(\cdot)$: The general structure of the social choice function for this case is

$$f(\theta) = (y_1(\theta), \ldots, y_n(\theta), t_1(\theta), \ldots, t_n(\theta)) \;\; \forall\, \theta \in \Theta$$

Two special cases of this example have received a great deal of attention in the literature - *bilateral trading* and *auctions*. In what follows, we discuss each of these two.

3.2 Example: Bilateral Trade

This is a special case of the previous example. In this setting, we have $n = 2$. Agent 1 is interpreted as the initial owner of the good (the "seller"), and agent 2 is the potential purchaser of the good (the "buyer"). A few facts are in order. Note that

1. If $\overline{\theta_1} < \underline{\theta_2}$ there are certain gains for both the agents from the trade regardless of θ_1 and θ_2.
2. If $\overline{\theta_2} < \underline{\theta_1}$ then it is certain there are no gains for both the agents from the trade.
3. If $\underline{\theta_2} < \overline{\theta_1}$ or $\underline{\theta_1} < \overline{\theta_2}$ or both, then there may or may not be gains for the agents from trade, depending on the realization of θ.

We will develop this example further in this paper as and when required.

3.3 Example: Single Unit - Single Item Auction

This is a special case of the example 3.1. For the purpose of illustration, we assume that there are $n+1$ agents instead of n agents. Note that there is no loss of generality in doing so. Let agent 0 be the owner of the good with the rest of the n agents interested in buying the good. In this case, the owner of the good is known as auctioneer and the rest of the agents are known as bidders. The distinguishing feature of this example which makes it a special case of the Example 3.1 is the following:

The type set of the seller is a singleton set, that is $\Theta_0 = \{\theta_0\}$, which is commonly known to all the buyers.

If $\theta_0 = 0$ then it implies that the auctioneer has no value for the good and he just wants to sell it off. We call such a scenario as single unit single item auction without reserve price. For this example, the various building blocks of the underlying mechanism design structure are the following.

$$\begin{aligned} X &= \Big\{(y_0, y_1 \ldots, y_n, t_0, t_1, \ldots, t_n) | y_0 = 0, t_0 \geq 0, y_i \in \{0,1\}, t_i \leq 0 \;\forall\, i = 1, \ldots, n, \\ &\qquad \sum_{i=0}^{n} y_i = 1, \sum_{i=0}^{n} t_i = 0\Big\} \\ \Theta_0 &= \{\theta_0 = 0\} \\ \Theta_i &= [\underline{\theta_i}, \overline{\theta_i}] \subset \mathbb{R} \;\forall\, i = 1, \ldots, n \\ u_0(x, \theta_0) &= u_0(y_0, y_1 \ldots, y_n, t_0, t_1, \ldots, t_n, \theta_0) = t_0 \\ u_i(x, \theta_i) &= u_i(y_0, y_1 \ldots, y_n, t_0, t_1, \ldots, t_n, \theta_i) = \theta_i y_i + t_i \;\; \forall\, i = 1, \ldots, n \\ f(\theta) &= (y_0(\theta), \ldots, y_n(\theta), t_0(\theta), \ldots, t_n(\theta)) \;\; \forall\, \theta \in \Theta \end{aligned}$$

3.4 Example: Single Unit - Single Item Auction with Reserve Price

This is the same example as the previous one except that now auctioneer has some positive value for the good, that is, $\theta_0 > 0$. The auctioneer announces a reserve price $r > 0$, which need not be the same as θ_0. All the other agents treat this reserve price as the valuation of the auctioneer for the good. This scenario is known as single unit - single item auction with reserve price. For this example, the various building blocks of the underlying mechanism design structure are the following.

$$\begin{aligned}
X &= \Big\{(y_0, y_1 \ldots, y_n, t_0, t_1, \ldots, t_n) | y_0 \in \{0,1\}, t_0 \geq 0, y_i \in \{0,1\}, t_i \leq 0 \; \forall \, i = 1, \ldots, n, \\
&\quad \sum_{i=0}^{n} y_i = 1, \sum_{i=0}^{n} t_i = 0 \Big\} \\
\Theta_0 &= \{\theta_0 > 0\} \\
\Theta_i &= \left[\underline{\theta_i}, \overline{\theta_i}\right] \subset \mathbb{R} \; \forall \, i = 1, \ldots, n \\
u_0(x, \theta_0) &= \theta_0 y_0 + t_0 \\
u_i(x, \theta_i) &= \theta_i y_i + t_i \; \forall \, i = 1, \ldots, n \\
f(\theta) &= (y_0(\theta), \ldots, y_n(\theta), t_0(\theta), \ldots, t_n(\theta))
\end{aligned}$$

3.5 Example: A Combinatorial Auction

This is a generalization of Example 3.3. Imagine a setting where an individual, say agent 0, is holding one unit of each of m different items and wants to sell these items. Let there be n agents who are interested in buying all the items or a non-empty subset (bundle) of the items. As before, the owner of the good is known as auctioneer and the rest of the agents are known as bidders. The auctioneer has no value for these items and just wants to sell them and this fact is a common knowledge among the bidders. Below are a few distinguishing features of this example which make it a generalization of Example 3.3.

1. The type set of the auctioneer is a singleton set, that is $\Theta_0 = \{\theta_0 = 0\}$. This means that auctioneer has no value for any of the bundle $S \subset M$.

2. The auctioneer's type θ_0 (also known as seller's valuation for the bundles) is commonly known to all the bidders.

3. The type of a bidder i consists of his valuation for every bundle of the items. Obviously, there are $2^m - 1$ possible bundles of the items, therefore, a type θ_i of a bidder i consists of $2^m - 1$ numbers.

For this example, the various building blocks of the underlying mechanism design structure are the following.

1. **Outcome Set** X: An outcome in this is a vector $x = (y_i(S), t_0, t_1, \ldots, t_n)_{i=1,\ldots,n, S \subset M}$, where $y_i(S) = 1$ if bidder i receives the bundle $S \subset M$, $y_i(S) = 0$ otherwise. t_0 is the monetary transfer received by the auctioneer and t_i is the monetary transfer received by the bidder i. The feasibility conditions require that each bidder i must be allocated at most one bundle, that is, $\sum_{S \subset M} y_i(S) \leq 1$. Also, each object $j \in M$ must be allocated to at most one bidder, that is,

$\sum_{S\subset M|j\in S}\sum_{i=1}^{n} y_i(S) \leq 1$. The set of feasible alternatives is then

$$\begin{aligned} X &= \{(y_i(S), t_0, t_1, \ldots, t_n)_{i\in N, S\subset M} | y_i(S) \in \{0,1\} \; \forall \, i \in N, S \subset M; \\ & \sum_{S\subset M} y_i(S) \leq 1 \; \forall \, i \in N; \sum_{S\subset M|j\in S} \sum_{i=1}^{n} y_i(S) \leq 1 \; \forall \, j \in M; \\ & t_0 \geq 0; t_i \leq 0 \; \forall \, i \in N; \sum_{i=0}^{n} t_i = 0 \Bigg\} \end{aligned}$$

2. **Type Set** Θ_i: In this example, the type set of the auctioneer is a singleton set $\Theta_0 = \{\theta_0 = 0\}$. However, the type θ_i of each of the bidders i is a tuple $(\theta_i(S))_{S\subset M}$, where $\theta_i(S) \in \left[\underline{V_i}, \overline{V_i}\right] \subset \mathbb{R} \; \forall \, S \subset M$ represents the value of agent i for bundle S. Here we are assuming that bidder i's valuation for each bundle S lies in the interval $\left[\underline{V_i}, \overline{V_i}\right]$. Therefore, the set Θ_i can be defined in the following manner:

$$\Theta_i = \left\{(\theta_i(S))_{S\subset M} | \theta_i(S) \in \left[\underline{V_i}, \overline{V_i}\right] \subset \mathbb{R} \; \forall \, S \subset M\right\} = \left[\underline{V_i}, \overline{V_i}\right]^{(2^m-1)}$$

3. **Utility Function** $u_i(\cdot)$: The utility function of the auctioneer and bidders can be given by

$$\begin{aligned} u_0 &= t_0 \\ u_i(x, \theta_i) &= t_i + \sum_{S\subset M} \theta_i(S) y_i(S) \;\; \forall \, i \in N \end{aligned}$$

4. **Social Choice Function** $f(\cdot)$: The general structure of the social choice function for this case is

$$f(\theta) = ((y_i(S, \theta))_{i\in N, S\subset M}, t_0(\theta), \ldots, t_n(\theta))$$

3.6 Need for a Mechanism

All the above examples illustrate the concept of social choice function (SCF). However, as we mentioned before, the problem of mechanism design does not end just by choosing some SCF. The social planner still needs to address the problem of information elicitation. Having decided about the SCF $f(\cdot)$, a trivial solution for the information elicitation problem seems to be to request the agents to reveal their types θ_i and then use them directly to compute the social outcome $x = f(\theta)$. However, given the fact that an outcome x yields a utility of $u_i(x, \theta_i)$ to agent i and agent i is a utility maximizing agent, the agents' preferences over the set X depend on the realization of their type profile $\theta = (\theta_1, \ldots, \theta_n)$. For any given realization θ, different agents may prefer different outcomes. Therefore, it is not surprising if agent i reveals untruthful type, say $\hat{\theta}_i$, to the social planner because doing so may help him drive the social outcome towards his most favorable choice, say $\hat{x}$. This phenomenon is known as the information revelation (or elicitation) problem and it is depicted in Figure 2. One way in which social planner can tackle this problem is the use of an appropriate mechanism.

4 Mechanisms

Definition 4.1 *A mechanism* $\mathscr{M} = ((S_i)_{i\in N}, g(\cdot))$ *is a collection of action sets* $(S_1, \ldots, S_n)$ *and an outcome function* $g : S \to X$*, where* $C = S_1 \times \ldots \times S_n$*.*

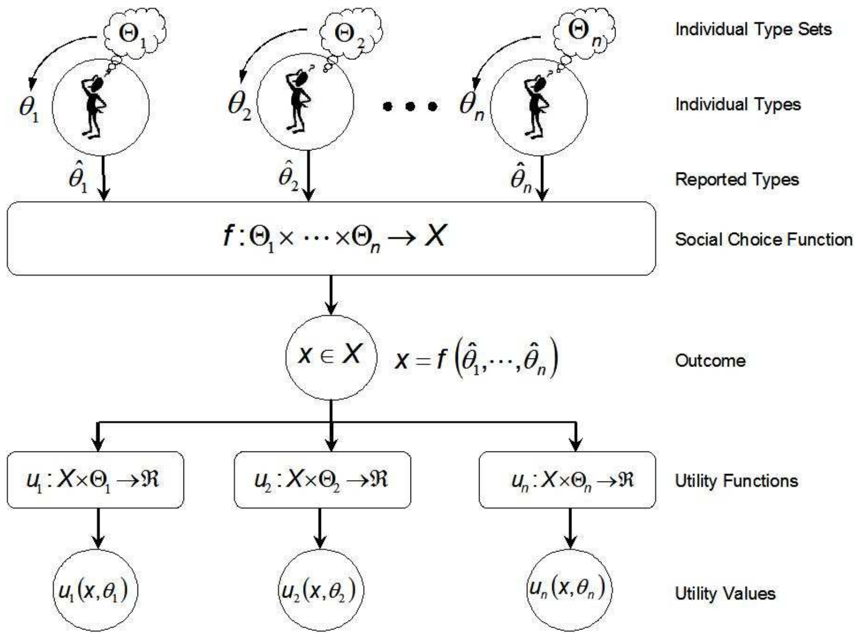

Figure 2: Information elicitation problem

The set S_i for each agent i describes the set of actions available to that agent. Based on his actual type θ_i, each agent i will choose some action, say $s_i \in S_i$. Once all the agents choose their actions, the social planner uses this profile of the actions $s = (s_1, \ldots, s_n)$ to pick a social outcome $x = g(c)$. This phenomenon is depicted in Figure 3. In view of the above definition, the trivial scheme of asking the agents to reveal their types becomes a special case; this special case is called a *direct revelation mechanism (DRM)*.

Definition 4.2 *Given a social choice function* $f : \Theta \to X$*, a mechanism* $\mathscr{D} = ((\Theta_i)_{i \in N}, f(\cdot))$ *is known as a direct revelation mechanism (DRM) corresponding to* $f(\cdot)$*.*

- Given a social choice function $f(\cdot)$, note that a direct revelation mechanism is a special case of a mechanism $\mathscr{M} = ((S_i)_{i \in N}, g(\cdot))$ with $S_i = \Theta_i \;\; \forall\, i \in N$ and $g = f$.
- Mechanisms that are not direct revelation mechanisms are typically referred to as *indirect mechanisms.*

The Figure 4 illustrates the idea behind direct revelation mechanism. In what follows, we present a few examples of both direct revelation mechanism and indirect mechanism.

4.1 Example: Fair Bilateral Trade

Consider the previous bilateral trade example Example 3.2. Let us assume that the broker agent invites the seller (agent 1) to report his type θ_1 directly to him in a confidential manner. θ_1 is the minimum amount at which agent 1 is willing to sell the item. The broker agent also invites the buyer (agent 2) to report his type θ_2 directly to him in a confidential manner. θ_2 is the maximum amount that agent 2 is willing to pay for the item. Let us assume that seller reports his type to be $\hat{\theta}_1$ which

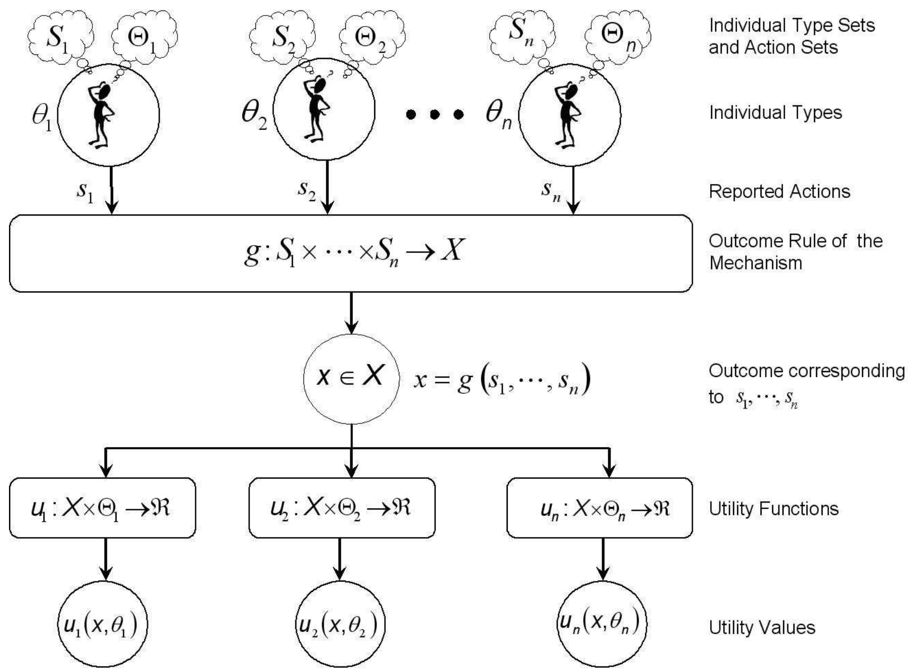

Figure 3: Abstract view of a mechanism

may be different from actual type θ_1 and buyer reports his type to be $\hat{\theta}_2$ which also may be different from actual type θ_2. Now the broker agent uses these reported types $\hat{\theta}_1$ and $\hat{\theta}_2$ in order to decide an outcome in the following manner. If $\hat{\theta}_1 < \hat{\theta}_2$ then the broker assigns the good to agent 2. The broker charges an amount equal to $\frac{\hat{\theta}_1+\hat{\theta}_2}{2}$ from agent 2 and pays the same amount to the agent 1. However, if $\hat{\theta}_2 < \hat{\theta}_1$ then no trade takes place.

In this example, the mechanism that is being used by the broker agent is a direct revelation mechanism $\mathscr{D} = (\Theta_1, \Theta_2, f(\cdot))$, where the social choice function $f(\cdot)$ is given by $f(\theta) = (y_1(\theta), y_2(\theta), t_1(\theta), t_2(\theta))$, where $\theta = (\theta_1, \theta_2)$. The functions $y_i(\cdot)$ are known as *winner determination rules* and the functions $t_i(\cdot)$ are known as *payment rules*. The winner determination and payment rules are given as follows

$$\begin{aligned} y_2(\theta) &= \begin{cases} 1 & : \quad \theta_1 < \theta_2 \\ 0 & : \quad \text{otherwise} \end{cases} \\ y_1(\theta) &= 1 - y_2(\theta) \\ t_1(\theta) &= y_1(\theta)\left(\frac{\theta_1 + \theta_2}{2}\right) \\ t_2(\theta) &= -t_1(\theta) \end{aligned}$$

4.2 Example: First-Price Sealed Bid Auction

Consider the Example 3.3 of single unit-single item auction without reserve price. Let us assume that the auctioneer himself is a social planner and he invites each bidder i to bid an amount $b_i \geq 0$ directly to him in a confidential manner. The bid b_i means that bidder i is ready to pay an amount b_i if he receives the good. The agent i decides bid b_i based on his actual type θ_i. The bids are opened by the

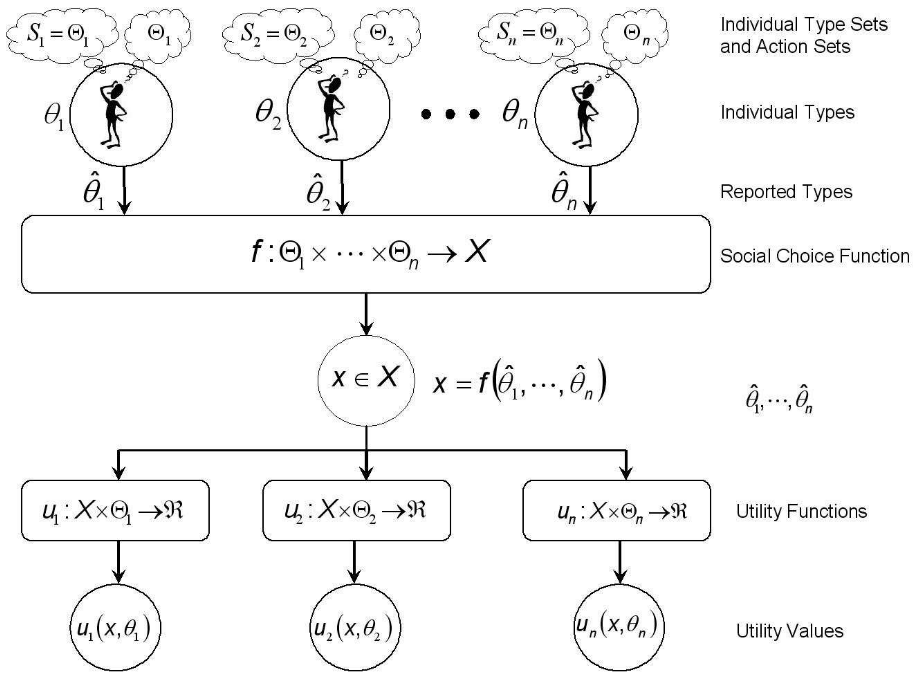

Figure 4: Abstract view of a direct revelation mechanism

auctioneer and the bidder with the highest bid gets the good and pays to the auctioneer an amount equal to his bid. The other bidders pay nothing. If there are several highest bids, we suppose that the lowest numbered of these bids gets the good. We could also break the tie by randomly selecting one of the highest bidders.

In this example, the mechanism that is being used by the auctioneer is an indirect mechanism $\mathscr{M} = ((S_i)_{i \in N}, g(\cdot))$, where $S_i \subset \mathbb{R}^+$ is the set of bids that bidder i can submit to the auctioneer and $g(\cdot)$ is the outcome rule given by $g(b) = (y_1(b), \ldots, y_n(b), t_1(b), \ldots, t_n(b))$, where $b = (b_1, \ldots, b_n)$. The functions $y_i(\cdot)$ are known as *winner determination rules* and the functions $t_i(\cdot)$ are known as *payment rules*. If we define $b^{(k)}$ to be the k^{th} highest element in $(b_1, \ldots, b_n)$ and $(b_{-i})^{(k)}$ to be the k^{th} highest element in $(b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n)$, then winner determination and payment rule can be written in following manner.

$$
\begin{aligned}
y_i(b) &= \begin{cases} 1 & : \text{ if } b_i = b^{(1)} \\ 0 & : \text{ otherwise} \end{cases} \\
t_i(b) &= -b_i y_i(b)
\end{aligned}
$$

A few remarks are in order with regard to this example.

- While writing the above winner determination and payment rule, we have assumed that any two bidders bidding the same bid value is a zero probability event.
- If $S_i = \Theta_i = [\underline{\theta_i}, \overline{\theta_i}] \quad \forall \ i \in N$ then this indirect mechanism becomes a direct revelation mechanism.

4.3 Example: Second-Price Sealed Bid (Vickrey) Auction

The setting is the same as the first-price auction. The only difference here is in terms of the allocation and payment rules invoked by the auctioneer. In the second-price sealed bid auction, the bidder with the highest bid gets the good and pays to the auctioneer an amount equal to the second highest bid. The winner determination and payment rules for this auction can be given as follows.

$$\begin{aligned} y_i(b) &= \begin{cases} 1 & : \quad \text{if } b_i = b^{(1)} \\ 0 & : \quad \text{otherwise} \end{cases} \\ t_i(b) &= -(b_{-i})^{(1)} y_i(b) \end{aligned}$$

This auction is also called as *Vickrey auction*, after the Nobel prize winning work of Vickrey [39].

4.4 Example: Generalized Vickrey Auction (GVA)

Once again, consider the example 3.5 of single unit-multi item auction without reserve price. Let us assume that auctioneer himself is a social planner and he invites each bidder i to report hid bid b_i directly to him in a confidential manner. In this example, the bid structure is as follows:

$$b_i = (b_i(S))_{S \subset M} ; b_i(S) \geq 0 \;\; \forall \, S \subset M$$

The bids are opened by the auctioneer and the bidders are allocated to the bundles in such a way that sum of the valuations of all the allocated bundles is maximized. Each bidder pays to the auctioneer an amount equal to his marginal contribution to the trade.

In this example, the mechanism that is being used by the auctioneer is an indirect mechanism $\mathscr{M} = ((S_i)_{i \in N}, g(\cdot))$, where $S_i \subset (\mathbb{R}^+)^{2^m - 1}$ is the set of bids that bidder i can submit to the auctioneer and $g(\cdot)$ is the outcome rule given by $g(b) = ((y_i^*(S, b))_{i \in N, S \subset M}, t_1(b), \ldots, t_n(b))$, where $b = (b_1, \ldots, b_n)$. The functions $y_i^*(\cdot, \cdot)$ are known as *winner determination rules* and the functions $t_i(\cdot)$ are known as *payment rules*. The winner determination rule $y_i^*(\cdot, \cdot)$ for this auction is basically solution of the following optimization problem.

Maximize

$$\sum_{i=1}^{n} \sum_{S \subset M} b_i(S) y_i(S, b)$$

subject to

(i) $\sum_{S \subset M} y_i(S, b) \leq 1 \quad \forall \, i \in N$

(ii) $\sum_{S \subset M | j \in S} \sum_{i=1}^{n} y_i(S, b) \leq 1 \quad \forall \, j \in M$

(iii) $y_i(S, b) \in \{0, 1\} \quad \forall \, i \in N, S \subset M$

The payment rule $t_i(\cdot)$ for this auction is given by the following relation

$$t_i(b) = \sum_{j \neq i} v_j(k^*(b), b_j) - \sum_{j \neq i} v_j(k^*_{-i}(b_{-i}), b_j)$$

where $v_j(k^*(b), b_j) = \sum_{S \subset M} b_j(S) y_j^*(S, b)$ is total value of the bundle which is allocated to the bidder j. The quantity $v_j(k^*_{-i}(b_{-i}), b_j) = \sum_{S \subset M} b_j(S) y_j^*(S, b_{-i})$ is the total value of the bundle that will be allocated to the bidder $j \neq i$ if the bidder i were not present into the system. It is easy to verify that if set M consists of just one item then above winner determination and payment rule will precisely be the same as winner determination and payment rule of the Vickrey auction, therefore, the name Generalized Vickrey Auction.

4.5 Bayesian Game Induced by a Mechanism

In view of the definition of the indirect mechanism and direct revelation mechanism, we can say that a social planner can either use an indirect mechanism $\mathscr{M}$ or a direct mechanism $\mathscr{D}$ to elicit the information about the agents' preferences in an indirect or a direct manner, respectively. As we assumed earlier, all the agents are rational and intelligent. Therefore, after knowing about the mechanism $\mathscr{M} = ((S_i)_{i \in N}, g(\cdot))$ chosen by the social planner, each agent i starts doing an analysis regarding which action s_i will result in his most favorable outcome and comes up with a strategy $s_i : \Theta_i \to S_i$ to take the action. This phenomenon indeed leads to a game among the agents. A mechanism $\mathscr{M} = ((S_i)_{i \in N}, g(\cdot))$ combined with possible types of the agents $(\Theta_1, \ldots, \Theta_n)$, probability density $\phi(\cdot)$, and utility functions $(u_1(\cdot), \ldots, u_n(\cdot))$ defines a Bayesian game of incomplete information that gets induced among the agents when the social planner invokes this mechanism as a means to solve the information elicitation problem. The induced Bayesian game Γ^b is given in the following manner:

$$\Gamma^b = (N, (S_i)_{i \in N}, (\Theta_i)_{i \in N}, \phi(\cdot), (\overline{u_i})_{i \in N})$$

where $\overline{u_i} : C \times \Theta \to \mathbb{R}$ is the utility function of agent i and is defined in following manner

$$\overline{u_i}(c, \theta) = u_i(g(c), \theta_i)$$

where $C = \bigtimes_{i \in N} S_i$, and $\Theta = \bigtimes_{i \in N} \Theta_i$.

Knowing the fact that choosing any mechanism $\mathscr{M} = ((S_i)_{i \in N}, g(\cdot))$ will induce the game among the agents, and the agents will respond to it in a way suggested by the corresponding equilibrium strategy of the game, the social planner now worries about whether or not the outcome of the game matches with the outcome of the social choice function $f(\theta)$ (if all the agents had revealed their true types when asked directly). This notion is captured in the definition that follows.

5 Implementing a Social Choice Function

Definition 5.1 *We say that the mechanism $\mathscr{M} = ((S_i)_{i \in N}, g(\cdot))$ implements the social choice function $f(\cdot)$ if there is a pure strategy equilibrium $s^*(\cdot) = (s_1^*(\cdot), \ldots, s_n^*(\cdot))$ of the Bayesian game Γ^b induced by $\mathscr{M}$ such that $g(s_1^*(\theta_1), \ldots, s_n^*(\theta_n)) = f(\theta_1, \ldots, \theta_n) \ \forall (\theta_1, \ldots, \theta_n) \in \Theta$.*

The Figure 5 explains the idea behind what we mean by mechanism implementing a social choice function. Depending on the underlying equilibrium concept, two ways of implementing an SCF $f(\cdot)$ are standard in the literature.

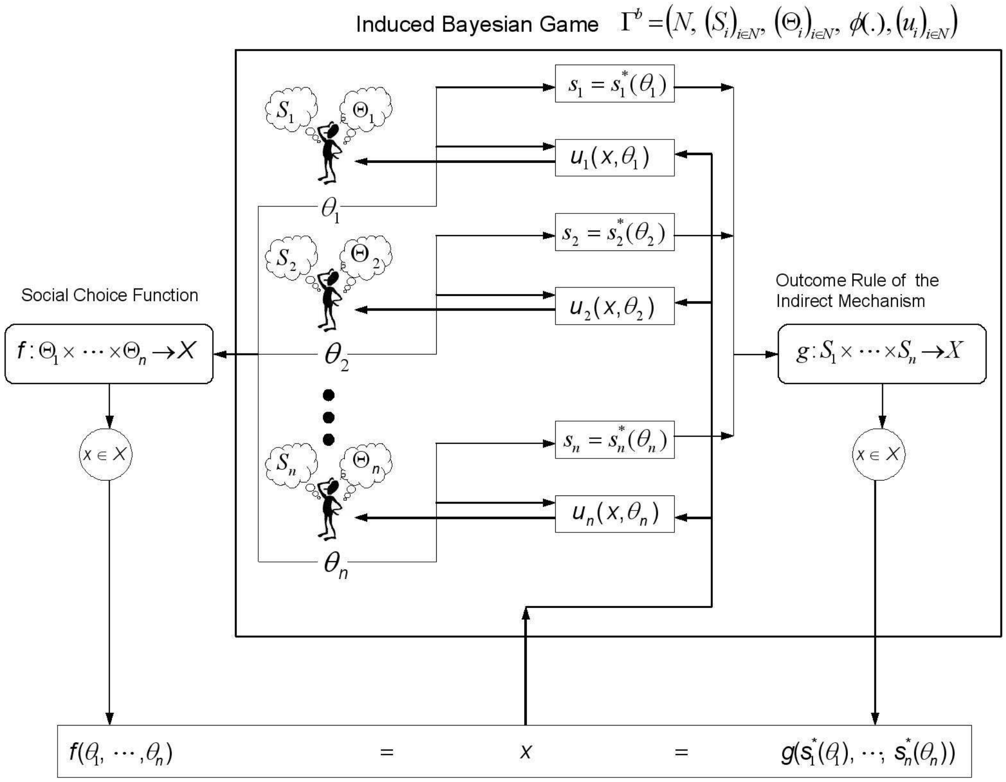

Figure 5: Mechanism $\mathscr{M} = ((S_i)_{i\in N}, g(\cdot))$ implements the social choice function $f(\cdot)$

5.1 Implementing a Social Choice Function in Dominant Strategy Equilibrium

First, we define the notion of a weakly dominant strategy equilibrium of the Bayesian game Γ^b.

Definition 5.2 (Weakly Dominant Strategy Equilibrium) *A pure strategy profile $s^d(\cdot) = \left(s_1^d(\cdot), \ldots, s_n^d(\cdot)\right)$ of the game Γ^b induced by the mechanism $\mathscr{M}$, is said to be a weakly dominant strategy equilibrium iff it satisfies the following condition.*

$$u_i(g(s_i^d(\theta_i), s_{-i}(\theta_{-i})), \theta_i) \geq u_i(g(s_i^{'}(\theta_i), s_{-i}(\theta_{-i})), \theta_i) \quad \forall i \in N, \forall \theta_i \in \Theta_i, \forall \theta_{-i} \in \Theta_{-i}, \forall s_i^{'}(\cdot) \in S_i, \forall s_{-i}(\cdot) \in S_{-i} \tag{1}$$

where S_i is the set of pure strategies of the agent i in the induced Bayesian game Γ^b, and S_{-i} is the set of pure strategy profiles of all the agents except agent i.

Definition 5.3 *We say that the mechanism $\mathscr{M} = ((S_i)_{i\in N}, g(\cdot))$ implements the social choice function $f(\cdot)$ in dominant strategy equilibrium if there is a weakly dominant strategy equilibrium $s^d(\cdot) = \left(s_1^d(\cdot), \ldots, s_n^d(\cdot)\right)$ of the game Γ^b induced by $\mathscr{M}$ such that*

$$g\left(s_1^d(\theta_1), \ldots, s_n^d(\theta_n)\right) = f\left(\theta_1, \ldots, \theta_n\right) \; \forall \left(\theta_1, \ldots, \theta_n\right) \in \Theta$$

5.2 Implementing a Social Choice Function in Bayesian Nash Equilibrium

Definition 5.4 *We say that the mechanism* $\mathscr{M} = ((S_i)_{i \in N}, g(\cdot))$ *implements the social choice function* $f(\cdot)$ *in Bayesian Nash equilibrium if there is a pure strategy Bayesian Nash equilibrium* $s^*(\cdot) = (s_1^*(\cdot), \ldots, s_n^*(\cdot))$ *of the game* Γ^b *induced by* $\mathscr{M}$ *such that*

$$g\left(s_1^*(\theta_1), \ldots, s_n^*(\theta_n)\right) = f\left(\theta_1, \ldots, \theta_n\right) \; \forall \left(\theta_1, \ldots, \theta_n\right) \in \Theta$$

Following is the definition of Bayesian Nash equilibrium of the Bayesian game Γ^b.

Definition 5.5 (Bayesian Nash Equilibrium) *A pure strategy profile* $s^*(\cdot) = (s_1^*(\cdot), \ldots, s_n^*(\cdot))$ *of the game* Γ^b *induced by the mechanism* $\mathscr{M}$*, is a Bayesian Nash equilibrium iff it satisfies the following condition.*

$$\begin{aligned} E_{\theta_{-i}}[u_i(g(s_i^*(\theta_i), s_{-i}^*(\theta_{-i})), \theta_i)|\theta_i] \quad \geq \quad & E_{\theta_{-i}}[u_i(g(s_i^{'}(\theta_i), s_{-i}^*(\theta_{-i})), \theta_i)|\theta_i] \\ & \forall i \in N, \forall \theta_i \in \Theta_i, \forall s_i^{'}(\cdot) \in S_i \end{aligned} \tag{2}$$

Following is a proposition that establishes the relationship between two equilibrium concepts defined above. The proof is straightforward.

Proposition 5.1 *A weakly dominant strategy equilibrium* $s^d(\cdot) = \left(s_1^d(\cdot), \ldots, s_n^d(\cdot)\right)$ *of the Bayesian game* Γ^b *induced by* $\mathscr{M}$*, is always a pure strategy Bayesian Nash equilibrium of the same Bayesian game* Γ^b*. If the condition (2) holds for all* $s_{-i}(\cdot) \in S_{-i}$ *also, then the proposition holds in the other direction also.*

Corollary 5.1 *If the mechanism* $\mathscr{M} = ((S_i)_{i \in N}, g(\cdot))$ *implements the social choice function* $f(\cdot)$ *in dominant strategy equilibrium, then it also implements* $f(\cdot)$ *in Bayesian Nash equilibrium.*

In what follows, we offer a few caveats to these definitions of implementing a social choice function.

1. The game Γ^b induced by the mechanism $\mathscr{M}$ may have more than one equilibrium, but the above definition requires only that *one of them* induces outcomes in accordance with the SCF $f(\cdot)$. Implicitly, then, the above definition assumes that, if multiple equilibria exist, the agents will play the equilibrium that the mechanism designer (social planner) wants.

2. Another implicit assumption of the above definition is that the game induced by the mechanism is a simultaneous move game, that is all the agents, after learning their types, choose their actions simultaneously. However, it is quite possible that the mechanism forces some agent(s) to lead the show, taking the action first followed by the actions of the remaining agents. In such a case, the game induced by the mechanism becomes a *Stackelberg game.*

6 Properties of a Social Choice Function

We have seen that a mechanism provides a solution to both the problem of information elicitation and the problem of preferences aggregation if it can implement the desired social choice function $f(\cdot)$. It is obvious that some SCFs are implementable and some are not. Before we look into the question of characterizing the space of implementable social choice functions, it is important to know which social choice function ideally a social planner would prefer to be implemented. In this section, we highlight a few properties of an SCF which ideally a social planner would wish the SCF to have.

Note that the fundamental characteristic of a social planner is that he is neutral to all the agents. Therefore, it is obvious for the social planner to be concerned about whether the outcome $f(\theta_1, \ldots, \theta_n)$ is socially fair or not. For this, a social planner would always like to use an SCF $f(\cdot)$ which satisfies as many desirable properties from the perspective of fairness as possible. A few important properties, which ideally a social planner would want an SCF $f(\cdot)$ to satisfy, are the following.

6.1 Ex-Post Efficiency

Definition 6.1 *The SCF $f : \Theta \to X$ is said to be ex-post efficient (or Paretian) if for any profile of agents' types $\theta = (\theta_1, \ldots, \theta_n)$, and any pair of alternatives $x, y \in X$, such that $u_i(x, \theta_i) \geq u_i(y, \theta_i)\ \forall\ i$ and $u_i(x, \theta_i) > u_i(y, \theta_i)$ for some i, we have $y \neq f(\theta_1, \ldots, \theta_n)$.*

An alternative definition of ex-post efficiency can be given in the following manner.

Definition 6.2 *The SCF $f : \Theta \to X$ is said to be ex-post efficient if for no profile of agents' type $\theta = (\theta_1, \ldots, \theta_n)$ does there exist an $x \in X$ such that $u_i(x, \theta_i) \geq u_i(f(\theta), \theta_i)\ \forall\ i$ and $u_i(x, \theta_i) > u_i(f(\theta), \theta_i)$ for some i.*

6.2 Non-Dictatorial SCF

We define this through a dictatorial SCF.

Definition 6.3 *The SCF $f : \Theta \to X$ is said to be dictatorial if for every profile of agents' type $\theta = (\theta_1, \ldots, \theta_n)$, we have $f(\theta_1, \ldots, \theta_n) \in \{x \in X | u_d(x, \theta_d) \geq u_d(y, \theta_d)\ \forall y \in X\}$, where d is a particular agent known as dictator.*

An SCF $f : \Theta \to X$ is said to be non-dictatorial if it is not dictatorial.

6.3 Incentive Compatibility (IC)

Definition 6.4 *The SCF $f(\cdot)$ is said to be incentive compatible (or truthfully implementable) if the direct revelation mechanism $\mathscr{D} = ((\Theta_i)_{i \in N}, f(\cdot))$ has a pure strategy equilibrium $s^*(\cdot) = (s_1^*(\cdot), \ldots, s_n^*(\cdot))$ in which $s_i^*(\theta_i) = \theta_i, \forall \theta_i \in \Theta_i, \forall i \in N$*

That is, truth telling by each agent constitutes an equilibrium of the game induced by $\mathscr{D}$. It is easy to verify that if an SCF $f(\cdot)$ is incentive compatible then the direct revelation mechanism $\mathscr{D} = ((\Theta_i)_{i \in N}, f(\cdot))$ can implement it. That is, directly asking the agents to report their types and plugging this information in $f(\cdot)$ to get the social outcome will solve both the problem of information elicitation and the problem of preferences aggregation.

Based on the type of equilibrium concept used, two types of incentive compatibility are given below.

6.4 Dominant Strategy Incentive Compatibility (DSIC)

Definition 6.5 *The SCF $f(\cdot)$ is said to be dominant strategy incentive compatible (or truthfully implementable in dominant strategies)*[1] *if the direct revelation mechanism $\mathscr{D} = ((\Theta_i)_{i \in N}, f(\cdot))$ has a* dominant strategy equilibrium *$s^d(\cdot) = (s_1^d(\cdot), \ldots, s_n^d(\cdot))$ in which $s_i^d(\theta_i) = \theta_i, \forall \theta_i \in \Theta_i, \forall i \in N$.*

[1] Strategy-proof, cheat-proof, straightforward are the alternative phrases used for this property.

That is, truth telling by each agent constitutes a dominant strategy equilibrium of the game induced by $\mathscr{D}$. Following is a necessary and sufficient condition for an SCF $f(\cdot)$ to be dominant strategy incentive compatible:

$$u_i\left(f\left(\theta_i, \theta_{-i}\right), \theta_i\right) \geq u_i(f(\hat{\theta}_i, \theta_{-i}), \theta_i), \forall i \in N, \forall \theta_i \in \Theta_i, \forall \theta_{-i} \in \Theta_{-i}, \forall \hat{\theta}_i \in \Theta_i \tag{3}$$

The above condition says that if the SCF $f(.)$ is DSIC, then, irrespective of what the other agents are doing, it is always in the best interest of agent i to report his true type θ_i.

6.5 Bayesian Incentive Compatibility (BIC)

Definition 6.6 *The SCF $f(\cdot)$ is said to be Bayesian incentive compatible (or truthfully implementable in Bayesian Nash equilibrium) if the direct revelation mechanism $\mathscr{D} = ((\Theta_i)_{i\in N}, f(\cdot))$ has a* Bayesian Nash equilibrium $s^*(\cdot) = (s_1^*(\cdot), \ldots, s_n^*(\cdot))$ *in which $s_i^*(\theta_i) = \theta_i, \forall \theta_i \in \Theta_i, \forall i \in N$.*

That is, truth telling by each agent constitutes a Bayesian Nash equilibrium of the game induced by $\mathscr{D}$. Following is a necessary and sufficient condition for an SCF $f(\cdot)$ to be Bayesian incentive compatible:

$$E_{\theta_{-i}}\left[u_i\left(f\left(\theta_i, \theta_{-i}\right), \theta_i\right) | \theta_i\right] \geq E_{\theta_{-i}}[u_i(f(\hat{\theta}_i, \theta_{-i}), \theta_i) | \theta_i], \forall i \in N, \forall \theta_i \in \Theta_i, \forall \hat{\theta}_i \in \Theta_i \tag{4}$$

The following proposition illustrates the relationship between these two notions of incentive compatibility of a social choice function. The proof of this proposition is quite straightforward.

Proposition 6.1 *If a social choice function $f(\cdot)$ is dominant strategy incentive compatible then it is also Bayesian incentive compatible.*

7 The Revelation Principle

This is one of the most fundamental results in the theory of mechanism design. This principle basically illustrates the relationship between an indirect mechanism $\mathscr{M}$ and a direct revelation mechanism $\mathscr{D}$ for any SCF $f(\cdot)$. This result enables us to restrict our inquiry about truthful implementation of an SCF to the class of direct revelation mechanisms only.

7.1 The Revelation Principle for Dominant Strategy Equilibrium

Proposition 7.1 *Suppose that there exists a mechanism $\mathscr{M} = (S_1, \ldots, S_n, g(\cdot))$ that implements the social choice function $f(\cdot)$ in dominant strategy equilibrium. Then $f(\cdot)$ is truthfully implementable in dominant strategy equilibrium (dominant strategy incentive compatible).*

Proof: If $\mathscr{M} = (S_1, \ldots, S_n, g(\cdot))$ implements $f(\cdot)$ in dominant strategy, then there exists a profile of strategies $s^d(\cdot) = \left(s_1^d(\cdot), \ldots, s_n^d(\cdot)\right)$ such that

$$g\left(s_1^d(\theta_1), \ldots, s_n^d(\theta_n)\right) = f\left(\theta_1, \ldots, \theta_n\right) \ \forall \left(\theta_1, \ldots, \theta_n\right) \in \Theta \tag{5}$$

and

$$\begin{aligned} u_i(g(s_i^d(\theta_i), s_{-i}(\theta_{-i})), \theta_i) \ \geq \ & u_i(g(s_i^{'}(\theta_i), s_{-i}(\theta_{-i})), \theta_i) \\ & \forall i \in N, \forall \theta_i \in \Theta_i, \forall \theta_{-i} \in \Theta_{-i}, \forall s_i^{'}(\cdot) \in S_i, \forall s_{-i}(\cdot) \in S_{-i} \end{aligned} \tag{6}$$

Condition (6) implies, in particular, that

$$u_i(g(s_i^d(\theta_i), s_{-i}^d(\theta_{-i})), \theta_i) \geq u_i(g(s_i^d(\hat{\theta}_i), s_{-i}^d(\theta_{-i})), \theta_i)$$
$$\forall i \in N, \forall \theta_i \in \Theta_i, \forall \hat{\theta}_i \in \Theta_i, \forall \theta_{-i} \in \Theta_{-i} \quad (7)$$

Conditions (5) and (7) together implies that

$$u_i\left(f\left(\theta_i, \theta_{-i}\right), \theta_i\right) \geq u_i(f(\hat{\theta}_i, \theta_{-i}), \theta_i), \forall i \in N, \forall \theta_i \in \Theta_i, \forall \theta_{-i} \in \Theta_{-i}, \forall \hat{\theta}_i \in \Theta_i$$

But this is precisely condition (3), the condition for $f(\cdot)$ to be truthfully implementable in dominant strategies.

Q.E.D.

The idea behind the revelation principle can be understood with the help of Figure 6. The set I_{DSE}

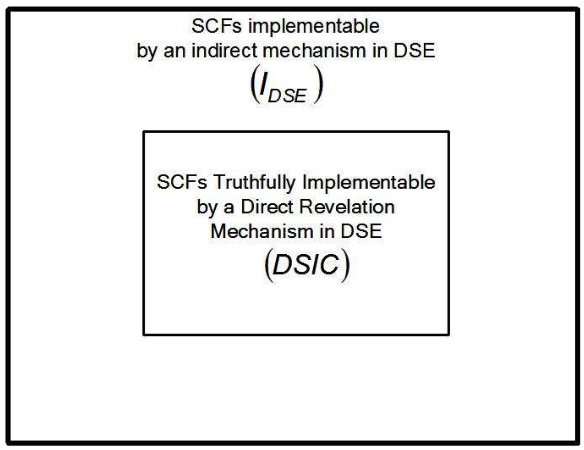

Figure 6: Revelation principle for dominant strategy equilibrium

consists of all the social choice functions that are implementable by some indirect mechanism $\mathscr{M}$ in dominant strategy equilibrium. The set $DSIC$ consists of all the social choice functions that are truthfully implementable by some direct mechanism $\mathscr{D}$ in dominant strategy equilibrium. Recall that a direct mechanism $\mathscr{D}$ can also be viewed as an indirect mechanism. Therefore, it is obvious that we must have

$$DSIC \subset I_{DSE} \quad (8)$$

In view of relation 8, we can say that the revelation principle for dominant strategy equilibrium basically says that $I_{DSE} \subset DSIC$, which further implies that $I_{DSE} = DSIC$.

Thus, the revelation principle for dominant strategy equilibrium says that an SCF $f(\cdot)$ is implementable by an indirect mechanism $\mathscr{M} = (S_1, \ldots, S_n, g(\cdot))$ in dominant strategy equilibrium iff it is truthfully implementable by the mechanism $\mathscr{D} = ((\Theta_i)_{i \in N}, f(\cdot))$ in dominant strategy equilibrium.

7.2 The Revelation Principle for Bayesian Nash Equilibrium

Proposition 7.2 *Suppose that there exists a mechanism $\mathscr{M} = (S_1, \ldots, S_n, g(\cdot))$ that implements the social choice function $f(\cdot)$ in Bayesian Nash equilibrium. Then $f(\cdot)$ is truthfully implementable in Bayesian Nash equilibrium (Bayesian incentive compatible).*

Proof: If $\mathscr{M} = (S_1, \ldots, S_n, g(\cdot))$ implements $f(\cdot)$ in Bayesian Nash equilibrium, then there exists a profile of strategies $s^*(\cdot) = (s_1^*(\cdot), \ldots, s_n^*(\cdot))$ such that

$$g\left(s_1^*(\theta_1), \ldots, s_n^*(\theta_n)\right) = f\left(\theta_1, \ldots, \theta_n\right) \ \forall \left(\theta_1, \ldots, \theta_n\right) \in \Theta \tag{9}$$

and

$$\begin{aligned} E_{\theta_{-i}}\left[u_i(g(s_i^*(\theta_i), s_{-i}^*(\theta_{-i})), \theta_i)|\theta_i\right] \ &\geq \ E_{\theta_{-i}}\left[u_i(g(s_i^{'}(\theta_i), s_{-i}^*(\theta_{-i})), \theta_i)|\theta_i\right] \\ &\forall i \in N, \forall \theta_i \in \Theta_i, \forall s_i^{'}(\cdot) \in S_i \end{aligned} \tag{10}$$

Condition (10) implies, in particular, that

$$\begin{aligned} E_{\theta_{-i}}\left[u_i(g(s_i^*(\theta_i), s_{-i}^*(\theta_{-i})), \theta_i)|\theta_i\right] \ &\geq \ E_{\theta_{-i}}\left[u_i(g(s_i^*(\hat{\theta}_i), s_{-i}^*(\theta_{-i})), \theta_i)|\theta_i\right] \\ &\forall i \in N, \forall \theta_i \in \Theta_i, \forall \hat{\theta}_i \in \Theta_i \end{aligned} \tag{11}$$

Conditions (9) and (11) together implies that

$$E_{\theta_{-i}}\left[u_i\left(f\left(\theta_i, \theta_{-i}\right), \theta_i\right)|\theta_i\right] \geq E_{\theta_{-i}}\left[u_i(f(\hat{\theta}_i, \theta_{-i}), \theta_i)|\theta_i\right], \forall i \in N, \forall \theta_i \in \Theta_i, \forall \hat{\theta}_i \in \Theta_i$$

But this is precisely condition (4), the condition for $f(\cdot)$ to be truthfully implementable in Bayesian Nash equilibrium.

Q.E.D.

In a way similar to the revelation principle for dominant strategy equilibrium, the revelation principle for Bayesian Nash equilibrium can be explained with the help of Figure 7. Following the similar line

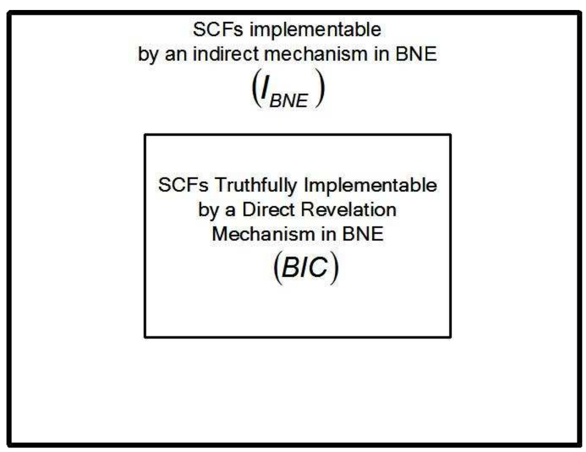

Figure 7: Revelation principle for Bayesian Nash equilibrium

of arguments, we can get the following relation for this case also

$$BIC \subset I_{BNE} \tag{12}$$

In view of relation 12, we can say that the revelation principle for Bayesian Nash equilibrium basically says that $I_{BNE} \subset BIC$, which further implies that $I_{BNE} = BIC$.

Thus, the revelation principle for Bayesian Nash equilibrium says that an SCF $f(\cdot)$ is implementable by an indirect mechanism $\mathscr{M} = (S_1, \ldots, S_n, g(\cdot))$ in Bayesian Nash equilibrium iff it is truthfully implementable by the mechanism $\mathscr{D} = ((\Theta_i)_{i \in N}, f(\cdot))$ in Bayesian Nash equilibrium.

In view of the above revelation principle, from now onwards, we will be just focusing on direct revelation mechanisms without loss of any generality.

8 The Gibbard-Satterthwaite Impossibility Theorem

Ideally, a social planner would prefer to implement a social choice function $f(\cdot)$ which is *ex-post efficient, non-dictatorial,* and *dominant strategy incentive compatible.* Now the question is: does there exist any such social choice function? The answer is *no*.

The Gibbard-Satterthwaite impossibility theorem shows that for a very general class of problems there is no hope of implementing any satisfactory social choice function in dominant strategies. This is an important landmark result which has shaped the course of research on incentives and implementation to a great extent and it was discovered independently by Gibbard in 1973 [40] and Satterthwaite in 1975 [41]. To understand the precise statement of this theorem, we need to build up a few concepts.

8.1 Preference Relation

In the mechanism design problem, the preference behavior of an agent i, over the set of outcomes X, can be summarized in the form of a *preference relation*, which is denoted by $\succsim_i$. The preference relation $\succsim_i$ is a binary relation on the set of alternatives X. We read $x \succsim_i y$ as "x is at least as preferred to agent i as y."

Definition 8.1 (Rational Preference Relation) *We say that the preference relation $\succsim_i$ is rational if it possesses the following three properties:*

1. *Reflexivity: For all $x \in X$, we have $x \succsim_i x$*
2. *Completeness: For all $x, y \in X$, we have that $x \succsim_i y$ or $y \succsim_i x$ (or both).*
3. *Transitivity: For all $x, y, z \in X$, if $x \succsim_i y$ and $y \succsim_i z$, then $x \succsim_i z$.*

Definition 8.2 (Strict-total Preference Relation) *We say that the preference relation $\succsim_i$ is strict-total if it possesses the following four properties:*

1. *Reflexivity: For all $x \in X$, we have $x \succsim_i x$*
2. *Completeness: For all $x, y \in X$, we have that $x \succsim_i y$ or $y \succsim_i x$ (or both).*
3. *Transitivity: For all $x, y, z \in X$, if $x \succsim_i y$ and $y \succsim_i z$, then $x \succsim_i z$.*
4. *Antisymmetry: For any $x, y \in X$ such that $x \neq y$, we have either $x \succsim_i y$ or $y \succsim_i x$ but not both.*

The set of all the rational preference relations and strict-total preference relations over the set X are denoted by $\mathscr{R}$ and $\mathscr{S}$, respectively. It is easy to see that $\mathscr{S} \subset \mathscr{R}$. The strict-total order relation is also known as *linear order relation* because it satisfies the properties of the usual "greater than or equal to" order in the real line.

8.2 Utility Function and Preference Relation

We have already seen that a given preference of an agent i, over the outcome set X, can also be described by means of a *utility function* $u_i : X \to \mathbb{R}$ which assigns a numerical value to each element in X. A utility function u_i always induces a *unique* preference relation $\succsim_i$ on X which can be described in following manner

$$x \succsim_i y \Leftrightarrow u_i(x) \geq u_i(y)$$

The following proposition establishes the relationship between these two ways of expressing the preferences of an agent i over the set X - *preference relation and utility function.*

Proposition 8.1

1. *If a preference relation $\succsim_i$ over X can be induced by some utility function $u_i(\cdot)$, then it will be a rational preference relation.*

2. *For every preference relation $\succsim_i \in \mathscr{R}$, there need not exist a utility function which will induce it. However, it is true when the set X is finite.*

3. *For a given preference relation $\succsim_i \in \mathscr{R}$, if there exists a utility function which induces it, then this need not be the unique one. Indeed, if the utility function $u_i(\cdot)$ induces $\succsim_i$, then $v(x) = f(u(x))$ is another utility function which will also induce $\succsim_i$, where $f : \mathbb{R} \to \mathbb{R}$ is a strictly increasing function.*

8.3 Set of Ordinal Preference Relations

In the mechanism design problem, for agent i, the preference over the set X is described in the form of a utility function $u_i : X \times \Theta_i \to \mathbb{R}$. That is, for every possible type $\theta_i \in \Theta_i$ of agent i, we can define a utility function $u_i(., \theta_i)$ over the set X. Let this utility function induce a rational preference relation $\succsim_i (\theta_i)$ over X. The set $\mathscr{R}_i = \{\succsim_i : \succsim_i = \succsim_i (\theta_i) \text{for some } \theta_i \in \Theta_i\}$ is known as the set of ordinal preference relations for agent i. It is easy to see that $\mathscr{R}_i \subset \mathscr{R} \;\forall\, i = 1, \ldots, n$.

Now we can state the Gibbard-Satterthwaite impossibility theorem without the proof. For proof, refer to Proposition 23.C.3 of [38].

Theorem 8.1 (Gibbard-Satterthwaite Impossibility Theorem) *Suppose that*

1. *The outcome set X is finite and contains at least three elements*

2. $\mathscr{R}_i = \mathscr{S} \;\forall\, i = 1, \ldots, n$

3. *$f(\Theta) = X$, that is, the image of SCF $f(\cdot)$ is the set X.*

Then the social choice function $f(\cdot)$ is truthfully implementable in dominant strategies iff it is dictatorial.

The Gibbard-Satterthwaite impossibility theorem gives a disappointing piece of news and the question facing a social planner is what kind of SCF to look for in the face of this impossibility result. There are two possible routes that one may take:

1. The first route is to focus on some restricted environment where at least one of the three requirements of the Gibbard-Satterthwaite impossibility theorem is not fulfilled. *Quasi-linear environment* is one such environment where the second condition of this theorem is not satisfied and in fact all the social choice functions in such environments are non-dictatorial.

2. The other route is to weaken the implementation concept and look for an SCF which is *ex-post efficient, non-dictatorial,* and *Bayesian incentive compatible.*

We now elaborate on each of these two routes.

9 Quasi-Linear Environments

This is a special and much studied class of environments where the Gibbard-Satterthwaite theorem does not hold. In this environment, an alternative $x \in X$ is a vector of the form $x = (k, t_1, \ldots, t_n)$, where k is an element of a set K, to be called as "project choice." The set K is a compact subset of a topological space. $t_i \in \mathbb{R}$ is a monetary transfer to agent i. If $t_i > 0$ then agent i will receive the money and if $t_i < 0$ then agent i will pay the money. We assume that we are dealing with a closed system in which n agents have no outside source of funding, i.e. $\sum_{i=1}^{n} t_i \leq 0$. This condition is known as *weak budget balance* condition. The set of alternatives X is therefore

$$X = \left\{ (k, t_1, \ldots, t_n) : k \in K, t_i \in \mathbb{R} \; \forall \; i = 1, \ldots, n, \sum_i t_i \leq 0 \right\}$$

A social choice function in this quasi-linear environment takes the form $f(\theta) = (k(\theta), t_1(\theta), \ldots, t_n(\theta))$ where $\forall \; \theta \in \Theta, \; k(\theta) \in K$ and $\sum_i t_i(\theta) \leq 0$.[2] For a direct revelation mechanism $\mathscr{D} = ((\Theta_i)_{i \in N}, f(\cdot))$ in this environment, the agent i's utility function takes the quasi-linear form

$$u_i(x, \theta_i) = u_i(k, t_1, \ldots, t_n, \theta_i) = v_i(k, \theta_i) + m_i + t_i$$

where m_i is agent i's initial endowment of the money and the function $v_i(\cdot)$ is known as agent i's valuation function. Recall condition 5 in the definition of mechanism design problem given in Section 2. This condition says that the utility function $u_i(\cdot)$ is common knowledge. In the context of a quasi-linear environment, this implies that for any given type θ_i of any agent i, the social planner and every other agent j have a way to figure out the function $v_i(., \theta_i)$. In many cases, the set Θ_i of the direct revelation mechanism $\mathscr{D} = ((\Theta_i)_{i \in N}, f(\cdot))$ is actually the set of all feasible valuation functions $v_i : K \to \mathbb{R}$ of agent i. That is, each possible function represents each possible type of agent i. Therefore, in such settings reporting a type is the same as reporting a valuation function.

As far as examples of quasi-linear environment are concerned, all the previously discussed examples, such as fair bilateral trade (Example 4.1), first price auction (Example 4.2), second-price auction (Example 4.3), and Generalized Vickrey Auction (Example 4.4) are all natural examples of the mechanism in quasi-linear environment.

In the quasi-linear environment, we can define two important properties of a social choice function.

Definition 9.1 (Allocative Efficiency (AE)) *We say that SCF $f(\cdot) = (k(\cdot), t_1(\cdot), \ldots, t_n(\cdot))$ is allocatively efficient if for each $\theta \in \Theta$, $k(\theta)$ satisfies the following condition*[3]

$$k(\theta) \in \underset{k \in K}{\arg\max} \sum_{i=1}^{n} v_i(k, \theta_i) \tag{13}$$

[2] Note that here we are using symbol k for both as an element of the set K and as a function going from Θ to K. It should be clear from the context as to which of these two we are referring.

[3] *We will keep using the symbol $k^*(\cdot)$ for a function $k(\cdot)$ that satisfies the Equation (13).*

Note that the above definition will make sense only when we ensure that for any given θ, the function $\sum_{i=1}^{n} v_i(.,\theta_i) : K \to \mathbb{R}$ attains a maximum over the set K. The simplest way to do this is to put a restriction that the function $v_i(.,\theta_i) : K \to \mathbb{R}$ is an upper semi-continuous function for each $\theta_i \in \Theta_i$ and for each $i = 1, \ldots, n$.

Definition 9.2 (Budget Balance (BB)) *We say that SCF $f(\cdot) = (k(\cdot), t_1(\cdot), \ldots, t_n(\cdot))$ is budget balanced if for each $\theta \in \Theta$, $t_1(\theta), \ldots, t_n(\theta)$ satisfy the following condition*[4]

$$\sum_{i=1}^{n} t_i(\theta) = 0 \tag{14}$$

The following lemma establishes an important relationship of these two properties of an SCF with the ex-post efficiency of the SCF.

Lemma 9.1 *A social choice function $f(\cdot) = (k(\cdot), t_1(\cdot), \ldots, t_n(\cdot))$ is ex-post efficient in quasi-linear environment if and only if it is allocatively efficient and budget balanced.*

Proof: Let us assume that $f(\cdot) = (k(\cdot), t_1(\cdot), \ldots, t_n(\cdot))$ is allocatively efficient and budget balanced. This implies that for any $\theta \in \Theta$, we have

$$\begin{aligned}
\sum_{i=1}^{n} u_i(f(\theta), \theta_i) &= \sum_{i=1}^{n} v_i(k(\theta), \theta_i) + \sum_{i=1}^{n} t_i(\theta) \\
&= \sum_{i=1}^{n} v_i(k(\theta), \theta_i) + 0 \\
&\geq \sum_{i=1}^{n} v_i(k, \theta_i) + \sum_{i=1}^{n} t_i \;\forall\; x = (k, t_1, \ldots, t_n) \\
&= \sum_{i=1}^{n} u_i(x, \theta_i) \;\forall\; (k, t_1, \ldots, t_n) \in X
\end{aligned}$$

That is if the SCF is allocatively efficient and budget balanced then for any type profile θ of the agent the outcome chosen by the social choice function will be such that it maximizes the total utility derived by all the agents. This will automatically imply that the SCF is ex-post efficient.

To prove the other part, we will first show that if $f(\cdot)$ is not allocatively efficient then it cannot be ex-post efficient and next we will show that if $f(\cdot)$ is not budget balanced then it cannot be ex-post efficient. These two facts together will imply that if $f(\cdot)$ is ex-post efficient then it will have to be allocatively efficient and budget balanced, thus completing the proof of the lemma.

To start with, let us assume that $f(\cdot)$ is not allocatively efficient. This means that $\exists\, \theta \in \Theta$, and $k \in K$ such that

$$\sum_{i=1}^{n} v_i(k, \theta_i) > \sum_{i=1}^{n} v_i(k(\theta), \theta_i)$$

This implies that there exists at least one agent j for whom $v_j(k, \theta_i) > v_j(k(\theta), \theta_i)$. Now consider the following alternative x

$$x = \Big(k, (t_i = t_i(\theta) + v_i(k(\theta), \theta_i) - v_i(k, \theta_i))_{i \neq j}, t_j = t_j(\theta)\Big)$$

[4] *Many authors prefer to call this property as strong budget balance and they refer the property of having $\sum_{i=1}^{n} t_i(\theta) \leq 0$ as weak budget balance. In this thesis, we will use the term budget balance to refer to strong budget balance.*

It is easy to verify that $u_i(x, \theta_i) = u_i(f(\theta), \theta_i) \ \forall \ i \neq j$ and $u_j(x, \theta_i) > u_j(f(\theta), \theta_i)$ implying that $f(\cdot)$ is not ex-post efficient.

Next, we assume that $f(\cdot)$ is not budget balanced. This means that there exists at least one agent j for whom $t_j(\theta) < 0$. Let us consider the following alternative x

$$x = \Big(k, (t_i = t_i(\theta))_{i \neq j}, t_j = 0\Big)$$

It is easy to verify that for the above alternative x, we have $u_i(x, \theta_i) = u_i(f(\theta), \theta_i) \ \forall \ i \neq j$ and $u_j(x, \theta_i) > u_j(f(\theta), \theta_i)$ implying that $f(\cdot)$ is not ex-post efficient.

Q.E.D.

The next lemma summarizes another fact about the social choice function in quasi-linear environment.

Lemma 9.2 *All social choice functions in quasi-linear environment are non-dictatorial.*

Proof: If possible assume that SCF $f(\cdot)$ is dictatorial. This means that for each $\theta \in \Theta$, we have

$$u_d(f(\theta), \theta_d) \geq u_d(x, \theta_d) \ \forall \ x \in X$$

where d is a dictatorial agent. However, because of quasi-linear environment, we have $u_d(f(\theta), \theta_d) = v_d(k(\theta), \theta_d) + t_d(\theta)$. Now consider the following alternative $x \in X$:

$$x = \begin{cases} (k(\theta), (t_i = t_i(\theta))_{i \neq d}, t_d = t_d(\theta) - \sum_{i=1}^{n} t_i(\theta)) & : \ \sum_{i=1}^{n} t_i(\theta) < 0 \\ (k(\theta), (t_i = t_i(\theta))_{i \neq d, j}, t_d = t_d(\theta) + \epsilon, t_j = t_j(\theta) - \epsilon) & : \ \sum_{i=1}^{n} t_i(\theta) = 0 \end{cases}$$

where $\epsilon > 0$ is any arbitrary number and j is any agent other than d. It is easy to verify that for the above constructed x, we have $u_d(x, \theta_d) > u_d(f(\theta), \theta_d)$ contradicting the assumption that $f(\cdot)$ is dictatorial.

Q.E.D.

In view of the Lemma 9.2, the social planner need not have to worry about the non-dictatorial property of the social choice function in quasi-linear environments and he can simply look for whether there exists any SCF which is both ex-post efficient and dominant strategy incentive compatible. Further in the light of Lemma 9.1, we can say that the social planner can look for an SCF which is allocatively efficient, budget balanced, and dominant strategy incentive compatible. Once again the question arises whether there exists any such social choice function which satisfies all these three properties - AE, BB, and DSIC. We explore the answer to this question in what follows.

10 Groves Mechanisms

The following theorem, due to Groves [42] confirms that in quasi-linear environment, there exist social choice functions which are both allocatively efficient and truthfully implementable in dominant strategies (dominant strategy incentive compatible).

Theorem 10.1 (Groves' Theorem) *Let the social choice function $f(\cdot) = (k^*(\cdot), t_1(\cdot), \ldots, t_n(\cdot))$ be allocatively efficient. This function can be truthfully implemented in dominant strategies if it satisfies the following payment structure (popularly known as Groves payment (incentive) scheme) :*

$$t_i(\theta) = \left[\sum_{j \neq i} v_j(k^*(\theta), \theta_j)\right] + h_i(\theta_{-i}) \ \ \forall \ i = 1, \ldots, n \tag{15}$$

where $h_i(\cdot)$ is any arbitrary function of θ_{-i} up to satisfying the feasibility condition $\sum_i t_i(\theta) \leq 0 \;\; \forall \, \theta \in \Theta$.

For proof of the Groves theorem, refer to Proposition 23.C.4 of [38]. Following are a few interesting implications of the above theorem.

1. Given the announcements θ_{-i} of agents $j \neq i$, agent i's transfer depends on his announced type only through his announcement's effect on the project choice $k^*(\theta)$.

2. The change in the monetary transfer of agent i when his type changes from θ_i to $\hat{\theta}_i$ is equal to the effect that the corresponding change in project choice has on total value of the rest of the agents. That is,

$$t_i(\theta_i, \theta_{-i}) - t_i(\hat{\theta}_i, \theta_{-i}) = \sum_{j \neq i} \left[v_j(k^*(\theta_i, \theta_{-i}), \theta_j) - v_j(k^*(\hat{\theta}_i, \theta_{-i}), \theta_j) \right]$$

   Another way of describing this is to say that the change in monetary transfer to agent i reflects exactly the externality he is imposing on the other agents.

After the famous result of Groves, a direct revelation mechanism in which the implemented SCF is allocatively efficient and satisfies the Groves payment scheme is called as *Groves Mechanism.*

Definition 10.1 (Groves Mechanisms) *A direct revelation mechanism, $\mathscr{D} = ((\Theta_i)_{i \in N}, f(\cdot))$ in which $f(\cdot) = (k(\cdot), t_1(\cdot), \ldots, t_n(\cdot))$ satisfies (13) and (15) is known as Groves mechanism.*[5]

In practice, Groves mechanisms are popularly known as Vickrey-Clarke-Groves (VCG) mechanisms because Clarke mechanism is a special case of Groves mechanism and Vickrey mechanism is a special case of Clarke mechanism. We will discuss this relationship later in this paper.

Groves theorem provides a sufficiency condition under which an allocatively efficient (AE) SCF will be DSIC. The following theorem due to Green and Laffont [43] provides a set of conditions under which the condition of Groves theorem also becomes a necessary condition for an AE SCF to be DSIC. In this theorem, we let $\mathscr{F}$ denote the set of all possible functions $f : K \to \mathbb{R}$.

Theorem 10.2 (First Characterization Theorem of Green-Laffont) *Suppose that for each agent $i = 1, \ldots, n$ $\{v_i(., \theta_i) : \theta_i \in \Theta_i\} = \mathscr{F}$; that is, every possible valuation function from K to $\mathbb{R}$ arises for some $\theta_i \in \Theta_i$. Then any allocatively efficient (AE) social choice function $f(\cdot)$ will be dominant strategy incentive compatible (DSIC) if and only if it satisfies the Groves payment scheme given by (15).*

Note that in the above theorem, every possible valuation function from K to $\mathbb{R}$ arises for some $\theta_i \in \Theta_i$. Therefore, in many cases, depending upon the structure of the compact set $\mathscr{K}$, it is quite possible that for some type profile $\theta = (\theta_1, \ldots, \theta_n)$, the maximum of the function $\sum_{i=1}^{n} v_i(., \theta_i)$ over the set $\mathscr{K}$ may not exist. In such cases, the set of AE social choice functions would be empty and the theorem will make no sense. One possible way to get away with this difficulty is to assume that the set $\mathscr{K}$ is a finite set. Another solution is to restrict the allowable valuation functions to the class of continuous functions. The following characterization theorem of Green and Laffont [43] fixes this problem by replacing the $\mathscr{F}$ with $\mathscr{F}_c$ where $\mathscr{F}_c$ denotes the set of all possible continuous functions $f : K \to \mathbb{R}$.

[5] *We will sometime abuse the terminology and simply refer to a SCF $f(\cdot)$ satisfying (13) and (15) as Groves mechanism.*

Theorem 10.3 (Second Characterization Theorem of Green-Laffont) *Suppose that for each agent $i = 1, \ldots, n$ $\{v_i(., \theta_i) : \theta_i \in \Theta_i\} = \mathscr{F}_c$; that is, every possible continuous valuation function from K to $\mathbb{R}$ arises for some $\theta_i \in \Theta_i$. Then any allocatively efficient (AE) social choice function $f(\cdot)$ will be dominant strategy incentive compatible (DSIC) if and only if it satisfies the Groves payment scheme given by (15).*

10.1 Groves Mechanisms and Budget Balance

Note that a Groves mechanism always satisfies the properties of AE and DSIC. Therefore, if a Groves mechanism is budget balanced then it will solve the problem of the social planner because it will then be ex-post efficient and dominant strategy incentive compatible. By looking at the definition of the Groves mechanism, one can conclude that it is the functions $h_i(\cdot)$ that decide whether or not the Groves mechanism is budget balanced. The natural question that arises now is whether there exists a way of defining functions $h_i(\cdot)$ such that the Groves mechanism is budget balanced. In what follows, we present one possibility and one impossibility result in this regard.

10.2 Possibility and Impossibility Results for Quasi-linear Environments

Green and Laffont [43] showed that in quasi-linear environment, if the set of possible types for each agent is sufficiently rich then ex-post efficiency and DSIC cannot be achieved together. The precise statement is given in the form of following theorem.

Theorem 10.4 (Green-Laffont Impossibility Theorem) *Suppose that for each agent $i = 1, \ldots, n$, $\{v_i(., \theta_i) : \theta_i \in \Theta_i\} = \mathscr{F}$; that is, every possible valuation function from K to $\mathbb{R}$ arises for some $\theta_i \in \Theta_i$. Then there is no social choice function which is ex-post efficient and DSIC.*

Thus, the above theorem says that if the set of possible types for each agent is sufficiently rich then there is no hope of finding a way to define the functions $h_i(\cdot)$ in Groves payment scheme so that we have $\sum_{i=1}^{n} t_i(\theta) = 0$. However, one special case in which a more positive result does obtain is summarized in the form of following possibility result.

Theorem 10.5 (Possibility Result for Budge Balance of Groves Mechanisms) *If there is at least one agent whose preferences are known (i.e. his type set is a singleton set) then it is possible to choose the functions $h_i(\cdot)$ so that $\sum_{i=1}^{n} t_i(\theta) = 0$.*

Proof Let agent i be such that his preferences are known, that is $\Theta_i = \{\theta_i\}$. In view of this condition, it is easy to see that for an allocatively efficient social choice function $f(\cdot) = (k^*(\cdot), t_1(\cdot), \ldots, t_n(\cdot))$, the allocation $k^*(\cdot)$ depends only on the types of the agents other than i. That is, the allocation $k^*(\cdot)$ is a mapping from Θ_{-i} to K. Let us define the functions $h_j(\cdot)$ in the following manner.

$$h_j(\theta_{-j}) = \begin{cases} h_j(\theta_{-j}) & : \quad j \neq i \\ -\sum_{r \neq i} h_r(\theta_{-r}) - (n-1)\sum_{r=1}^{n} v_r(k^*(\theta), \theta_r) & : \quad j = i \end{cases}$$

It is easy to see that under the above definition of the functions $h_i(\cdot)$, we will have $t_i(\theta) = -\sum_{j \neq i} t_j(\theta)$.

10.3 Clarke (Pivotal) Mechanisms: Special Case of Groves Mechanisms

A special case of Groves mechanism was discovered independently by Clarke in 1971 [44] and is known as *Clarke*, or *pivotal* mechanism. It is a special case of Groves mechanism in the sense of using a particular form for the function $h_i(\cdot)$. In Clarke mechanism, the function $h_i(\cdot)$ is given by the following relation

$$h_i(\theta_{-i}) = -\sum_{j \neq i} v_j(k^*_{-i}(\theta_{-i}), \theta_j) \;\; \forall \, \theta_{-i} \in \Theta_{-i}, \forall \, i = 1, \ldots, n \tag{16}$$

where $k^*_{-i}(\theta_{-i}) \in K_{-i}$ is the choice of a project which is allocatively efficient if there were only the $n-1$ agents $j \neq i$. Formally, $k^*_{-i}(\theta_{-i})$ must satisfy the following condition.

$$\sum_{j \neq i} v_j(k^*_{-i}(\theta_{-i}), \theta_j) \geq \sum_{j \neq i} v_j(k, \theta_j) \; \forall \, k \in K_{-i} \tag{17}$$

where the set K_{-i} is the set of project choices available when agent i is absent. Substituting the value of $h_i(\cdot)$ from Equation (16) in Equation (15), we get the following expression for agent i's transfer in the Clarke mechanism

$$t_i(\theta) = \left[\sum_{j \neq i} v_j(k^*(\theta), \theta_j)\right] - \left[\sum_{j \neq i} v_j(k^*_{-i}(\theta_{-i}), \theta_j)\right] \tag{18}$$

10.4 Clarke Mechanisms and Weak Budget Balance

Recall from the definition of Groves mechanisms that, for weak budget balance, we should choose the functions $h_i(\theta_{-i})$ in such a way that the weak budget balance condition $\sum_{i=1}^{n} t_i(\theta) \leq 0$ is satisfied. In this sense, the Clarke mechanism is a useful special-case because it achieves weak budget balance under fairly general settings. In order to understand these general sufficiency conditions, we define following quantities

$$\begin{aligned} B^*(\theta) &= \left\{ k \in K | k \in \begin{array}{c} \arg\max \\ k \in K \end{array} \sum_{j=1}^{n} v_j(k, \theta_j) \right\} \\ B^*(\theta_{-i}) &= \left\{ k \in K_{-i} | k \in \begin{array}{c} \arg\max \\ k \in K_{-i} \end{array} \sum_{j \neq i} v_j(k, \theta_j) \right\} \end{aligned}$$

where $B^*(\theta)$ is the set of project choices that are allocatively efficient when all the agents are present there in the system. Similarly, $B^*(\theta_{-i})$ is the set of project choices that are allocatively efficient if there were $n-1$ agents $j \neq i$. It is obvious that $k^*(\theta) \in B^*(\theta)$ and $k^*_{-i}(\theta_{-i}) \in B^*(\theta_{-i})$.

Using the above quantities, we define the following properties of a direct revelation mechanism in quasi-linear environment.

Definition 10.2 (No Single Agent Effect) *We say that mechanism $\mathscr{M}$ has no single agent effect if for each agent i, each $\theta \in \Theta$, and each $k^*(\theta) \in B^*(\theta)$, we have a $k \in K_{-i}$ such that*

$$\sum_{j \neq i} v_j(k, \theta_j) \geq \sum_{j \neq i} v_j(k^*(\theta), \theta_j)$$

In view of the above properties, we have the following proposition that gives sufficiency condition for Clarke mechanism to be weak budget balanced.

Proposition 10.1 *If the Clarke mechanism has no single agent effect, then transfer of each agent would be non-positive, that is, $t_i(\theta_i) \leq 0 \; \forall \; \theta \in \Theta; \; \forall \; i = 1, \ldots, n$. In such a situation, the Clarke mechanism would satisfy the weak budget balance property.*

Proof: Note that by virtue of no single agent effect, for each agent i, each $\theta \in \Theta$, and each $k^*(\theta) \in B^*(\theta)$, $\exists k \in K_{-i}$ such that

$$\sum_{j \neq i} v_j(k, \theta_j) \geq \sum_{j \neq i} v_j(k^*(\theta), \theta_j)$$

However, by definition of $k^*_{-i}(\theta_{-i})$, given by Equation (17), we have

$$\sum_{j \neq i} v_j(k^*_{-i}(\theta_{-i}), \theta_j) \;\; \geq \;\; \sum_{j \neq i} v_j(k, \theta_j) \; \forall \; k \in K_{-i}$$

Combining the above two facts, we get

$$\begin{aligned} \sum_{j \neq i} v_j(k^*_{-i}(\theta_{-i}), \theta_j) \;\; &\geq \;\; \sum_{j \neq i} v_j(k^*(\theta), \theta_j) \\ \Rightarrow 0 \;\; &\geq \;\; t_i(\theta) \\ \Rightarrow 0 \;\; &\geq \;\; \sum_{i=1}^{n} t_i(\theta) \end{aligned}$$

Q.E.D.

Now we can make following assertions with respect to these definitions. In what follows is an interesting corollary of the above proposition.

Corollary 10.1

1. $t_i(\theta) = 0$ *iff* $k^*(\theta) \in B^*(\theta_{-i})$. *That is, agent i's monetary transfer is zero iff his announcement does not change the project decision relative to what would be allocatively efficient for agents $j \neq i$ in isolation.*

2. $t_i(\theta) < 0$ *iff* $k^*(\theta) \notin B^*(\theta_{-i})$. *That is, agent i's monetary transfer is negative iff his announcement changes the project decision relative to what would be allocatively efficient for agents $j \neq i$ in isolation. In such situation, the agent i is known to be "pivotal" to the efficient project choice and he pays a tax equal to his effect on the other agents.*

The following picture summarizes the conclusions of this section by showing how the space of social choice functions looks like in quasi-linear environment. In what follows we discuss a few examples of direct revelation mechanisms in quasi-linear environment and study the various properties of the underlying SCF.

10.5 Fair Bilateral Trade is BB+AE but not DSIC

Consider the fair bilateral trade Example 4.1 in which the good is always allocated to the agent who values it the most. Therefore, the SCF is allocatively efficient. The SCF here is also budget balanced, which is quite clear from the definition of SCF itself. However, it is easy to see that this SCF is not dominant strategy incentive compatible. For example, let us assume that $\theta_1 = 50$ and $\theta_2 = 100$. Then, it is easy to see that for any $50 < \hat{\theta}_1 < 100$

$$u_1(f(\hat{\theta}_1, \theta_2), \theta_1) > u_1(f(\theta_1, \theta_2), \theta_1)$$

This violates the required condition of the dominant strategy incentive compatibility.

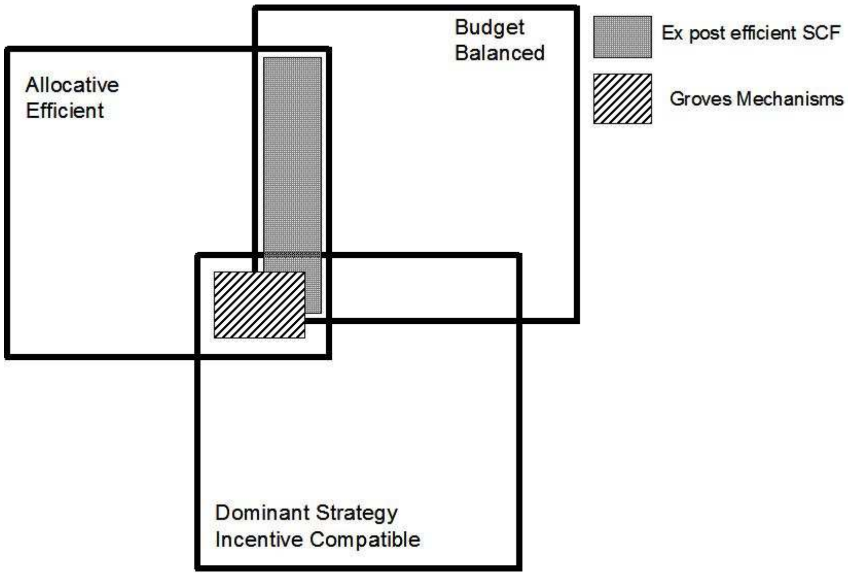

Figure 8: Space of social choice functions in quasi-linear environment

10.6 First-Price Sealed Bid Auction is AE but neither DSIC nor BB

Consider the Example 4.2 of first-price auction. Let us assume that $S_i = \Theta_i = [\underline{\theta_i}, \overline{\theta_i}] \quad \forall \, i \in N$. In such a case, the first-price auction becomes a direct revelation mechanism $\mathscr{D} = ((\Theta_i)_{i \in N}, f(\cdot))$, where $f(\cdot)$ is an SCF which is same as outcome rule of the first-price auction. Note that under the SCF $f(\cdot)$, the good is always allocated to the bidder who values it the most. Therefore, the SCF $f(\cdot)$ is allocatively efficient. It is an easy fact to note that the SCF $f(\cdot)$ is not budget balanced, because the auctioneer is not considered as one of the agents. Moreover, the SCF used here is not DSIC because truth telling is not a dominant strategy for the bidders. In order to show this, let us assume that there are two bidders and for some instance we have $\theta_1 = 50$ and $\theta_2 = 100$. Then, it is easy to see that for any $50 < \hat{\theta}_2 < 100$

$$u_2(f(\theta_1, \hat{\theta}_2), \theta_2) > u_2(f(\theta_1, \theta_2), \theta_2)$$

This violates the required condition of the dominant strategy incentive compatibility.

10.7 Second-Price Sealed Bid (Vickrey) Auction is AE + DSIC but not BB

Consider Example 4.3 of second-price (Vickrey) auction. Once again we assume that $S_i = \Theta_i = [\underline{\theta_i}, \overline{\theta_i}] \quad \forall \, i \in N$. In such a case, the second-price auction becomes a direct revelation mechanism $\mathscr{D} = ((\Theta_i)_{i \in N}, f(\cdot))$, where $f(\cdot)$ is an SCF which is same as outcome rule of the second-price auction. We have already shown that Vickrey auction is a special case of GVA in which the auctioneer is selling just single unit of a single item. Given that the Vickrey auction a special case of the GVA, we can assert that the SCF $f(\cdot)$ used by the auctioneer in Vickrey auction is AE + DSIC but not BB because the SCF used in GVA is also AE + DSIC and not BB, which we are going to show next.

10.8 Generalized Vickrey Auction is AE + DSIC but not BB

Consider Example 4.4 of generalized Vickrey auction. Let us assume that $S_i = \Theta_i = \left[\underline{V_i}, \overline{V_i}\right]^{(2^m-1)} \ \ \forall\, i \in N$. In such a case, the GVA auction becomes a direct revelation mechanism $\mathscr{D} = ((\Theta_i)_{i \in N}, f(\cdot))$, where $f(\cdot)$ is an SCF which is same as outcome rule of the GVA auction. It is easy to see that the under this SCF $f(\cdot)$, the bundles are allocated in such a way that total value of all the agents gets maximized. This implies that the SCF in this example is allocatively efficient. Also, note that the payment structure is the same as the Clarke mechanism. Thus, the direct revelation mechanism used by the auctioneer in GVA is a Clarke mechanism. Therefore, the SCF is automatically DSIC. However, note that SCF certainly not budget balanced because we are not including the auctioneer into the system. If auctioneer is also part of the system then by virtue of theorem 10.5, we can claim that the SCF would be BB also.

The Table 1 summarizes the properties of the SCFs discussed in above examples. Figure 9 summarizes the relationship among various SCFs discussed in above examples.

SCF	**AE**	**BB**	**DSIC**
Fair Bilateral Trade	✓	✓	×
First-Price Auction	✓	×	×
Vickrey Auction	✓	×	✓
GVA	✓	×	✓

Table 1: Properties of social choice functions in quasi-linear environment

11 A Variant of Quasi-linear Environment

In this section, we consider a variant of the quasi-linear environment which is commonly encountered in applications. In this environment, we make the following change in the settings of quasi-linear environment: we assume that $K = \mathbb{R}$. Note, this new set of project choices is no more a compact set. Next, we assume that for each agent i and each of his type $\theta_i \in \Theta_i$, we have

1. $v_i(., \theta_i) : K \to \mathbb{R}$ is a twice continuously differentiable function
2. $\frac{\partial^2 v_i(k,\theta_i)}{\partial k^2} < 0$
3. $\frac{\partial^2 v_i(k,\theta_i)}{\partial k\, \partial \theta_i} > 0$
4. $\theta_i \in \left[\underline{\theta}_i,\ \overline{\theta_i}\right] \subset \mathbb{R}$

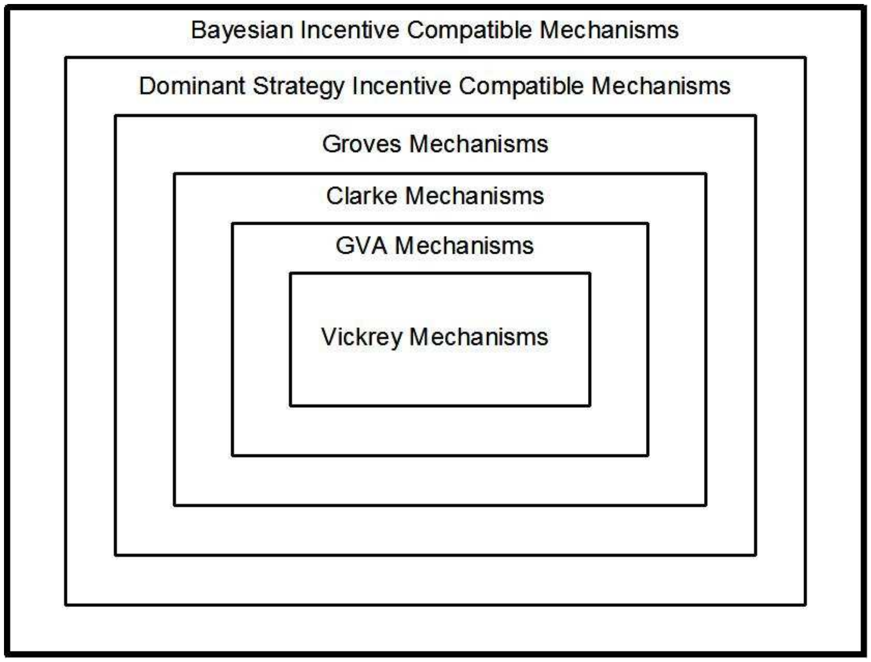

Figure 9: Space of BIC and DSIC social choice functions in quasi-linear environment

It is easy to verify that even in this modified environment, there is no special choice function which is dictatorial. We now define a special class $\mathscr{C}$ of the social choice functions in this environment.

$$\mathscr{C} = \left\{ f : \Theta \to \mathbb{R}^{n+1} \,|\, f \text{ is a continuously differentiable function} \right\}$$

The social choice functions of this class are commonly encountered in applications and they have some very interesting properties. In what follows we characterize the social choice functions of this class $\mathscr{C}$ from the perspective of various properties satisfied by them.

Proposition 11.1 (Characterization of AE SCFs) *If an SCF $f \in \mathscr{C}$ is AE then $\forall\, i = 1, \ldots, n$ and $\forall\, \theta \in \Theta$, $k(\theta)$ is non-decreasing in θ_i.*

Proposition 11.2 (Characterization of DSIC SCFs) *An SCF $f \in \mathscr{C}$ is DSIC iff $\forall\, i = 1, \ldots, n$ and $\forall\, \theta \in \Theta$, we have*

1. *$k(\theta)$ is non-decreasing in θ_i*
2. *$t_i(\theta_i, \theta_{-i}) = t_i(\underline{\theta}_i, \theta_{-i}) - \int_{\underline{\theta}_i}^{\theta_i} \frac{\partial v_i(k(s,\theta_{-i}),s)}{\partial k} \frac{\partial k(s,\theta_{-i})}{\partial s} ds$*

Proposition 11.3 (Characterization of AE+DSIC SCFs) *An AE SCF $f \in \mathscr{C}$ is DSIC iff it satisfies the Groves payment scheme, that is, $\forall\, i = 1, \ldots, n$ and $\forall\, \theta \in \Theta$, we have*

$$t_i(\theta) = \left[\sum_{j \neq i} v_j(k^*(\theta), \theta_j) \right] + h_i(\theta_{-i}) \;\; \forall\, i = 1, \ldots, n \tag{19}$$

where $h_i(\cdot)$ is any arbitrary function of θ_{-i}.

12 Bayesian Implementation

Recall that we mentioned two possible routes to get around the Gibbard-Satterthwaite impossibility theorem. The first was to focus on restricted environments like quasi-linear environment, and the second one was to weaken the implementation concept and look for an SCF which is ex-post efficient, non-dictatorial, and Bayesian incentive compatible. In this section, our objective is to explore the second route.

Throughout this section, we will once again be working within the quasi-linear environment. As we saw earlier, the quasi-linear environments have a nice property that every social choice function in these environments is non-dictatorial. Therefore, while working within quasi-linear environment, we do not have to worry about non-dictatorial part of the social choice function. We can just investigate whether there exists any SCF in quasi-linear environment, which is both *ex-post efficient* and *BIC*, or equivalently which has three properties - *AE, BB,* and *BIC*. Recall that in the previous section, we have already addressed the question whether there exists any SCF in quasi-linear environment which is *AE, BB*, and *DSIC* and we found that hardly any function satisfies all these three properties. On the contrary, in this section, we will show that a wide range of SCF in quasi-linear environment satisfy three properties - AE, BB, and BIC.

12.1 (AE + BB + BIC) is Possible: Expected Externality Mechanisms

The following theorem, due to d'Aspremont and Gérard-Varet [45] and Arrow [46] confirms that in quasi-linear environment, there exist social choice functions which are both ex-post efficient (allocatively efficient + budget balance) and truthfully implementable in Bayesian Nash equilibrium (Bayesian incentive compatible). We refer this theorem by *dAGVA* theorem.

Theorem 12.1 (The dAGVA Theorem) *Let the social choice function $f(\cdot) = (k^*(\cdot), t_1(\cdot), \ldots, t_n(\cdot))$ be allocatively efficient and the agents' types be statistically independent of each other (i.e. the density $\phi(\cdot)$ has the form $\phi_1(\cdot) \times \ldots \times \phi_n(\cdot)$). This function can be truthfully implemented in Bayesian Nash equilibrium if satisfies the following payment structure, popularly known as dAGVA payment (incentive) scheme.*

$$t_i(\theta) = E_{\tilde{\theta}_{-i}}\left[\sum_{j \neq i} v_j(k^*(\theta_i, \tilde{\theta}_{-i}), \tilde{\theta}_j)\right] + h_i(\theta_{-i}) \;\; \forall\, i = 1, \ldots, n; \quad \forall theta \in Theta \tag{20}$$

where $h_i(\cdot)$ is any arbitrary function of θ_{-i}. Moreover, it is always possible to choose the functions $h_i(\cdot)$ such that $\sum_{i=1}^{n} t_i(\theta) = 0$.

Proof: Let the social choice function $f(\cdot) = (k^*(\cdot), t_1(\cdot), \ldots, t_n(\cdot))$ be allocatively efficient, i.e. it satisfies the condition (13), and also satisfies the dAGVA payment scheme (20). Consider

$$E_{\theta_{-i}}\left[u_i(f(\theta_i, \theta_{-i}), \theta_i)|\theta_i\right] \;\; = \;\; E_{\theta_{-i}}\left[v_i(k^*(\theta_i, \theta_{-i}), \theta_i) + t_i(\theta_i, \theta_{-i})|\theta_i\right]$$

Since θ_i and θ_{-i} are statistically independent, the expectation can be taken without conditioning on θ_i. This will give us

$$\begin{aligned}
E_{\theta_{-i}}\left[u_i(f(\theta_i,\theta_{-i}),\theta_i)|\theta_i\right] &= E_{\theta_{-i}}\left[v_i(k^*(\theta_i,\theta_{-i}),\theta_i) + h_i(\theta_{-i}) + E_{\tilde{\theta}_{-i}}\left[\sum_{j\neq i} v_j(k^*(\theta_i,\tilde{\theta}_{-i}),\tilde{\theta}_j)\right]\right] \\
&= E_{\theta_{-i}}\left[\sum_{j=1}^{n} v_j(k^*(\theta_i,\theta_{-i}),\theta_j)\right] + E_{\theta_{-i}}[h_i(\theta_{-i})]
\end{aligned}$$

Since $k^*(\cdot)$ is satisfies the condition (13),

$$\sum_{j=1}^{n} v_j(k^*(\theta_i,\theta_{-i}),\theta_j) \geq \sum_{j=1}^{n} v_j(k^*(\hat{\theta}_i,\theta_{-i}),\theta_j) \;\forall\, \hat{\theta}_i \in \Theta_i$$

Thus we get

$$E_{\theta_{-i}}\left[\sum_{j=1}^{n} v_j(k^*(\theta_i,\theta_{-i}),\theta_j)\right] + E_{\theta_{-i}}[h_i(\theta_{-i})] \geq E_{\theta_{-i}}\left[\sum_{j=1}^{n} v_j(k^*(\hat{\theta}_i,\theta_{-i}),\theta_j)\right] + E_{\theta_{-i}}[h_i(\theta_{-i})] \;\forall\, \hat{\theta}_i \in \Theta_i$$

Again by making use of statistical independence we can rewrite the above inequality in the following form

$$E_{\theta_{-i}}\left[u_i(f(\theta_i,\theta_{-i}),\theta_i)|\theta_i\right] \geq E_{\theta_{-i}}\left[u_i(f(\hat{\theta}_i,\theta_{-i}),\theta_i)|\theta_i\right] \;\forall\, \hat{\theta}_i \in \Theta_i$$

This shows that when agents $j \neq i$ announce their types truthfully, agent i finds the truth telling is his optimal strategy, thus proving that the SCF is BIC. We now show that the functions $h_i(\cdot)$ can be chosen to guarantee $\sum_{i=1}^{n} t_i(\theta) = 0$. Let us define,

$$\begin{aligned}
\xi_i(\theta_i) &= E_{\tilde{\theta}_{-i}}\left[\sum_{j\neq i} v_j(k^*(\theta_i,\tilde{\theta}_{-i}),\tilde{\theta}_j)\right] \;\forall\, i = 1,\ldots,n \\
h_i(\theta_{-i}) &= -\left(\frac{1}{n-1}\right)\sum_{j\neq i}\xi_j(\theta_j) \;\;\forall\, i = 1,\ldots,n
\end{aligned}$$

In view of the above definitions, we can say that

$$\begin{aligned}
t_i(\theta) &= \xi_i(\theta_i) - \left(\frac{1}{n-1}\right)\sum_{j\neq i}\xi_j(\theta_j) \\
\Rightarrow \sum_{i=1}^{n} t_i(\theta) &= \sum_{i=1}^{n}\xi_i(\theta_i) - \left(\frac{1}{n-1}\right)\sum_{i=1}^{n}\sum_{j\neq i}\xi_j(\theta_j) \\
\Rightarrow \sum_{i=1}^{n} t_i(\theta) &= \sum_{i=1}^{n}\xi_i(\theta_i) - \left(\frac{1}{n-1}\right)\sum_{i=1}^{n}(n-1)\xi_j(\theta_j) \\
\Rightarrow \sum_{i=1}^{n} t_i(\theta) &= 0
\end{aligned}$$

Q.E.D.

The budget balanced payment structure of the agents in the above mechanism can be given a nice graph theoretic interpretation. Imagine a directed graph $G = (V, A)$ where V is the set of $n+1$ vertices, numbered $0, 1, \ldots, n$, and A is the set of $[n + n(n-1)]$ directed arcs. The vertices starting from 1 through n correspond to the n agents involved into the system and the vertex number 0 corresponds to the social planner. The set A consists of two types of the directed arcs:

1. Arcs $0 \to i \ \ \forall i = 1, \ldots, n$
2. Arcs $i \to j \ \ \forall i, j \in \{1, 2, \ldots, n\}\,; i \neq j$

Each of the arcs $0 \to i$ carries a flow of $t_i(\theta)$ and each of the arcs $i \to j$ carries a flow of $\frac{\xi_i(\theta_i)}{n-1}$. Thus the total outflow from a node $i \in \{1, 2, \ldots, n\}$ is $\xi_i(\theta_i)$ and total inflow to the node i from nodes $j \in \{1, 2, \ldots, n\}$ is $-h_i(\theta_{-i}) = \left(\frac{1}{n-1}\right) \sum_{j \neq i} \xi_j(\theta_j)$ Thus for any node i, $t_i(\theta) + h_i(\theta_{-i})$ is the net outflow which it is receiving from node 0 in order to respect the flow conservation constraint. Thus, if $t_i(\cdot)$ is positive then the agent i receives the money from the social planner and if it is negative, then the agent pays the money to the social planner. However, by looking at flow conservation equation for node 0, we can say that total payment received by the planner from the agents and total payment made by the planner to the agents will add up to zero. In graph theoretic interpretation, the flow from node i to node j can be justified as follows. Each agent i first evaluates the expected total valuation that would be generated together by all his rival agents in his absence, which turns out to be $\xi_i(\theta_i)$. Now, agent i divides it equally among the rival agents and pays to every rival agent an amount equivalent to this. The idea can be better understood with the help of Figure 10 which depicts the three agents case.

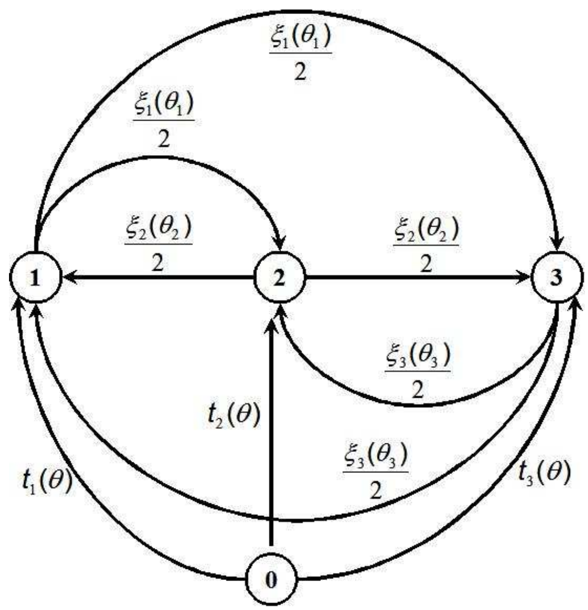

Figure 10: Payment structure in budget balance expected externality mechanism

After the results of d'Aspremont and Gérard-Varet [45] and Arrow [46], a direct revelation mechanism in which SCF is allocatively efficient and satisfies the dAGVA payment scheme is called as *dAGVA mechanism/expected externality mechanism/expected Groves mechanism.*

Definition 12.1 (dAGVA/expected externality/expected Groves Mechanisms) *A direct revelation mechanism, $\mathscr{D} = ((\Theta_i)_{i \in N}, f(\cdot))$ in which $f(\cdot) = (k(\cdot), t_1(\cdot), \ldots, t_n(\cdot))$ satisfies (13) and (20) is known as dAGVA/expected externality/expected Groves Mechanism.*[6]

In view of the definition of dAGVA mechanisms, the Figure 8 can be enriched by including the space of dAGVA mechanisms. This is shown in Figure 11.

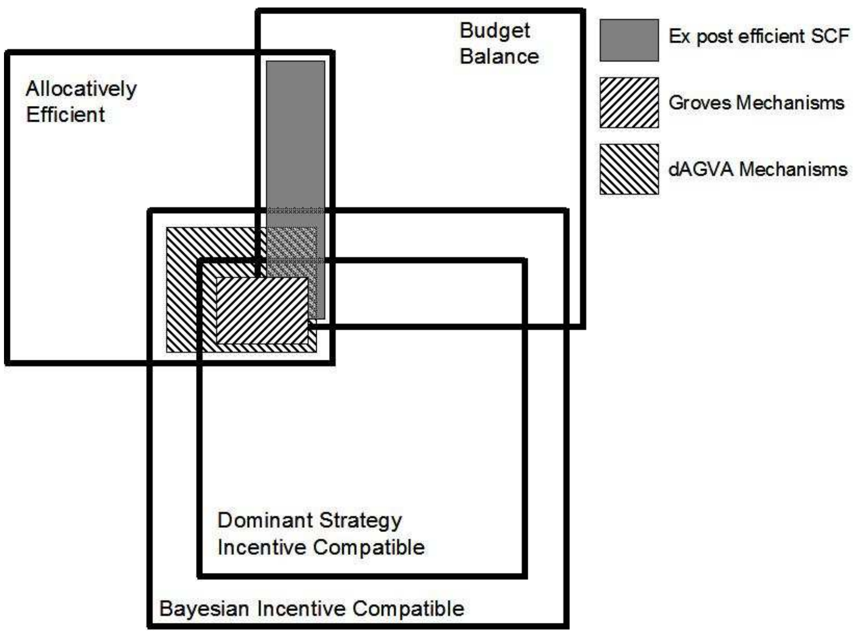

Figure 11: Space of social choice functions in quasi-linear environment

12.2 BIC in Linear Environment

The linear environment is a special, but often-studied, subclass of quasi-linear environment. This environment is a restricted version of the quasi-linear environment in following sense.

1. Each agent i's type lies in an interval $\Theta_i = [\underline{\theta}_i, \overline{\theta_i}] \subset \mathbb{R}$ with $\underline{\theta}_i < \overline{\theta_i}$
2. Agents' types are statistically independent, that is, the density $\phi(\cdot)$ has the form $\phi_1(\cdot) \times \ldots \times \phi_n(\cdot)$
3. $\phi_i(\theta_i) > 0 \; \forall \; \theta_i \in [\underline{\theta}_i, \overline{\theta_i}] \; \forall \; i = 1, \ldots, n$
4. Each agent i's utility function takes the following form

$$u_i(x, \theta_i) = \theta_i v_i(k) + m_i + t_i$$

[6] *We will sometime abuse the terminology and simply refer to a SCF $f(\cdot)$ satisfying (13) and (20) as dAGVA/expected externality/expected Groves Mechanisms.*

The linear environment has very interesting properties in terms of characterizing the BIC social choice functions. Before we present Myerson's characterization theorem for BIC social choice functions in linear environment, we would like to define following quantities with regard to any social choice function $f(\cdot) = (k(\cdot), t_1(\cdot), \ldots, t_n(\cdot))$ in this environment.

- Let $\overline{t_i}(\hat{\theta}_i) = E_{\theta_{-i}}[t_i(\hat{\theta}_i, \theta_{-i})]$ be agent i's expected transfer given that he announces his type to be $\hat{\theta}_i$ and that all agents $j \neq i$ truthfully reveal their types.
- Let $\overline{v_i}(\hat{\theta}_i) = E_{\theta_{-i}}[v_i(\hat{\theta}_i, \theta_{-i})]$ be agent i's expected "benefits" given that he announces his type to be $\hat{\theta}_i$ and that all agents $j \neq i$ truthfully reveal their types.
- Let $U_i(\hat{\theta}_i|\theta_i) = E_{\theta_{-i}}[u_i(f(\hat{\theta}_i, \theta_{-i}), \theta_i)|\theta_i]$ be agent i's expected utility when his type is θ_i, he announces his type to be $\hat{\theta}_i$, and that all agents $j \neq i$ truthfully reveal their types. It is easy to verify from previous two definitions that

$$U_i(\hat{\theta}_i|\theta_i) = \theta_i \overline{v_i}(\hat{\theta}_i) + \overline{t_i}(\hat{\theta}_i)$$

- Let $U_i(\theta_i) = U_i(\theta_i|\theta_i)$ be the agent i's expected utility conditional on his type being θ_i when he and all other agents report their true types. It is easy to verify that

$$U_i(\theta_i) = \theta_i \overline{v_i}(\theta_i) + \overline{t_i}(\theta_i)$$

In view of the paraphernalia developed, we now present Myerson's [37] theorem for characterizing the BIC social choice functions in this environment.

Theorem 12.2 (Myerson's Characterization Theorem) *In linear environment, a social choice function $f(\cdot) = (k(\cdot), t_1(\cdot), \ldots, t_n(\cdot))$ is BIC if and only if, for all $i = 1, \ldots, n$,*

1. $\overline{v_i}(\cdot)$ *is non-decreasing*
2. $U_i(\theta_i) = U_i(\underline{\theta_i}) + \int_{\underline{\theta_i}}^{\theta_i} \overline{v_i}(s)ds \;\; \forall \, \theta_i$

For proof of the above theorem, refer to Proposition 23.D.2 of [38]. The above theorem shows that to identify all BIC social choice functions in linear environment, we can proceed as follows: First identify which functions $k(\cdot)$ lead every agent i's expected benefit function $\overline{v_i}(\cdot)$ to be non-decreasing. Then, for each such function identify transfer functions $\overline{t_1}(\cdot), \ldots, \overline{t_n}(\cdot)$ that satisfy the second condition of the above proposition. Substituting for $U_i(\cdot)$ in the second condition above, we get that expected transfer functions are precisely those which satisfy, for $i = 1, \ldots, n$,

$$\overline{t_i}(\theta_i) = \overline{t_i}(\underline{\theta_i}) + \underline{\theta_i}\,\overline{v_i}(\underline{\theta_i}) - \theta_i \overline{v_i}(\theta_i) + \int_{\underline{\theta_i}}^{\theta_i} \overline{v_i}(s)ds$$

for some constant $\overline{t_i}(\underline{\theta_i})$. Finally, choose any set of transfer functions $t_1(\cdot), \ldots, t_n(\cdot)$ such that $E_{\theta_{-i}}[t_i(\theta_i, \theta_{-i})] = \overline{t_i}(\theta_i)$ for all θ_i. In general, there are many such functions, $t_i(\cdot, \cdot)$; one, for example, is simply $t_i(\theta_i, \theta_{-i}) = \overline{t_i}(\theta_i)$

In what follows we discuss a few examples where the environment is linear and analyze the BIC property of the social choice function by means of Myerson's characterization theorem.

12.3 Fair Bilateral Trade in Linear Environment

Once again consider the Example 4.1 of fair bilateral trade. Recall that each agent i's type lies in an interval $\Theta_i = [\underline{\theta_i}, \overline{\theta_i}]$. Let us impose the additional conditions on the environment to make it linear. We assume that

1. Agents' types are statistically independent, that is, the density $\phi(\cdot)$ has the form $\phi_1(\cdot) \times \phi_2(\cdot)$
2. Let each agent i draw his type from the set $[\underline{\theta_i}, \overline{\theta_i}]$ by means of a uniform distribution, that is $\phi_i(\theta_i) = 1/(\overline{\theta_i} - \underline{\theta_i}) \; \forall \, \theta_i \in [\underline{\theta_i}, \overline{\theta_i}] \; \forall \, i = 1, 2$

Note that the utility function of the agents in this example are given by

$$u_i(f(\theta), \theta_i) = \theta_i y_i(\theta) + t_i(\theta) \; \forall \, i = 1, 2$$

Thus, viewing $y_i(\theta) = v_i(k(\theta))$ will confirm that these utility functions also satisfy the fourth condition required for linear environment. Now we can apply Myerson's characterization theorem to test the Bayesian incentive compatibility of the SCF involved here. It is easy to see that $\overline{v_1}(\theta_1) = \overline{y_1}(\theta_1) = 1 - \Phi_2(\theta_1)$ is not a non-decreasing function. Therefore, we can claim that fair bilateral trade is not BIC.

12.4 First-Price Sealed Bid Auction in Linear Environment

Once again consider the Example 4.2 of first-price sealed bid auction. Let us assume that $S_i = \Theta_i = [\underline{\theta_i}, \overline{\theta_i}] \quad \forall \; i \in N$. In such a case, the first-price auction becomes a direct revelation mechanism $\mathscr{D} = ((\Theta_i)_{i \in N}, f(\cdot))$, where $f(\cdot)$ is an SCF which is same as outcome rule of the first-price auction. Let us impose the additional conditions on the environment to make it linear. We assume that

1. Bidders' types are statistically independent, that is, the density $\phi(\cdot)$ has the form $\phi_1(\cdot) \times \ldots \times \phi_n(\cdot)$
2. Let each bidder draw his type from the set $[\underline{\theta_i}, \overline{\theta_i}]$ by means of a uniform distribution, that is $\phi_i(\theta_i) = 1/(\overline{\theta_i} - \underline{\theta_i}) \; \forall \, \theta_i \in [\underline{\theta_i}, \overline{\theta_i}] \; \forall \, i = 1, \ldots, n$

Note that the utility function of the agents in this example are given by

$$u_i(f(\theta), \theta_i) = \theta_i y_i(\theta) + t_i(\theta) \; \forall \, i = 1, \ldots, n$$

Thus, viewing $y_i(\theta) = v_i(k(\theta))$ will confirm that these utility functions also satisfy the fourth condition required for linear environment. Now we can apply Myerson's characterization theorem to test the Bayesian incentive compatibility of the SCF involved here. It is easy to see that for any bidder i, we have

$$\begin{aligned}
\overline{v_i}(\theta_i) &= E_{\theta_{-i}}[v_i(\theta_i, \theta_{-i})] \\
&= E_{\theta_{-i}}[y_i(\theta_i, \theta_{-i})] \\
&= 1.P\left((\theta_{-i})_{(n-1)} \leq \theta_i\right) + 0.\left(1 - P\left(\theta_i < (\theta_{-i})_{(n-1)}\right)\right) \\
&= P\left((\theta_{-i})_{(n-1)} \leq \theta_i\right) \qquad (21)
\end{aligned}$$

where $P\left((\theta_{-i})_{(n-1)} \leq \theta_i\right)$ is the probability that the given type θ_i of the bidder i is the highest among all the bidders' types. This implies that in the presence of independence assumptions made above, $\overline{v_i}(\theta_i)$ is a non-decreasing function.

We know that for first-price sealed bid auction, $t_i(\theta) = -\theta_i y_i(\theta)$. Therefore, we can claim that for first-price sealed bid auction, we have

$$\overline{t_i}(\theta_i) \quad = \quad -\theta_i \overline{v_i}(\theta_i) \; \forall \; \theta_i \in \Theta_i$$

The above values of $\overline{v_i}(\theta_i)$ and $\overline{t_i}(\theta_i)$ can be used to compute $U_i(\theta_i)$ in following manner.

$$U_i(\theta_i) \quad = \quad \theta_i \overline{v_i}(\theta_i) + \overline{t_i}(\theta_i) = 0 \;\; \forall \; \theta_i \in [\underline{\theta_i}, \overline{\theta_i}] \tag{22}$$

The above equation can be used to test the second condition of the Myerson's theorem, which require

$$U_i(\theta_i) = U_i(\underline{\theta_i}) + \int_{\underline{\theta_i}}^{\theta_i} \overline{v_i}(s) ds$$

In view of the Equations (21) and (22), it is easy to see that this second condition of Myerson's characterization theorem is not being met by the SCF used in the first-price sealed bid auction. Therefore, we can finally claim that first-price sealed bid auction is not BIC in linear environment.

12.5 Second-Price Sealed Bid Auction in Linear Environment

Once again consider the example 4.3 of second-price sealed bid auction. Let us assume that $S_i = \Theta_i = [\underline{\theta_i}, \overline{\theta_i}] \;\; \forall \; i \in N$. In such a case, the second-price auction becomes a direct revelation mechanism $\mathscr{D} = ((\Theta_i)_{i \in N}, f(\cdot))$, where $f(\cdot)$ is an SCF which is same as outcome rule of the second-price auction. We have already seen that this SCF $f(\cdot)$ is DSIC in quasi-linear environment and linear environment is a special case of quasi-linear environment, therefore, it is DSIC in the linear environment also. Moreover, we know that DSIC implies BIC. Therefore, we can directly claim that SCF used in the Vickrey auction is BIC in linear environment.

Table 2 summarizes the properties of the SCFs discussed in above examples in linear environments.

SCF	**AE**	**BB**	**DSIC**	**BIC**
Fair Bilateral Trade	✓	✓	×	×
First-Price Auction	✓	×	×	×
Vickrey Auction	✓	×	✓	✓

Table 2: Properties of social choice functions in linear environment

13 The Revenue Equivalence Theorem

There are four basic types of auctions when a single indivisible item is to be sold:

1. **English auction**: This is also called oral auction, open auction, open cry auction, and ascending bid auction. Here, the price starts at a low level and is successively raised until only one bidder remains in the fray. This can be done in several ways: (a) an auctioneer announces prices, (b) bidders call the bids themselves, or (b) bids are submitted electronically. At any point of time, each bidder knows the level of the current best bid. The winning bidder pays the last going price.

2. **Dutch auction**: This is also called a descending bid auction. Here, the auctioneer announces an initial (high) price and then keeps lowering the price iteratively until one of the bidders accepts the current price. The winner pays the current price.

3. **First price sealed bid auction**: Recall that in this auction, potential buyers submit sealed bids and the highest bidder is awarded the item. The winning bidder pays the price that he has bid.

4. **Second price sealed bid auction**: This is the classic Vickrey auction. Recall that potential buyers submit sealed bids and the highest bidder is awarded the item. The winning bidder pays a price equal to the second highest bid (which is also the highest losing bid).

When a single indivisible item is to be bought or procured, the above four types of auctions can be used in a *reverse* way. These are called *reverse auctions* or *procurement auctions*. In this section, we would be discussing the revenue equivalence theorem as it applies to selling. The procurement version can be analyzed on similar lines.

13.1 The Benchmark Model

There are four assumptions underlying the derivation of the revenue equivalence theorem: (1) risk neutrality of bidders (2) bidders have independent private values (3) bidders are symmetric (4) payments depend on bids alone. These are described below in more detail.

13.1.1 Risk Neutrality of Bidders

A bidder is said to be:

- *risk-averse* if his utility is a *concave* function of his wealth; that is an increment in the wealth at a lower level of wealth leads to an increment in utility that is higher than the increase in utility due to an identical increment in wealth at a higher level of wealth;

- *risk-loving* if his utility is a *convex* function of his wealth; that is an increment in the wealth at a lower level of wealth leads to an increment in utility that is lower than the increase in utility due to an identical increment in wealth at a higher level of wealth; and

- *risk-neutral* if his utility is a *linear* function of his wealth; that is an increment in the wealth at a lower level of wealth leads to the same increment in the utility as an identical increment would yield at a higher level of wealth.

It is assumed in the benchmark model that all the bidders are risk neutral. This immediately implies that the utility functions are linear.

13.1.2 Independent Private Values Model

In the independent private values model, each bidder knows precisely how highly he values the item. He has no doubt about the true value of the item to him. However, each bidder does not know anyone else's valuation of the item. Instead, he perceives any other bidder's valuation as a draw from some known probability distribution. Also, each bidder knows that the other bidders and the seller regard his own valuation as being drawn from some probability distribution. More formally, let $N = \{1, 2, \cdots, n\}$ be the set of bidders. There is a probability distribution Φ_i from which bidder i draws his valuation v_i. Only bidder i observes his own valuation v_i, but all other bidders and the seller know the distribution Φ_i. Any one bidder's valuation is statistically independent from any other bidder's valuation.

An apt example of this assumption is provided by the auction of an antique in which the bidders are consumers buying for their own use and not for resale. Another example is Government contract bidding when each bidder known his own production cost if he wins the contract.

Common Value Model

A contrasting model is the *common value model*. Here, if V is the unobserved true value of the item, then the bidders' perceived values $v_i, i = 1, 2, \cdots, n$ are independent draws from some probability distribution $H(v_i|V)$. All the bidders know the distribution H. An example is provided by the sale of an antique that is being bid for by dealers who intend to resell it. The item has one single objective value, namely its market price. However, no one knows the true value. The bidders, perhaps having access to different information, have different guesses about how much the item is objectively worth. Another example is that of sale of mineral rights to a particular tract of land. The objective value here is the amount of mineral actually lying beneath the ground. However no one knows its true value.

Suppose a bidder were somehow to learn another bidder's valuation. If the situation is described by the common value model, then the above provides useful information about the likely true value of the item and the bidder would probably change his own valuation in the light of this. If the situation is described by the independent private value model, the bidder knows his own mind and learning about another's valuation will not cause him to change his own valuation (although he may, for strategic reasons, change his bid).

Real world auction situations are likely to contain aspects of both the independent private values model and the common value model. It is assumed in the benchmark model that the independent private values assumption holds.

13.1.3 Symmetry

This assumption implies that the all bidders have the same set of possible valuations and further they draw their valuations using the same probability distribution Φ. That is $\phi_1 = \Phi_2 = \ldots = \Phi_n$.

13.1.4 Dependence of Payments on Bids Alone

It is assumed that the payment to be made by the winner to the auctioneer is a function of bids alone.

13.2 Statement of the Revenue Equivalence Theorem

It is clear that the benchmark model leads to a Bayesian game. We have already studied the structure of this game while discussing the first price and second price auctions earlier. Bayesian Nash equilibrium is a natural outcome of such a game: each bidder bids an amount that is some function of his own valuation, such that, given that everyone else chooses his bid in this way, no individual bidder could do better by bidding differently. This Bayesian Nash equilibrium turns out to be a weakly dominant strategy equilibrium for the English auction and second price auction.

Theorem 13.1 (Revenue Equivalence Theorem) *: Consider a seller or an auctioneer trying to sell a single indivisible item in which n bidders are interested. For the benchmark model (bidders are risk neutral, bidders have independent private values, bidders are symmetric, and payments depend only on bids), all the four basic auction types (English auction, Dutch auction, first price auction, and second price auction) yield the same average revenue to the seller.*

The result looks counter intuitive: for example, it might seen that receiving the highest bid in a first price scaled bid auction must be better for the seller than receiving the second highest bid, as in second price auction. However, it is to be noted that bidders act differently in different auction situations. In particular, they bid more aggressively in a second price auction than in a first price auction.

13.3 Proof of the Revenue Equivalence Theorem

The proof proceeds in three parts. In Part 1, we show that the first price auction and the second price auction yield the same expected revenue in their respective equilibria. In Part 2, we show that the Dutch auction and the first price auction produce the same outcome. In Part 3, we show that the English auction and the second price auction yield the same outcome. Before taking up these parts, we first state and prove an important proposition.

A Proposition on Revenue Equivalence of Two Auctions

Assume that $y_i(\theta)$ is the probability of agent i getting the object when the vector of announced types is $\theta = (\theta_1, \ldots, \theta_n)$. Expected payoff to the buyer i with a type profile $\theta = (\theta_1, \cdots, \theta_n)$ will be $y_i(\theta)\theta_i + t_i(\theta)$. The set of allocations is given by

$$K = \left\{ (y_1, \cdots, y_n) : y_i \in [0,1] \forall i = 1, \cdots, n; \sum_{i=1}^{n} y_i \leq 1 \right\}$$

As earlier, let $\overline{y_i}(\hat{\theta}_i) = E_{\theta_i}[y_i(\hat{\theta}_i, \theta_i)]$ be the probability that agent i gets the object conditional to announcing his type as $\hat{\theta}_i$, with the rest of the agents announcing their types truthfully. Similarly, $\overline{t_i}(\hat{\theta}_i) = E_{\theta_i}[t_i(\hat{\theta}_i, \theta_i)]$ denotes the expected payment received by agent i conditional to announcing his type as $\hat{\theta}_i$, with the rest of the agents announcing their types truthfully. Let $\overline{v_i}(\hat{\theta}_i) = \overline{y_i}(\hat{\theta}_i)$. Then,

$$U_i(\theta_i) \quad = \quad \overline{y_i}(\theta_i)\theta_i + \overline{t_i}(\theta_i)$$

denotes the payoff to agent i when all the buying agents announce their types truthfully. We now state and prove an important proposition.

Proposition 13.1 *Consider an auction scenario with:*

1. *n risk-neutral bidders (buyers) $1, 2, \cdots, n$*
2. *The valuation of bidder i $(i = 1, \cdots, n)$ is a real interval $[\underline{\theta_i}, \overline{\theta_i}] \subset R$ with $\underline{\theta_i} < \overline{\theta_i}$*

3. *The valuation of bidder* i $(i = 1, \cdots, n)$ *is drawn from* $[\underline{\theta_i}, \overline{\theta_i}]$ *with a strictly positive density* $\phi_i(.) > 0$*. Let* $\Phi_i(.)$ *be the cumulative distribution function.*

4. *The bidders' types are statistically independent.*

Suppose that a given pair of Bayesian Nash equilibria of two different auction procedures are such that:

- *For every bidder* i*, for each possible realization of* $(\theta_1, \cdots, \theta_n)$*, bidder* i *has an identical probability of getting the good in the two auctions.*
- *Every bidder* i *has the same expected payoff in the two auctions when his valuation for the object is at its lowest possible level.*

Then the two auctions generate the same expected revenue for the seller.

Proof: By the revelation principle, it is enough we investigate two Bayesian incentive compatible social choice functions in this auction setting. It is enough we show that two Bayesian incentive compatible social choice functions having (a) the same allocation functions $(y_1(\theta), \cdots, y_n(\theta))$ $\forall \theta \in \Theta$, and (b) the same values of $U_1(\underline{\theta_1}), \cdots, U_n(\underline{\theta_n})$ will generate the same expected revenue to the seller.

We first derive an expression for the seller's expected revenue given any Bayesian incentive compatible mechanism. Expected revenue to the seller

$$= \sum_{i=1}^{n} E_\theta[-t_i(\theta)] \tag{23}$$

Now, we have:

$$\begin{aligned} E_\theta[-t_i(\theta)] &= E_{\theta_i}[-E_{\theta_{-i}}[t_i(\theta)]] \\ &= \int_{\underline{\theta_i}}^{\overline{\theta_i}} [\overline{y_i}(\theta_i)\theta_i - U_i(\theta_i)]\phi_i(\theta_i)d\theta_i \\ &= \int_{\underline{\theta_i}}^{\overline{\theta_i}} \left[\overline{y_i}(\theta_i)\theta_i - U_i(\underline{\theta_i})] - \int_{\underline{\theta_i}}^{\theta_i} \overline{y_i}(s)ds\right] \phi_i(\theta_i)d\theta_i \end{aligned}$$

The last step is an implication of Myerson's characterization of Bayesian incentive compatible functions in linear environment. The above expression is now equal to

$$= \left[\int_{\underline{\theta_i}}^{\overline{\theta_i}} \left(\overline{y_i}(\theta_i)\theta_i - \int_{\underline{\theta_i}}^{\theta_i} \overline{y_i}(s)ds\right) \phi_i(\theta_i)d\theta_i\right] - U_i(\underline{\theta_i})$$

Now, applying integration by parts with $\int_{\underline{\theta_i}}^{\theta_i} \overline{y_i}(s)ds$ as the first function, we get

$$\int_{\underline{\theta_i}}^{\overline{\theta_i}} \left(\int_{\underline{\theta_i}}^{\theta_i} \overline{y_i}(s)ds\right) \phi_i(\theta_i)d\theta_i$$

$$= \int_{\underline{\theta_i}}^{\overline{\theta_i}} \overline{y_i}(\theta_i)d\theta_i - \int_{\underline{\theta_i}}^{\overline{\theta_i}} \overline{y_i}(\theta_i)\Phi_i(\theta_i)d\theta_i$$

$$= \int_{\underline{\theta_i}}^{\overline{\theta_i}} \overline{y_i}(\theta_i)[1 - \Phi_i(\theta_i)]d\theta_i$$

Therefore we get

$$\begin{aligned} E[-\overline{t_i}(\theta_i)] &= -U_i(\underline{\theta_i}) + \left[\int_{\underline{\theta_i}}^{\overline{\theta_i}} \overline{y_i}(\theta_i) \left\{\theta_i - \frac{1 - \Phi_i(\theta_i)}{\phi_i(\theta_i)}\right\} \phi_i(\theta_i) d\theta_i\right] \\ &= -U_i(\underline{\theta_i}) + \left[\int_{\underline{\theta_1}}^{\overline{\theta_1}} \cdots \int_{\underline{\theta_n}}^{\overline{\theta_i}} y_i(\theta_1, \cdots, \theta_n) \right. \\ &\quad \left. \times \left(\theta_i - \frac{1 - \Phi_i(\theta_i)}{\phi_i(\theta_i)}\right) \left(\prod_{j=1}^{n} \phi_j(\theta_j)\right) d\theta_n \cdots d\theta_1 \right] \end{aligned}$$

since

$$\overline{y_i}(\theta_i) = \int_{\underline{\theta_1}}^{\overline{\theta_1}} \cdots \int_{\underline{\theta_n}}^{\overline{\theta_n}} y_i(\theta_1, \cdots, \theta_n) \quad \underbrace{d\theta_n \cdots d\theta_1}_{\text{without } d\theta_i}$$

Therefore the expected revenue of the seller

$$\begin{aligned} = & \left[\int_{\underline{\theta_1}}^{\overline{\theta_1}} \cdots \int_{\underline{\theta_n}}^{\overline{\theta_n}} \sum_{i=1}^{n} y_i(\theta_1, \cdots, \theta_n) \left(\theta_i - \frac{1 - \Phi_i(\theta_i)}{\phi_i(\theta_i)}\right)\right] \left(\prod_{j=1}^{n} \phi_j(\theta_j)\right) d\theta_n \cdots d\theta_1 \\ & - \sum_{i=1}^{n} U_i(\underline{\theta_i}) \end{aligned}$$

By looking at the above expression, we see that any two Bayesian incentive compatible social choice functions that generate the same functions $(y_1(\theta), \cdots, y_n(\theta))$ and the same values of $(U_1(\underline{\theta_1}), \cdots, U_n(\underline{\theta_n}))$ generate the same expected revenue for the seller.

This proves the proposition. With this in hand, we now prove the revenue equivalence theorem in three stages.

Part 1: Revenue Equivalence of First Price Auction and Second Price Auction

The first price auction and the second price auction satisfy the conditions of the above proposition.

- In both the auctions, the bidder with the highest valuation wins the auction
- bidders' valuations are drawn from $[\underline{\theta_i}, \overline{\theta_i}]$ and a bidder with valuation at the lower limit of the interval has a payoff of zero in both the auctions.

Thus the proposition can be applied to the equilibria of the two auctions: Note that in the case of first price auction, it is a Bayesian Nash equilibrium while in the case of second price auction, it is a weakly dominant strategy equilibrium. In fact, it can be shown in any *symmetric* auction setting (where the bidders' valuations are independently drawn from identical distributions) that the conditions of the above proposition will be satisfied by any Bayesian Nash equilibrium of first price auction and the weakly dominant strategy equilibrium of the second price scaled bid auction.

Part 2: Equivalence of Dutch Auction and First Price Auction

To prove this, consider the situation facing a bidder in these two auctions. In each case, the bidder must choose how high to bid without knowing the other bidders' decisions. If he wins, the price he pays equals his own bid. This result is true irrespective of which of the assumptions in the benchmark model apply. Note that the equilibrium in the underlying Bayesian game in the two cases here is a Bayesian Nash equilibrium.

Part 3: Equivalence of English Auction and Second Price Auction

First we analyze the English auction. Note that a bidder drops out as soon the going price exceeds his valuation. The second last bidder drops out as soon as the price exceeds his own valuation. This leaves only one bidder in the fray and he wins the auction. Note that the winning bidder's valuation is the highest among all the bidders and he earns some payoff in spite of the monopoly power of the seller. Only the winning bidder knows how much payoff he receives because only he knows his own valuation. Suppose the valuations of the n bidders are $v_{(1)}, v_{(2)}, \cdots, v_{(n)}$. Since the bidders are symmetric, these valuations are draws from the same distribution and without loss of generality, assume that these are in descending order. The winning bidder gets a payoff of $v_{(1)} - v_{(2)}$.

Next we analyze the second price auction. In the second price auction, the bidder's choice of bid determines only whether or not he wins; the amount he pays if he wins is beyond his control. Suppose the bidder considers lowering his bid below his valuation. The only case in which this changes the outcome occurs when this lowering of his bid results in his bid now being lower than someone else's. Because he would have earned non-negative payoff if he won, lowering his bid cannot make him better off. Suppose the bidder considers raising his bid above his valuation. The only case in which this changes the outcome occurs when some other bidder has submitted a bid higher than the first bidder's valuation but lower than his new bid. Thus raising the bid causes this bidder to win but he must pay more for the item than it is worth to him; raising his bid beyond his valuation cannot make him better off. Thus, each bidder's equilibrium best response strategy is to bid his own valuation for the item. The payment here is equal to the actual valuation of the bidder with the second highest valuation (i.e., realization of the second order statistic). Thus the expected payment and payoff are the same in English auction and the second price auction. This establishes Part 3 and therefore proves the revenue equivalence theorem.

Note that the outcomes of the English auction and the second price auction satisfy a weakly dominant strategy equilibrium. That is, each bidder has a well defined best bid regardless of how high he believes his rivals will bid. In the second price auction, the weakly dominant strategy is to bid time valuation. In the English auction, the weakly dominant strategy is to remain in the bidding until the price reaches the bidder's own valuation.

13.4 Some Observations on the Revenue Equivalence Theorem

We now make a few important observations.

- The theorem does not imply that the outcomes of the four auction forms are always exactly the same. They are only equal on average.
  - Note that in the English auction or the second price auction, the price exactly equals the valuation of the bidder with the second highest valuation, $v_{(2)}$. In Dutch auction or the first price auction, the price is the expectation of the second highest valuation conditional on

the winning bidder's own valuation. The above two prices will be equal only by accident, however they are equal on average.

- Bidding logic is very simple in the English auction and second price auction. In the former, a bidder remains in bidding until the price reaches his valuation. In the latter, he submits a sealed bid equal to his own valuation.
- On the other hand, the bidding logic is quite complex in the Dutch auction and the first price auction. Here the bidder bids some amount less than his true valuation. Exactly how much less depends upon the probability distribution of the other bidders' valuations and the number of competing bidders. Finding the Nash equilibrium bid is a non-trivial computational problem.
- The revenue equivalence theorem is devoid of empirical predictions about which type of auction will be chosen by the seller in any particular set of circumstances. However when the assumptions of the benchmark model are relaxed, particular auction forms emerge as being superior.
- The variance of revenue is lower in English auction or second price auction than in Dutch auction or first price auction. Hence if the seller were risk averse, he would choose English or second price rather than Dutch or first price.
- For more details on the revenue equivalence theorem, the reader is referred to the classic survey paper by McAfee and McMillan [47], and the books by Milgrom [48] and Krishna [49].

14 Concept of Individual Rationality

Note that if a social choice function is BIC then each agent i finds it in his best interest to tell truth if all the other agents are also doing so. However, note that neither agent i nor the social planner has any control over the types that are being revealed by the agents other than i. Therefore, agent i may wonder what if any one or more of his rival agents announce untruthful types. In such a situation, telling truth should not result in any kind of loss to the agent i. Otherwise, either agent i will also start lying or, alternatively, he may quit the mechanism itself because participation in the mechanism is voluntary. Thus, in order to avoid such a situation, the social planner needs to ensure that each agent, despite what the rival agents are reporting, will be better off telling truth than not participating in the mechanism. These constraints are known as *participation constraints* or *individual rationality constraints*. Thus, individual rationality adds one more dimension to the desirable properties of a social choice function.

There are three stages at which participation constraints may be relevant in any particular application.

14.1 Ex-Post Individual Rationality Constraints

These constraints become relevant when any agent i is given a choice to withdraw from the mechanism at the ex-post stage that arise after all the agents have announced their types and an outcome in X has been chosen. Let $\overline{u_i}(\theta_i)$ be the utility that agent i receives by withdrawing from the mechanism when his type is θ_i. Then, to ensure agent i's participation, we must satisfy the following *ex-post participation (or individual rationality) constraints*

$$u_i(f(\theta_i, \theta_{-i}), \theta_i) \geq \overline{u_i}(\theta_i) \;\; \forall \, (\theta_i, \theta_{-i}) \in \Theta$$

14.2 Interim Individual Rationality Constraints

Let the agent i be allowed to withdraw from the mechanism only at an interim stage that arises after the agents have learned their type but before they have chosen their actions in the mechanism. In such a situation, the agent i will participate in the mechanism only if his interim expected utility $U_i(\theta_i|f) = E_{\theta_{-i}}[u_i(f(\theta_i,\theta_{-i}),\theta_i)|\theta_i]$ from social choice function $f(\cdot)$, when his type is θ_i, is greater than $\overline{u_i}(\theta_i)$. Thus, *interim participation (or individual rationality) constraints* for agent i require that

$$U_i(\theta_i|f) = E_{\theta_{-i}}[u_i(f(\theta_i,\theta_{-i}),\theta_i)|\theta_i] \geq \overline{u_i}(\theta_i) \;\; \forall\, \theta_i \in \Theta_i$$

14.3 Ex-Ante Individual Rationality Constraints

Let agent i be allowed to refuse to participate in the mechanism only at ex-ante stage that arises before the agents learn their type. In such a situation, the agent i will participate in the mechanism only if his ex-ante expected utility $U_i(f) = E_\theta[u_i(f(\theta_i,\theta_{-i}),\theta_i)]$ from social choice function $f(\cdot)$ is greater than $E_{\theta_i}[\overline{u_i}(\theta_i)]$. Thus, *ex-ante participation (or individual rationality) constraints* for agent i require that

$$U_i(f) = E_\theta[u_i(f(\theta_i,\theta_{-i}),\theta_i)] \geq E_{\theta_i}[\overline{u_i}(\theta_i)]$$

The following proposition establishes a relationship among the three different participation constraints discussed above.

Proposition 14.1 *For any social choice function $f(\cdot)$, we have*

$$f(\cdot) \textit{ is ex-post IR} \Rightarrow f(\cdot) \textit{ is interim IR} \Rightarrow f(\cdot) \textit{ is ex-ante IR}$$

The next proposition establishes the individual rationality of Clarke mechanism. First, we provide two definitions.

Definition 14.1 (Choice Set Monotonicity) *We say that a mechanism $\mathscr{M}$ is choice set monotone if the the set of feasible outcomes X (weakly) increases as additional agents are introduced into the system. An implication of this property is $K_{-i} \subset K \;\forall\, i = 1,\ldots,n$.*

Definition 14.2 (No Negative Externality) *Consider a choice set monotone mechanism $\mathscr{M}$. We say that the mechanism $\mathscr{M}$ has no negative externality if for each agent i, each $\theta \in \Theta$, and each $k^*_{-i}(\theta_{-i}) \in B^*(\theta_{-i})$, we have*

$$v_i(k^*_{-i}(\theta_{-i}),\theta_i) \geq 0$$

It is easy to verify that the mechanisms in all the previous examples - *fair bilateral trade, first-price sealed bid auction, Vickrey auction*, and *GVA* satisfy all the three properties described above.

Proposition 14.2 (Ex-Post Individual Rationality of Clarke Mechanism) *Let us consider a Clarke mechanism in which*

1. $\overline{u_i}(\theta_i) = 0 \;\forall \theta_i \in \Theta_i;\; \forall\, i = 1,\ldots,n$
2. *The mechanism satisfies choice set monotonicity property*
3. *The mechanism satisfies no negative externality property*

Then the Clarke mechanism is ex-post individual rational.

Proof: Recall that utility $u_i(f(\theta), \theta_i)$ of an agent i in Clarke mechanism is given by

$$\begin{aligned} u_i(f(\theta), \theta_i) &= v_i(k^*(\theta), \theta_i) + \left[\sum_{j \neq i} v_j(k^*(\theta), \theta_j)\right] - \left[\sum_{j \neq i} v_j(k^*_{-i}(\theta_{-i}), \theta_j)\right] \\ &= \left[\sum_{j} v_j(k^*(\theta), \theta_j)\right] - \left[\sum_{j \neq i} v_j(k^*_{-i}(\theta_{-i}), \theta_j)\right] \end{aligned}$$

By virtue of choice set monotonicity, we know that $k^*_{-i}(\theta_{-i}) \in K$. Therefore, we have

$$\begin{aligned} u_i(f(\theta), \theta_i) &\geq \left[\sum_{j} v_j(k^*_{-i}(\theta_{-i}), \theta_j)\right] - \left[\sum_{j \neq i} v_j(k^*_{-i}(\theta_{-i}), \theta_j)\right] \\ &= v_i(k^*_{-i}(\theta_{-i}), \theta_i) \\ &\geq 0 = \overline{u_i}(\theta_i) \end{aligned}$$

The last step follows due to the fact that mechanism has no negative externality.

Q.E.D.

Let us now investigate the individual rationality of the social choice functions of the examples discussed earlier.

14.4 Individual Rationality in Fair Bilateral Trade

Let us consider the example of fair bilateral trade. Let us assume that utility $\overline{u_i}(\theta_i)$ derived by the agents i from not participating into the trade, when his type is θ_i, is as follows.

$$\begin{aligned} \overline{u_1}(\theta_1) &= \theta_1 \ \forall \, \theta_1 \in [\underline{\theta_i}, \overline{\theta_i}] \\ \overline{u_2}(\theta_2) &= 0 \ \forall \, \theta_2 \in [\underline{\theta_i}, \overline{\theta_i}] \end{aligned}$$

In view of the above definitions, it is now easy to see that the SCF used in this example is ex-post IR.

14.5 Individual Rationality in First-Price Sealed Bid Auction

Let us consider the example of first-price sealed bid auction. If for each possible type θ_i, the utility $\overline{u_i}(\theta_i)$ derived by the agents i from not participating into the auction is 0, then it is easy to see that the SCF used in this example would be ex-post IR.

14.6 Individual Rationality in Vickrey Auction

Let us consider the example of second-price sealed bid auction. If for each possible type θ_i, the utility $\overline{u_i}(\theta_i)$ derived by the agents i from not participating into the auction is 0, then it is easy to see that the SCF used in this example would be ex-post IR. Moreover, the ex post IR of this example also follows directly from Proposition 14.2 because this is a special case of Clarke mechanism satisfying all the required conditions in the proposition.

14.7 Individual Rationality in GVA

Let us consider the example of generalized Vickrey auction. If for each possible type θ_i, the utility $\overline{u_i}(\theta_i)$ derived by the agents i from not participating into the auction is 0, then by virtue of Proposition 14.2, we can claim that SCF used here would be ex-post IR.

By including the above facts regarding individual rationality in Tables 1 and 2, we get Tables 3 and 4.

SCF	**AE**	**BB**	**DSIC**	**Ex-Post IR**
Fair Bilateral Trade	✓	✓	×	✓
First-Price Auction	✓	×	×	✓
Vickrey Auction	✓	×	✓	✓
GVA	✓	×	✓	✓

Table 3: Properties of social choice functions in quasi-linear environment

SCF	**AE**	**BB**	**DSIC**	**BIC**	**Ex-Post IR**
Fair Bilateral Trade	✓	✓	×	×	✓
First-Price Auction	✓	×	×	×	✓
Vickrey Auction	✓	×	✓	✓	✓

Table 4: Properties of social choice functions in linear environment

14.8 The Myerson-Satterthwaite Theorem

In the previous section, we saw that we did not even have a single example where we have all the desired properties in a SCF - AE, BB, BIC, and IR. This provides a motivation to study the feasibility of having all these properties in a social choice function.

The Myerson-Satterthwaite theorem is one disappointing news in this direction, which tells us that in a bilateral trade setting, whenever the gains from the trade are possible but not certain, then there is no SCF which satisfies AE, BB, BIC, and Interim IR all together. The precise statement of the theorem is as follows.

Theorem 14.1 (Myerson-Satterthwaite Impossibility Theorem) *Consider a bilateral trade setting in which the buyer and seller are risk neutral, the valuations θ_1 and θ_2 are drawn independently from the intervals $[\underline{\theta_1}, \overline{\theta_1}] \subset \mathbb{R}$ and $[\underline{\theta_2}, \overline{\theta_2}] \subset \mathbb{R}$ with strict positive densities, and $(\underline{\theta_1}, \overline{\theta_1}) \bigcap (\underline{\theta_2}, \overline{\theta_2}) \neq \emptyset$. Then there is no Bayesian incentive compatible social choice function that is ex-post efficient and gives every buyer type and every seller type non-negative expected gains from participation.*

For proof of the above theorem, refer to Proposition 23.E.1 of [38].

15 Optimal Mechanism Design

So far, we have studied the various properties of a mechanism. We have also studied a set of possibility and impossibility results. An obvious problem that faces a social planner is to decide which direct revelation mechanism (or equivalently, social choice function) is *optimal* for a given problem. In the rest of this paper, our objective is to familiarize the reader with a couple of techniques which social planner can adopt to design an optimal direct revelation mechanism for a given problem at hand.

One notion of optimality in multi-agent systems is that of *Pareto efficiency*. We now define three different notions of efficiency: ex-ante, interim, and ex-post.

Definition 15.1 (Ex-Ante Efficiency) *For any given set of social choice functions F, and any member $f(\cdot) \in F$, we say that $f(\cdot)$ is ex-ante efficient in F if there is no other $\hat{f}(\cdot) \in F$ having the following two properties*

$$
\begin{aligned}
E_\theta[u_i(\hat{f}(\theta), \theta_i)] &\geq E_\theta[u_i(f(\theta), \theta_i)] \;\; \forall\, i = 1, \ldots, n \\
E_\theta[u_i(\hat{f}(\theta), \theta_i)] &> E_\theta[u_i(f(\theta), \theta_i)] \;\; \text{for some } i
\end{aligned}
$$

Definition 15.2 (Interim Efficiency) *For any given set of social choice functions F, and any member $f(\cdot) \in F$, we say that $f(\cdot)$ is interim efficient in F if there is no other $\hat{f}(\cdot) \in F$ having the following two properties*

$$
\begin{aligned}
E_{\theta_{-i}}[u_i(\hat{f}(\theta), \theta_i)|\theta_i] &\geq E_{\theta_{-i}}[u_i(f(\theta), \theta_i)|\theta_i] \;\; \forall\, i = 1, \ldots, n, \; \forall\, \theta_i \in \Theta_i \\
E_{\theta_{-i}}[u_i(\hat{f}(\theta), \theta_i)|\theta_i] &> E_{\theta_{-i}}[u_i(f(\theta), \theta_i)|\theta_i] \;\; \text{for some } i \text{ and some } \theta_i \in \Theta_i
\end{aligned}
$$

Definition 15.3 (Ex-Post Efficiency) *For any given set of social choice functions F, and any member $f(\cdot) \in F$, we say that $f(\cdot)$ is ex-post efficient in F if there is no other $\hat{f}(\cdot) \in F$ having the following two properties*

$$
\begin{aligned}
u_i(\hat{f}(\theta), \theta_i) &\geq u_i(f(\theta), \theta_i) \;\; \forall\, i = 1, \ldots, n, \; \forall\, \theta \in \Theta \\
u_i(\hat{f}(\theta), \theta_i) &> u_i(f(\theta), \theta_i) \;\; \text{for some } i \text{ and some } \theta \in \Theta
\end{aligned}
$$

Using the above definition of ex-post efficiency, we can say that a social choice function $f(\cdot)$ is ex-post efficient in the sense of definition 6.1 if and only if it is ex-post efficient in the sense of definition 15.3 when we take $F = \{f : \Theta \to X\}$.

The following proposition establishes a relationship among these three different notions of efficiency.

Proposition 15.1 *Given any set of feasible social choice functions F and $f(\cdot) \in F$, we have*

$$
f(\cdot) \text{ is ex-ante efficient} \Rightarrow f(\cdot) \text{ is interim efficient} \Rightarrow f(\cdot) \text{ is ex-post efficient}
$$

For proof of the above proposition, refer to Proposition 23.F.1 of [38]. Also, compare the above proposition with the Proposition 14.1.

With this setup, we now try to formalize the design objectives of a social planner. For this, we need to define the concept known as *social welfare function.*

Definition 15.4 (Social Welfare Function) *A social welfare function (SWF) is a function* $w : \mathbb{R}^n \rightarrow \mathbb{R}$ *that aggregates the profile* $(u_1, \ldots, u_n) \in \mathbb{R}^n$ *of individual utility values of the agents into a social utility.*

15.1 Social Welfare Function and Mechanism Design

Consider a mechanism design problem and a direct revelation mechanism $\mathscr{D} = ((\Theta_i)_{i \in N}, f(\cdot))$ proposed for it. Let $(\theta_1, \ldots, \theta_n)$ be the actual type profile of the agents and assume for a moment that they will all reveal their true types when requested by the planner. In such a case, the social utility that would be realized by the social planner for every possible type profile θ of the agents is given by:

$$w(u_1(f(\theta), \theta_1), \ldots, u_n(f(\theta), \theta_n)) \tag{24}$$

However, recall the implicit assumption behind a mechanism design problem, namely, that the agents are autonomous and they would report a type as dictated by their rational behavior. Therefore, the assumption that all the agents will report their true types is not true in general. In general, rationality implies that the agents report their types according to a strategy suggested by a Bayesian Nash equilibrium $s^*(\cdot) = (s_1^*(\cdot), \ldots, s_n^*(\cdot))$ of the underlying Bayesian game. In such a case, the social utility that would be realized by the social planner for every possible type profile θ of the agents is given by

$$w(u_1(f(s^*(\theta)), \theta_1), \ldots, u_n(f(s^*(\theta)), \theta_n)) \tag{25}$$

In some instances, the above Bayesian Nash equilibrium may turn out to be a dominant strategy equilibrium. Better still, truth revelation by all agents could turn out to be a Bayesian Nash equilibrium or a dominant strategy equilibrium.

15.2 Optimal Mechanism Design Problem

In view of the above notion of social welfare function, it is clear that the objective of a social planner would be to look for a social choice function $f(\cdot)$ that would maximize the expected social utility for a given social welfare function $w(\cdot)$. However, being the social planner, it is always expected of him to be fair to all the agents. Therefore, the social planner would first put a few fairness constraints on the set of social choice functions which he can probably choose from. The fairness constraints may include any combination of all the previously studied properties of a social choice function, such as ex-post efficiency, incentive compatibility, and individual rationality. This set of social choice functions is known as *set of feasible social choice functions* and is denoted by F. Thus, the problem of a social planner can now be cast as an optimization problem where the objective is to maximize the expected social utility and the constraint is that the social choice function must be chosen from the feasible set F. This problem is known as the *optimal mechanism design* problem and the solution of the problem is some social choice function $f^*(\cdot) \in F$ which is used to define the optimal mechanism $\mathscr{D}^* = ((\Theta_i)_{i \in N}, f^*(\cdot))$ for the problem that is being studied.

Depending on whether the agents are loyal or autonomous entities, the optimal mechanism design problem may take two different forms.

$$\underset{f(\cdot) \in F}{\text{maximize}} \; E_\theta \left[w(u_1(f(\theta), \theta_1), \ldots, u_n(f(\theta), \theta_n)) \right] \tag{26}$$

$$\underset{f(\cdot) \in F}{\text{maximize}} \; E_\theta \left[w(u_1(f(s^*(\theta)), \theta_1), \ldots, u_n(f(s^*(\theta)), \theta_n)) \right] \tag{27}$$

The problem (26) is relevant when the agents are loyal and always reveal their true types whereas the problem (27) is relevant when the agents are rational. At this point of time, one may ask how to define the set of feasible social choice functions F. There is no unique definition of this set. The set of feasible social choice functions is a subjective judgment of the social planner. The choice of the set F depends on what all fairness properties the social planner would wish to have in the optimal social choice function $f^*(\cdot)$. If we define

$$\begin{aligned}
F_{DSIC} &= \{f : \Theta \to X | f(\cdot) \text{ is dominant strategy incentive compatible}\} \\
F_{BIC} &= \{f : \Theta \to X | f(\cdot) \text{ is Bayesian incentive compatible}\} \\
F_{ExPostIR} &= \{f : \Theta \to X | f(\cdot) \text{ is ex-post individual rational}\} \\
F_{IntIR} &= \{f : \Theta \to X | f(\cdot) \text{ is interim individual rational}\} \\
F_{ExAnteIR} &= \{f : \Theta \to X | f(\cdot) \text{ is ex-ante individual rational}\} \\
F_{Ex-AnteEff} &= \{f : \Theta \to X | f(\cdot) \text{ is ex-ante efficient}\} \\
F_{IntEff} &= \{f : \Theta \to X | f(\cdot) \text{ is interim efficient}\} \\
F_{Ex-PostEff} &= \{f : \Theta \to X | f(\cdot) \text{ is ex post efficient}\}
\end{aligned}$$

The set of feasible social choice functions F may be either any one of the above sets or intersection of any combination of the above sets. For example, the social planner may choose $F = F_{BIC} \bigcap F_{IntIR}$. In the literature, this particular feasible set is known as *incentive feasible set* due to Myerson [50]. Also, note that if the agents are loyal then the sets F_{DSIC} and F_{BIC} will be equal to the whole set of all the social choice functions.

If the environment is quasi-linear, then we can also define the set of allocatively efficient social choice functions F_{AE} and the set of budget balanced social choice functions F_{BB}. In such an environment, we will have $F_{Ex-PostEff} = F_{AE} \bigcap F_{BB}$.

15.3 Myerson's Optimal Auction: An Example of Optimal Mechanism

Let us consider Example 3.3 of single unit - single item auction without reserve price and discuss an optimal mechanism developed by Myerson [37]. The objective function here is to maximize the auctioneer's revenue.

Recall that each bidder i's type lies in an interval $\Theta_i = [\underline{\theta_i}, \overline{\theta_i}]$. We impose the following additional conditions on the environment.

1. The auctioneer and the bidders are risk neutral
2. Bidders' types are statistically independent, that is, the joint density $\phi(\cdot)$ has the form $\phi_1(\cdot) \times \ldots \times \phi_n(\cdot)$

3. $\phi_i(\cdot) > 0 \; \forall \; i = 1, \ldots, n$

4. We generalize the outcome set X relative to that considered in Example 3.3 by allowing a random assignment of the good. Thus, we now take $y_i(\theta)$ to be buyer i's probability of getting the good when the vector of announced types is $\theta = (\theta_1, \ldots, \theta_n)$. Thus, the new outcome set is given by

$$\begin{aligned} X &= \{(y_0, y_1 \ldots, y_n, t_0, t_1, \ldots, t_n) | y_0 = 0, t_0 \geq 0, y_i \in [0,1], t_i \leq 0 \; \forall \; i = 1, \ldots, n, \\ &\quad \sum_{i=0}^{n} y_i \leq 1, \sum_{i=0}^{n} t_i = 0 \Big\} \end{aligned}$$

Recall that the utility functions of the agents in this example are given by

$$u_i(f(\theta), \theta_i) = u_i(y_0(\theta), \ldots, y_n(\theta), t_0(\theta), \ldots, t_n(\theta), \theta_i) = \theta_i y_i(\theta) + t_i(\theta) \; \forall \; i = 1, \ldots, n$$

Thus, viewing $y_i(\theta) = v_i(k(\theta))$ in conjunction with the second and third conditions above, we can claim that the underlying environment here is linear.

In the above example, we assume that the auctioneer is the social planner and he is looking for an optimal direct revelation mechanism to sell the good. Myerson's [37] idea was that the auctioneer must use a social choice function which is Bayesian incentive compatible and interim individual rational and at the same time fetches the maximum revenue to the auctioneer. Thus, in this problem, the set of feasible social choice functions is given by $F = F_{BIC} \bigcap F_{InterimIR}$. The objective function in this case would be to maximize the total expected revenue of the seller which would be given by

$$E_\theta \left[w(u_1(f(\theta), \theta_1), \ldots, u_n(f(\theta), \theta_n)) \right] = -E_\theta \left[\sum_{i=1}^{n} t_i(\theta) \right]$$

Note that in above objective function we have used $f(\theta)$ not $f(s^*(\theta))$. This is because in the set of feasible social choice functions we are considering only BIC social choice functions and for these functions we have $s^*(\theta) = \theta \quad \forall \; \theta \in \Theta$. Thus, Myerson's optimal auction design problem can be formulated as the following optimization problem.

$$\underset{f(\cdot) \in F}{\text{maximize}} \quad - E_\theta \left[\sum_{i=1}^{n} t_i(\theta) \right] \tag{28}$$

where

$$F = \{ f(\cdot) = (y_1(\cdot), \ldots, y_n(\cdot), t_1(\cdot), \ldots, t_n(\cdot)) | f(\cdot) \text{ is BIC and interim IR} \}$$

By invoking Myerson's Characterization Theorem 12.2 for BIC SCF in linear environment, we can say that an SCF $f(\cdot)$ in the above context would be BIC iff it satisfies the following two conditions

1. $\overline{y_i}(\cdot)$ is non-decreasing for all $i = 1, \ldots, n$

2. $U_i(\theta_i) = U_i(\underline{\theta_i}) + \int_{\underline{\theta_i}}^{\theta_i} \overline{y_i}(s) ds \;\; \forall \; \theta_i \in \Theta_i; \; \forall \; i = 1, \ldots, n$

Also, we can invoke the definition of interim individual rationality to claim that the an SCF $f(\cdot)$ in the above context would be interim IR iff it satisfies the following conditions

$$U_i(\theta_i) \geq 0 \;\; \forall \theta_i \in \Theta_i; \; \forall \; i = 1, \ldots, n$$

where

- $\overline{t_i}(\hat{\theta}_i) = E_{\theta_{-i}}[t_i(\hat{\theta}_i, \theta_{-i})]$ be bidder i's expected transfer given that he announces his type to be $\hat{\theta}_i$ and that all bidders $j \neq i$ truthfully reveal their types.
- $\overline{y_i}(\hat{\theta}_i) = E_{\theta_{-i}}[y_i(\hat{\theta}_i, \theta_{-i})]$ is the probability that bidder i would receive the object given that he announces his type to be $\hat{\theta}_i$ and all bidders $j \neq i$ truthfully reveal their types.
- $U_i(\theta_i) = \theta_i \overline{y_i}(\theta_i) + \overline{t_i}(\theta_i)$

In view of the above paraphernalia, problem (28) can be rewritten as follows.

$$\underset{(y_i(\cdot), U_i(\cdot))_{i \in N}}{\text{maximize}} \quad \sum_{i=1}^{n} \int_{\underline{\theta_i}}^{\overline{\theta_i}} \left(\theta_i \overline{y_i}(\theta_i) - U_i(\theta_i)\right) \phi_i(\theta_i) d\theta_i \tag{29}$$

subject to

(i) $\overline{y_i}(\cdot)$ is non-decreasing $\forall\, i = 1, \ldots, n$

(ii) $y_i(\theta) \in [0,1], \sum_{i=1}^{n} y_i(\theta) \leq 1\ \forall i = 1, \ldots, n, \forall\, \theta \in \Theta$

(iii) $U_i(\theta_i) = U_i(\underline{\theta_i}) + \int_{\underline{\theta_i}}^{\theta_i} \overline{y_i}(s) ds \ \ \forall\, \theta_i \in \Theta_i;\ \forall\, i = 1, \ldots, n$

(iv) $U_i(\theta_i) \geq 0\ \forall\, \theta_i \in \Theta_i;\ \forall\, i = 1, \ldots, n$

We first note that if constraint (iii) is satisfied then constraint (iv) will be satisfied iff $U_i(\underline{\theta_i}) \geq 0\ \ \forall\, i = 1, \ldots, n$. As a result, we can replace the constraint (iv) with

(iv') $U_i(\underline{\theta_i}) \geq 0\ \forall\, i = 1, \ldots, n$

Next, substituting for $U_i(\theta_i)$ in the objective function from constraint (iii), we get

$$\sum_{i=1}^{n} \int_{\underline{\theta_i}}^{\overline{\theta_i}} \left(\theta_i \overline{y_i}(\theta_i) - U_i(\underline{\theta_i}) - \int_{\underline{\theta_i}}^{\theta_i} \overline{y_i}(s) ds \right) \phi_i(\theta_i) d\theta_i$$

Integrating by parts the above expression, the auctioneer's problem can be written as one of choosing the $y_i(\cdot)$ functions and the values $U_1(\underline{\theta_1}), \ldots, U_n(\underline{\theta_n})$ to maximize

$$\int_{\underline{\theta_1}}^{\overline{\theta_1}} \cdots \int_{\underline{\theta_n}}^{\overline{\theta_n}} \left[\sum_{i=1}^{n} y_i(\theta_i) J_i(\theta_i) \right] \left[\prod_{i=1}^{n} \phi_i(\theta_i) \right] d\theta_n \ldots d\theta_1 - \sum_{i=1}^{n} U_i(\underline{\theta_i})$$

subject to constraints (i), (ii), and (iv'), where

$$J_i(\theta_i) = \left(\theta_i - \frac{1 - \Phi_i(\theta_i)}{\phi_i(\theta_i)} \right) = \left(\theta_i - \frac{\overline{\Phi_i}(\theta_i)}{\phi_i(\theta_i)} \right)$$

where, we define $\overline{\Phi_i}(\theta_i) = 1 - \Phi_i(\theta_i)$. It is evident that solution must have $U_i(\underline{\theta_i}) = 0$ for all $i = 1, \ldots, n$. Hence, the auctioneer's problem reduces to choosing functions $y_i(\cdot)$ to maximize

$$\int_{\underline{\theta_1}}^{\overline{\theta_1}} \cdots \int_{\underline{\theta_n}}^{\overline{\theta_n}} \left[\sum_{i=1}^{n} y_i(\theta_i) J_i(\theta_i) \right] \left[\prod_{i=1}^{n} \phi_i(\theta_i) \right] d\theta_n \ldots d\theta_1$$

subject to constraints (i) and (ii).

Let us ignore constraint (i) for the moment. Then inspection of the above expression indicates that $y_i(\cdot)$ is a solution to this relaxed problem iff for all $i = 1, \ldots, n$, we have

$$y_i(\theta) = \begin{cases} 0 & : \text{ if } J_i(\theta_i) < \max\{0, \max_{h \neq i} J_h(\theta_h)\} \\ 1 & : \text{ if } J_i(\theta_i) > \max\{0, \max_{h \neq i} J_h(\theta_h)\} \end{cases} \tag{30}$$

Note that $J_i(\theta_i) = \max\{0, \max_{h \neq i} J_h(\theta_h)\}$ is a zero probability event.

In other words, if we ignore the constraint (i) then $y_i(\cdot)$ is a solution to this relaxed problem iff the good is allocated to a bidder who has highest non-negative vale for $J_i(\theta_i)$. Now, recall the definition of $\overline{y_i}(\cdot)$. It is easy to write down the following expression

$$\overline{y_i}(\theta_i) \quad = \quad E_{\theta_{-i}}\left[y_i(\theta_i, \theta_{-i})\right] \tag{31}$$

Now, if we assume that $J_i(\cdot)$ is non-decreasing in θ_i then it is easy to see that above solution $y_i(\cdot)$, given by (30), will be non-decreasing in θ_i, which in turn implies, by looking at expression (31), that $\overline{y_i}(\cdot)$ is non-decreasing in θ_i. Thus, the solution to this relaxed problem actually satisfies constraint (i) under the assumption that $J_i(\cdot)$ is non-decreasing. Assuming that $J_i(\cdot)$ is non-decreasing, the solution given by (30) seems to be the solution of the optimal mechanism design problem for single unit- single item auction. The condition that $J_i(\cdot)$ is non-decreasing in θ_i is met by most of the distribution functions such as Uniform and Exponential.

So far we have computed the allocation rule for the optimal mechanism and now we turn out attention towards the payment rule. The optimal payment rule $t_i(\cdot)$ must be chosen in such a way that it satisfies

$$\overline{t_i}(\theta_i) = E_{\theta_{-i}}[t_i(\theta_i, \theta_{-i})] = U_i\left(\theta_i\right) - \theta_i \overline{y_i}\left(\theta_i\right) = \int_{\underline{\theta_i}}^{\theta_i} \overline{y_i}(s)ds - \theta_i \overline{y_i}\left(\theta_i\right) \tag{32}$$

Looking at the above formula, we can say that if the payment rule $t_i(\cdot)$ satisfies the following formula (33), then it would also satisfy the formula (32).

$$t_i(\theta_i, \theta_{-i}) = \int_{\underline{\theta_i}}^{\theta_i} y_i(s, \theta_{-i})ds - \theta_i y_i\left(\theta_i, \theta_{-i}\right) \quad \forall\, \theta \in \Theta \tag{33}$$

The above formula can be rewritten more intuitively, as follows. For any vector θ_{-i}, let we define

$$z_i(\theta_{-i}) = \inf\{\theta_i | J_i(\theta_i) > 0 \text{ and } J_i(\theta_i) \geq J_j(\theta_j)\ \forall\, j \neq i\}$$

Then $z_i(\theta_{-i})$ is the infimum of all winning bids for bidder i against θ_{-i}, so

$$y_i(\theta_i, \theta_{-i}) = \begin{cases} 1 & : \text{ if } \theta_i > z_i(\theta_{-i}) \\ 0 & : \text{ if } \theta_i < z_i(\theta_{-i}) \end{cases}$$

This gives us

$$\int_{\underline{\theta_i}}^{\theta_i} y_i(s, \theta_{-i})ds = \begin{cases} \theta_i - z_i(\theta_{-i}) & : \text{ if } \theta_i \geq z_i(\theta_{-i}) \\ 0 & : \text{ if } \theta_i < z_i(\theta_{-i}) \end{cases}$$

Finally, the formula (33) becomes

$$t_i(\theta_i, \theta_{-i}) = \begin{cases} -z_i(\theta_{-i}) & : \text{ if } \theta_i \geq z_i(\theta_{-i}) \\ 0 & : \text{ if } \theta_i < z_i(\theta_{-i}) \end{cases}$$

That is bidder i must pay only when he gets the good, and then he pays the amount equal to his lowest possible winning bid.

A few interesting observations are worth mentioning here.

1. When the various bidders have differing distribution function $\Phi_i(\cdot)$ then, the bidder who has the largest value of $J_i(\theta_i)$ is *not* necessarily the bidder who has bid the highest amount for the good. Thus Myerson's optimal auction need not be allocatively efficient and therefore, need not be ex-post efficient.

2. If the bidders are symmetric, that is,

   - $\Theta_1 = \ldots = \Theta_n = \Theta$
   - $\Phi_1(\cdot) = \ldots = \Phi_n(\cdot) = \Phi(\cdot)$

   then the allocation rule would be precisely the same allocation rule of first-price and second-price auctions. In such a case the object would be allocated to the highest bidder. In such a situation, the optimal auction would also become allocatively efficient. Also, note that in such a case the payment rule that we described above would coincide with the payment rules in second-price auction. In other words, the second price (Vickrey) auction would be the optimal auction when the bidders are symmetric. Therefore, many a time, the optimal auction is also known as *modified Vickrey auction.*

15.4 Recent Work in Optimal Mechanism Design

Optimal mechanism design is an active area of research since past couple of years. The recent contributions in this direction include Krishna and Perry [51], where the authors consider the same problem of single unit - single item auction without reserve price and design an optimal mechanism for it. The objective function of Krishna and Perry's model (KP model for short) is the same as Myerson's optimal auction but the set of feasible social choice function is different. Krishna and Perry insists on having allocative efficiency also in addition to interim IR and BIC. Thus in the KP model the set F is given by $F = F_{AE} \bigcap F_{BIC} \bigcap F_{InterimIR}$. Krishna and Perry prefer to call such a mechanism as *efficient mechanism.*

In a recent paper, Malakhov and Vohra [52] offer a comparison between the optimal auction problem with continuous types and with discrete types. This paper highlights connections between the discrete and continuous approaches to optimal auction design with single and multi-dimensional types. This paper provides an interpretation of an optimal auction design problem in terms of a linear program that is an instance of a parametric shortest path problem on a lattice.

In another recent paper, Malakhov and Vohra [52] consider the problem of a seller with limited supply facing a group of agents whose private information is two dimensional. Each agent has a constant marginal value for the good up to some capacity, thereafter it is zero. Both the marginal value and the capacity are private information. The authors have shown that an optimal auction design problem for this setting can be interpreted in terms of a linear program that is an instance of a parametric shortest path problem on a lattice. The authors have described the revenue maximizing Bayesian incentive compatible auction for this environment.

16 To Probe Further

For a more detailed treatment of mechanism design, the readers are requested to refer to textbooks, such as the ones by Mas-Colell, Whinston, and Green [38] (Chapter 23), Green and Laffont [43], and Laffont [53]. Readers can also refer to two survey articles on mechanism theory and implementation theory written by Jackson in the recent past [54, 55].

The current paper is not to be treated as a survey on auctions in general. There are widely popular books (for example, by Milgrom [48] and Krishna [49]) and surveys on auctions (for example, [47, 56, 57, 58, 59]) which deal with auctions in a comprehensive way.

The current paper is also not to be treated as a survey on combinatorial auctions (currently an active area of research). Exclusive surveys on combinatorial auctions include the articles by de Vries and Vohra [60, 61], Pekec and Rothkopf [62], and Narahari and Pankaj Dayama [30]. Cramton, Ausubel, and Steinberg [26] have recently brought out a comprehensive edited volume containing expository and survey articles on varied aspects of combinatorial auctions.

References

[1] N. Nisan. Algorithms for selfish agents, mechanism design for distributed computation. In *16th Symposium on Theoretical Aspects of Computer Science (STACS'99)*, pages 1–15, Heidelberg, 1999.

[2] T. Roughgarden. *Selfish Routing*. PhD thesis, Graduate School, Cornell University, May 2002.

[3] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. The MIT Press, 2005.

[4] J. Feigenbaum, C. H. Papadimitriou, R. Sami, and S. Shenker. A BGP-based mechanism for lowest-cost routing. In *21st ACM Symposium on Principles of Distributed Computing (PODC'02)*, pages 173–182, New York, June 30-July 4 2002.

[5] J. Feigenbaum and S. Shenker. Distributed algorithmic mechanism design: Recent results and future directions. In *6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications (MobiCom'02)*, pages 1–13, New York, 2002.

[6] J. Hershberger and S. Suri. Vickrey prices and shortest paths: What is an edge worth? In *42nd Annual IEEE Symposium on Foundations of Computer Science (FOCS'01)*, pages 252–259, 2001.

[7] N. Nisan and A. Ronen. Computationally feasible VCG mechanisms. In *2nd ACM Conference on Electronic Commerce (EC'00)*, pages 242–252, New York, 2000.

[8] N. Nisan and A. Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35:166–196, 2001.

[9] L. Anderegg and Eidenbenz. Ad hoc-VCG: A truthful and cost-efficient routing protocol for mobile ad hoc networks with selfish agents. In *9th ACM Annual International Conference on Mobile Computing and Networking (MobiCom'03)*, pages 245–259, San Diego, California, USA, September 14-19 2003.

[10] S. Eidenbenz, P. Santi, and G. Resta. COMMIT: A sender centric truthful and energy-efficient routing protocol for ad hoc networks. In *Workshop on Wireless, Mobile, and Ad hoc Networks*

(WMAN) in conjunction with 19th IEEE International Parallel and Distributed Processing Symposium (IPDPS'05), 2005.

[11] W. Wang and Xiang-Yang Li. Truthful multicast in selfish wireless networks. In *10th ACM Annual International Conference on Mobile Computing and Networking (MobiCom'04)*, Pennsylvania, 2004.

[12] N. Rama Suri. Design of incentive compatible broadcast protocols for wireless ad-hoc networks. Technical report, Master's Dissertation, Department of Computer Science and Automation, Indian Institute of Science, Bangalore, India, May 2006.

[13] D. Grosu and A. Chronopoulos. A load balancing mechanism with verification. In *17th International Parallel and Distributed Processing Symposium (IPDPS'03)*, Nice, France, 2003.

[14] D. Grosu and A. Chronopoulos. Algorithmic mechanism design for load balancing in distributed systems. In *18th International Parallel and Distributed Processing Symposium (IPDPS'04)*, Santa Fe, New Mexico, 2004.

[15] A. Das and D. Grosu. Combinatorial auction-based protocols for resource allocation in grids. In *19th International Parallel and Distributed Processing Symposium (IPDPS'05)*, Denver, CA, 2005.

[16] M. Wellman, W. Walsh, P. Wurman, and J. MacKie-Mason. Some economics of market based distributed scheduling. In *18th International Conference on Distributed Computing Systems (ICDCS'98)*, Amsterdam, The Netherlands, 1998.

[17] R. Buyya. *Economic-based Distributed Resource Management and Scheduling for Grid Computing.* PhD thesis, School of Computer Science and Software Engineering, Monash University, Australia, 2002.

[18] H. Prakash V. and Y. Narahari. A strategy-proof auction mechanism for grid scheduling with selfish entities. In *2nd International Conference on Web Information Systems and Technologies (WEBIST'05)*, pages 178–183, Satubal, Portugal, April 2005.

[19] M. Naor. Cryptography and mechanism design. In 8^{th} *Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, Certosa di Pontignano, University of Siena, Italy, July 2001.

[20] H. R. Varian. Economic mechanism design for computerized agents. In *1st USENIX Workshop on Electronic Commerce*, 1995.

[21] M. Naor, B. Pinkas, and R. Sumner. Privacy preserving auctions and mechanism design. In *1st ACM Conference on Electronic Commerce (EC'99)*, 1999.

[22] F. Brandt. Social choice and preference protection - towards fully private mechanism design. In *4th ACM Conference on Electronic Commerce (EC'03)*, pages 220–221, 2003.

[23] A. Altman and M. Tennenholtz. Ranking systems: The pagerank axioms. In *6th ACM Conference on Electronic Commerce (EC'05)*, pages 1–8, 2005.

[24] A. Altman and M. Tennenholtz. Quantifying incentive compatibility of ranking systems. Technical report, Technion, Israel, March 2006.

[25] Z. Gyongyi and H. Garcia-Molina. Link spam alliances. In *31st International Conference on Very Large Data Bases (VLDB'05)*, Trondheim, Norway, 2005.

[26] P. Cramton. Simultaneous ascending auctions. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*. The MIT Press, Cambridge, Massachusetts, USA, 2005.

[27] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. In *2nd Workshop on Sponsored Search Auctions in conjunction with the ACM Conference on Electronic Commerce (EC'06)*, Ann Arbor, MI, June 2006.

[28] D. Garg. *Design of Innovative Mechanisms for Contemporary Game Theoretic Problems in Electronic Commerce*. PhD thesis, Department of Computer Science and Automation, Indian Institute of Science, Bangalore, India, May 2006.

[29] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *7th ACM Conference on Electronic Commerce (EC'06)*, Ann Arbor, Michigan, USA, June 2006.

[30] Y. Narahari and P. Dayama. Combinatorial auctions for electronic business. *Sadhana: Indian academy Proceedings in Engineering Sciences*, 30(2-3):179–212, 2005.

[31] M. Eso, J. Kalagnanam, L. Ladanyi, and Y.G. Li. Winner determination in bandwidth exchanges. Technical report, IBM TJ Watson Research Center, 2001.

[32] T.S. Chandrashekar, Y. Narahari, C. H. Rosa, D. Kulkarni, J.D. Tew, and P. Dayama. Auction based mechanisms for electronic procurement. *IEEE Transactions on Automation Science and Engineering*, 3(6), 2006.

[33] C. Caplice and Y. Sheffi. Combinatorial auctions for truckload transportation. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*. The MIT Press, Cambridge, Massachusetts, USA, 2005.

[34] W. E. Walsh, M. P. Wellman, and F. Ygge. Combinatorial auctions for supply chain formation. In *2nd ACM Conference on Electronic Commerce (EC'00)*, pages 260–269, Minneapolis, Minnesota, 2000.

[35] W. E. Walsh and M. P. Wellman. Decentralized supply chain formation: A market protocol and competitive equilibrium analysis. *Journal of Artificial Intelligence Research*, 19:513–567, 2003.

[36] S. Karabuk and D.S. Wu. Incentive schemes for semiconductor capacity allocation: A game theoretic analysis. Technical report, Department of Industrial and Systems Engineering, Lehigh University, August 2004.

[37] R. B. Myerson. Optimal auction design. *Math. Operations Res.*, 6(1):58–73, Feb. 1981.

[38] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.

[39] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, March 1961.

[40] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–601, 1973.

[41] M. A. Satterthwaite. Strategy-proofness and arrow's conditions: Existence and correspondence theorem for voting procedure and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.

[42] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[43] J. R. Green and J. J. Laffont. *Incentives in Public Decision Making*. North-Holland Publishing Company, Amsterdam, 1979.

[44] E. Clarke. Multi-part pricing of public goods. *Public Choice*, 11:17–23, 1971.

[45] C. d'Aspremont and L.A. Gérard-Varet. Incentives and incomplete information. *Journal of Public Economics*, 11:25–45, 1979.

[46] K. Arrow. The property rights doctrine and demand revelation under incomplete information. In M. Boskin, editor, *Economics and Human Welfare*. Academic Press, New York, 1979.

[47] P. R. McAfee and J. McMillan. Auctions and bidding. *Journal of Economic Literature*, 25(2):699–738, June 1987.

[48] P. Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004.

[49] V. Krishna. *Auction Theory*. Academic Press, 2002.

[50] R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, Massachusetts, 1997.

[51] V. Krishna and M. Perry. Efficient mechanism design. May 2000.

[52] A. Malakhov and R. V. Vohra. Single and multi-dimensional optimal auctions - a network approach. December 2004.

[53] J.J. Laffont. *Fundamentals of Public Economics*. The MIT Press, Cambridge, 1988.

[54] M. O. Jackson. A crash course in implementation theory. *Social Choice and Welfare*, 18:655–708, 2001.

[55] M. O. Jackson. Mechanism theory. In Ulrich Derigs, editor, *Optimizations and Operations Research*. Oxford, U.K., 2003.

[56] P. Milgrom. Auctions and bidding: A primer. *Journal of Economic Perspectives*, 3(3):3–22, 1989.

[57] Paul Klemperer. *Auctions: Theory and Practice. The Toulouse Lectures in Economics*. Princeton University Press, 2004.

[58] E. Wolfstetter. Auctions: An introduction. *Economic Surveys*, 10:367–421, 1996.

[59] Jayant Kalagnanam and David C Parkes. Auctions, bidding, and exchange design. In David Simchi-Levi, David S Wu, and Max Shen, editors, *Handbook of Quantitative Supply Chain Analysis: Modeling in the E-Business Era*, Int. Series in Operations Research and Management Science. Kluwer Academic Publishers, 2005.

[60] S. de Vries and R. V. Vohra. Combinatorial auctions: A survey. *INFORMS Journal of Computing*, 15(1):284–309, 2003.

[61] Sven de Vries and Rakesh H. Vohra. Design of combinatorial auctions. In *Handbook of Quantitative Supply Chain Analysis: Modeling in the E-Business Era*, pages 247–292. International Series in Operations Research and Management Science, Kluwer Academic Publishers, Norwell, MA, USA, 2005.

[62] A. Pekec and M.H. Rothkopf. Combinatorial auction design. *Management Science*, 49:1485–1503, 2003.