Design of Incentive Compatible Mechanisms for Stackelberg Problems

Dinesh Garg[1] and Yadati Narahari[⋆2]

[1] Department of Computer Science and Automation, Indian Institute of Science, Bangalore - 560 012, India
dgarg@csa.iisc.ernet.in

[2] Department of Computer Science and Automation, Indian Institute of Science, Bangalore - 560 012, India
Tel: +91-80-22932773
hari@csa.iisc.ernet.in

Abstract. This paper takes the first steps towards designing incentive compatible mechanisms for hierarchical decision making problems involving selfish agents. We call these Stackelberg problems. These are problems where the decisions or actions in successive layers of the hierarchy are taken in a sequential way while decisions or actions within each layer are taken in a simultaneous manner. There are many immediate applications of these problems in distributed computing, grid computing, network routing, ad hoc networks, electronic commerce, and distributed artificial intelligence. We consider a special class of Stackelberg problems called SLRF (Single Leader Rest Followers) problems and investigate the design of incentive compatible mechanisms for these problems. In developing our approach, we are guided by the classical theory of mechanism design. To illustrate the design of incentive compatible mechanisms for Stackelberg problems, we consider first-price and second-price electronic procurement auctions with reserve prices. Using the proposed framework, we derive some interesting results regarding incentive compatibility of these two mechanisms.

1 Mechanism Design and Stackelberg Problems

The *Theory of Mechanism Design* is an important discipline in the area of Welfare Economics. The area of Welfare Economics is concerned with settings where a 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 theory of mechanism design aims at studying how this privately held information can be elicited [2, 4, 7]. The state-of-the-art literature on mechanism design theory deals with situations where individuals are symmetric, that is to say, no single individual dominates the decision process. However, there are situations arising in Welfare Economics, Sociology, Engineering, Operations

⋆ Corresponding author.

X. Deng and Y. Ye (Eds.): WINE 2005, LNCS 3828, pp. 718–727, 2005.

Research, Control Theory, and Computer Science where individuals take decisions in a hierarchical manner. A simple example is that of a situation wherein one of the individuals (or a group of individuals), called *leader(s)*, has the ability to enforce his preference on the other individual(s), called *follower(s)*. In such problems, the policy maker first invites the leaders to reveal their privately held information in a simultaneous manner. After receiving this information the policy maker broadcasts it among the followers and the followers respond to this by revealing their preferences in a simultaneous manner. After receiving the preferences from all the individuals, the policy maker aggregates the information into a social decision. The problem faced by the individuals in such a situation can be naturally modeled as a *Stackelberg game* following the seminal work of Stackelberg [15]. Following are some interesting examples where one can see these problems arising naturally.

- Task allocation in parallel/distributed systems
- Scheduling in grids
- Internet routing
- Admission, routing, and scheduling in telecom networks
- Flow control, routing, and sequencing in manufacturing systems
- Auctions in electronic commerce

In a recent work, Roughgarden [14] considered the problem of job shop scheduling, where fraction of the jobs are scheduled by a centralized authority (leader) and the remaining jobs are scheduled by selfish users (followers). He modeled this scheduling problem as Stackelberg game and showed that it is NP-hard to compute Stackelberg strategies. The underlying Stackelberg game in this problem is complete information game and there is no privately held information by the players which we require to elicit truthfully. Thus, the problem considered in this paper is about computing the optimal strategies of the leader and the followers and not about designing a mechanism for Stackelberg problem.

1.1 Contributions and Outline of the Paper

The major contributions of this paper are as follows.

- We investigate the mechanism design problem for SLRF games [1]. In this new framework, we define the notion of *Bayesian Stackelberg Incentive Compatible (BaSIC)* social choice functions. To the best of our knowledge, this is the first time mechanisms are being investigated in the context of Stackelberg problems.
- To illustrate our approach, we investigate the Bayesian Stackelberg incentive compatibility of *first-price and second-price procurement auctions with reserve prices.* We obtain two key results in this regard. The first result shows that in the first-price auction with reserve prices, the social choice function is BaSIC for the buyer but not for the sellers. The second result shows that in the second-price auction with reserve prices, the social choice function is BaSIC for the sellers but not for the buyer.

[1] We keep using the phrases SLRF games and SLRF problems interchangeably.

The organization of the paper is as follows. Section 2 presents a crisp review of relevant concepts in Stackelberg games. A more detailed treatment can be found in [1]. In Section 3, we motivate the Stackelberg mechanism design problems by means of two examples - *first-price procurement auction with reserve prices (F-PAR) and second-price procurement auction with reserve prices (S-PAR).* We then describe the problem of designing incentive compatible mechanisms for SLRF problems in Section 4. In Section 5, we state two important results concerning the Bayesian Stackelberg incentive compatibility of the two mechanisms F-PAR and S-PAR. Due to paucity of space, we are unable to provide the proofs for these results. Interested readers are urged to look into our recent technical report [3].

2 Stackelberg Games

2.1 Stackelberg Games with Incomplete Information

To begin with, we consider the following *noncooperative finite game with incomplete information* in its strategic form (also called a Bayesian game [11]):

$$\Gamma^b = (N, (C_i)_{i \in N}, (\Theta_i)_{i \in N}, (\phi_i)_{i \in N}, (u_i)_{i \in N})$$

where $N = \{1, 2, \ldots, n\}$ is a nonempty set of players, and, for each i in N, C_i is a nonempty set of actions available to player i; Θ_i is a nonempty set of possible types of player i; $\phi_i : \Theta_i \mapsto \Delta\Theta_{-i}$ is a belief function which gives the subjective probability of player i about the types of the other players for a given type of his own; and $u_i : C \times \Theta \mapsto \Re$ is the utility function of player i, where $C = \times_{j \in N} C_j$, $\Theta = \times_{j \in N} \Theta_j$, and $\Theta_{-i} = \times_{j \in N - \{i\}} \Theta_j$. Note that ΔS for any set S is the set of all probability distributions over S. A pure strategy s_i for player i in the Bayesian game Γ^b is defined to be a function from Θ_i to C_i.

In the above description of the game Γ^b, it is an implicit assumption that all the players choose their actions simultaneously. However, it is possible to impose an additional structure of hierarchical decision making on this game where agents choose their actions in a sequential manner as suggested by the hierarchy. The hierarchy is defined as a sequence of nonempty pairwise disjoint subsets of players, $H = H_1, H_2, \ldots, H_h$, where h represents the total number of levels in the hierarchy. In this setup, after learning their types, first, all the players in hierarchy level 1, i.e. H_1, choose their actions simultaneously. The actions (but not the types) chosen by all these players are announced publicly to the rest of the players. Next, all the players in hierarchy level 2, i.e. H_2, choose their actions simultaneously and again the chosen actions by all these players are announced publicly to the rest of the players. This process continues until all the players announce their actions.

A Bayesian game Γ^b together with hierarchical decision making can be called a *Bayesian Stackelberg (BS) game* and would have the following representation.

$$\Gamma^b_s = ((H_j)_{j=1,\ldots,k}, (C_i)_{i \in N}, (\Theta_i)_{i \in N}, (\phi_i)_{i \in N}, (u_i)_{i \in N})$$

If ($h = 2$), that is, the players are divided into two levels of hierarchy - H_1 and H_2, the games are referred to as *Leader-Follower Games*. The players in H_1 are called *leaders* and the players in H_2 are called *followers*. The followers, having observed the actions taken by the leaders, choose their actions in a simultaneous manner. Within the class of leader-follower games, there is an interesting subclass of games where H_1 is a singleton set. Such games are called *single leader rest followers (SLRF) games*. In an *SLRF Bayesian Stackelberg game*, one player is declared as the leader and after learning her type, she first takes her action. The action taken by the leader becomes common knowledge among the followers but her type remains unknown to the followers. Followed by the action of the leader, all the followers, *who have already learned their types*, take their actions simultaneously.

2.2 Pure Strategy Bayesian Stackelberg Equilibrium for SLRF Games

In this section, we would like to characterize the solution of the SLRF Bayesian Stackelberg games. The natural choice for the solution of such games is a combination of Stackelberg equilibrium and Bayesian Nash equilibrium. We call such a solution as Bayesian Stackelberg (BS) equilibrium. In what follows is a characterization of BS equilibrium.

1. **The Set of Followers' Optimal Response Strategy Profiles.** Let us assume that after learning her type $\theta_l \in \Theta_l$, the leader takes an action $c_l \in C_l$. For any such action $c_l \in C_l$, the set $R(c_l)$ below, which is a set of pure strategy profiles of the followers, is called as the set of *followers' optimal response (or rational reaction) strategy profiles.*

$$R(c_l) = \{s_{-l} \in \times_{j \in N-\{l\}} \times_{\Theta_j} C_j | v_{\theta_j}(c_l, c_j, s_{-l,j}) \leq v_{\theta_j}(c_l, s_{-l}) \\ \forall c_j \in C_j \; \forall \theta_j \in \Theta_j \; \forall j \in N - \{l\}\} \quad (1)$$

   where (c_l, s_{-l}) is an action-strategy profile of the players in which the leader takes an action c_l and the followers take actions as suggested by the corresponding pure strategy for them in the profile (s_{-l}). Similarly, $(c_l, c_j, s_{-l,j})$ is an action-strategy profile in which the leader takes an action c_l, follower $j \in N$ takes an action c_j and rest of the followers take actions as suggested by the corresponding pure strategy for them in the profile (s_{-l}). The quantities $v_{\theta_j}(c_l, s_{-l})$ and $v_{\theta_j}(c_l, c_j, s_{-l,j})$ are the expected payoffs to the player j when his type is θ_j, and the players follow the action-strategy profile (c_l, s_{-l}) and $(c_l, c_j, s_{-l,j})$, respectively. Note that risk averse and risk neutral followers will always play an optimal response against any action taken by the leader.
2. **Secure Strategy Set of Leader.** Assuming that $R(c_l)$ is nonempty for each $c_l \in C_l$, we call a strategy $s^*_l \in S_l$ of the leader l to be a secure strategy if it satisfies the following security constraint for each $\theta_l \in \Theta_l$:

$$s^*_l(\theta_l) = \arg \max_{c_l \in C_l} \min_{s_{-l} \in R(c_l)} \sum_{\theta_{-l} \in \Theta_{-l}} \phi_l(\theta_{-l}|\theta_l) u_l(c_l, s_{-l}(\theta_{-l}), (\theta_l, \theta_{-l})) \quad (2)$$

Note that a risk averse and risk neutral leader will always play a secure strategy. An implicit assumption behind the above relation is that the game Γ_s^b is a finite game. However, if we allow the sets Θ_i to be infinite, then the relation (2) gets modified in the following manner:

$$s_l^*(\theta_l) = \arg \max_{c_l \in C_l} \min_{s_{-l} \in R(c_l)} E_{\theta_{-l}}[u_l(c_l, s_{-l}(\theta_{-l}), (\theta_l, \theta_{-l}))|\theta_l] \quad (3)$$

Further if we allow the sets $(C_i)_{i \in N}$ to be infinite then in the above relation, we need to replace max and min by sup and inf respectively.

3. **Bayesian Stackelberg Equilibrium.** A strategy profile $s^* = \left(s_l^*, t_{-l}^*\right)$ is said to be a Bayesian Stackelberg equilibrium if s_l^* is a secure strategy for the leader and $t_{-l}^* : C_l \mapsto \cup_{c_l \in C_l} R(c_l)$ is a rational reaction strategy of followers against s_l^*, that is $t_{-l}^*(s_l^*(\theta_l)) \in R(s_l^*(\theta_l))\ \forall \theta_l \in \Theta_l$. One can also define the Bayesian Stackelberg equilibrium in mixed strategies for SLRF Bayesian Stackelberg games. However, we omit this because it is not required for this paper.

3 Stackelberg Mechanism Design Problems: Motivating Example

Consider an electronic procurement marketplace where a buyer b registers himself and wishes to procure a single indivisible object. There are n potential sellers, indexed by $i = 1, 2, \ldots, n$, who also register themselves with the marketplace. We make the following assumptions, which are quite standard in the existing literature on auction theory. A comprehensive discussion about these assumptions can be found in [9, 8, 17].

A1: Risk Neutral Bidders: The buyer and all the n sellers are risk neutral.

A2: Independent Private Value (IPV) Model: Each seller i, and buyer b draw their valuations θ_i, and θ_b, respectively (which can be viewed as their types) from distribution $F_i(.)$ and $F_b(.)$, respectively. $F_i(.), i = 1, 2, \ldots, n$ and $F_b(.)$ are mutually independent. Let $\Theta_i, i = 1, 2, \ldots, n$ and Θ_b denote the set of all possible types of the sellers and buyer, respectively. This implies that $F_i(.), i = 1, 2, \ldots, n$ and $F_b(.)$ are probability distribution functions of the random variables $\Theta_i, , i = 1, 2, \ldots, n$ and Θ_b, respectively.

A3: Symmetry among Sellers: The sellers are symmetric in following sense:

$\Theta_1 = \Theta_2 = \ldots = \Theta_n = \Theta\ ; F_1(.) = F_2(.) = \ldots = F_n(.) = F(.)$

A4: Properties of $F(.)$ and Θ: We assume that $F(.)$ and Θ satisfy:

$\Theta = [\,\underline{\theta}\,, \overline{\theta}\,], \underline{\theta} > 0$

$F(.)$ is twice continuously differentiable

$f(\theta) = F'(\theta) > 0; \forall\ \underline{\theta} \leq \theta \leq \overline{\theta}$

The marketplace first invites the buyer to report his type. Based on his actual type θ_b, the buyer first reports his type, say, $\hat{\theta}_b \in \Theta_b$ to the marketplace. The declared type of the buyer, that is $\hat{\theta}_b$, is treated as the price above which the buyer is not willing to buy the object. This price is known as reserve price. The marketplace publicly announces this reserve price among all the sellers. Now, the sellers are invited to submit their bids (or types) to the marketplace. Based on actual type θ_i, each seller i bids (or reports) his type say $\hat{\theta}_i \in \Theta_i$. After receiving the bids from all the sellers, the marketplace determines the winning seller, the amount that will be paid to him, and the amount that will be paid by the buyer. These are called as *winner determination* and *payment rules*. In E-commerce, such a trading institution is known as *procurement auction with reserve prices (PAR)*. It is easy to see that the above problem is a Stackelberg problem and designing the winner determination and payment rules for this problem in a way that there is no incentive for buyer and sellers to reveal untruthful valuation is essentially a problem of designing an incentive compatible mechanism for Stackelberg problem. Depending on what winner determination and payment rules are employed by the marketplace, it may affect the incentive compatibility property. Following are two well known and existing mechanisms for PAR.

1. **First-Price Procurement Auction with Reserve Prices (F-PAR):** In this setting, the marketplace first discards all the received bids that fall above the reserve price announced by the buyer. Next, the seller whose bid is the lowest among the remaining bids is declared as the winner. The winner transfers the object to the buyer and the buyer pays to the winning seller an amount equal to his bid, that is $\hat{\theta}_i$. If there is no bid below the reserve price then no deal is struck. On the other hand, if there is a tie among the winning bids then the winner is chosen randomly, where each of the lowest valued bids has an equal chance of winning.
2. **Second-Price Procurement Auction with Reserve Prices (S-PAR):** In this setting, the winner determination rule is the same as F-PAR but the payment rules are slightly different. The winning seller transfers the object to the buyer and the buyer pays to him an amount equal to second lowest valued bid, if such a bid exists, otherwise an amount equal to the reserve price. Further, if there is no bid below the reserve price then no deal is struck. If there is a tie among winning bids, the winner is chosen randomly, where each of the lowest bids has an equal chance of winning.

It is easy to see that if the buyer announces a reserve price of $\hat{\theta}_b = \overline{\theta}$, then the procurement auction with reserve price will simply become the classical version of procurement auction with no reserve price.

4 Mechanism Design for SLRF Problems

A mechanism can be viewed as an institution, which a social planner deploys, to elicit the information from the agents about their types and then aggregate this information into a social outcome. Formally, a mechanism for an SLRF problem is a collection of action sets $(C_1, \ldots, C_n)$ and an outcome function $g : \times_{i \in N} C_i \mapsto X$, that is $\mathcal{M}_{\text{SLRF}} = ((C_i)_{i \in N}, g(.))$. A mechanism $\mathcal{M}_{\text{SLRF}}$ combined with possible types of the agents $(\Theta_1, \ldots, \Theta_n)$, probability density $\phi(.)$, Bernoulli utility functions $(u_1(.), \ldots, u_n(.))$, and description of leader agent l defines a Bayesian Stackelberg game Γ_s^b which is induced among the agents when the social planner invokes this mechanism as a means to solve the SLRF problem. The induced Bayesian Stackelberg game Γ_s^b is given by:

$$\Gamma_s^b = (\{l\}, N - \{l\}, (C_i)_{i \in N}, (\Theta_i)_{i \in N}, \phi(.), (\overline{u_i})_{i \in N})$$

where $\overline{u_i} : C \times \Theta \mapsto \Re$ is the utility function of agent i and is defined in the following manner: $\overline{u_i}(c, \theta) = u_i(g(c), \theta_i)$, where, we recall that $C = \times_{i \in N} C_i$, and $\Theta = \times_{i \in N} \Theta_i$.

4.1 Social Choice Function

A social choice function is a function $f : \times_{i \in N} \Theta_i \mapsto X$, 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' type $\theta = (\theta_1, \ldots, \theta_n)$. The set X is known as collective choice set or outcome set. For example, in the context of PAR, an outcome may be represented by a vector $x = (y_b, y_1, \ldots, y_n, t_b, t_1, \ldots, t_n)$, where $y_b = 1$ if the buyer receives the object, $y_b = 0$ otherwise, and t_b is the monetary transfer received by the buyer. Similarly, $y_i = -1$ if the seller i is the winner, $y_i = 0$ otherwise, and t_i is the monetary transfer received by the seller i. The set of feasible alternatives is then

$$X = \{(y_b, y_1, \ldots, y_n, t_b, t_1, \ldots, t_n) | y_b + \sum_{i=1}^{n} y_i = 0, t_b + \sum_{i=1}^{n} t_i \leq 0\}$$

In view of the above description, the general structure of the social choice function for PAR is

$$f(\theta_b, \theta) = (y_b(\theta_b, \theta), y_1(\theta_b, \theta), \ldots, y_n(\theta_b, \theta), t_b(\theta_b, \theta), t_1(\theta_b, \theta), \ldots, t_n(\theta_b, \theta)) \quad (4)$$

where $\theta = (\theta_1, \ldots, \theta_n)$. Note that $y_b(.)$, and $y_i(.)$ depend on the winner determination rule whereas $t_b(.)$ and $t_i(.)$ depend on the payment rule.

4.2 Implementing a Social Choice Function in Bayesian Stackelberg Equilibrium

We say that the mechanism $\mathcal{M}_{\text{SLRF}} = ((C_i)_{i \in N}, g(.))$ implements the social choice function $f : \times_{i \in N} \Theta_i \mapsto X$ in Bayesian Stackelberg equilibrium if there is a pure strategy Bayesian Stackelberg equilibrium $s^* = \left(s_l^*, t_{-l}^*\right)$ of the game Γ_s^b induced by $\mathcal{M}_{\text{SLRF}}$ such that $g\left(s_l^*(\theta_l), (t_{-l}^*(s_l^*(\theta_l)))(\theta_{-l})\right) = f\left(\theta_l, \theta_{-l}\right) \ \forall \left(\theta_l, \theta_{-l}\right) \in \times_{i \in N} \Theta_i$.

By making use of the definition of Bayesian Stackelberg equilibrium, we can say that $s^* = \left(s_l^*, t_{-l}^*\right)$ is a pure strategy Bayesian Stackelberg equilibrium of the game Γ_s^b induced by the mechanism $\mathcal{M}_{\text{SLRF}}$ iff leader Plays a secure strategy and followers play an optimal response.

4.3 Bayesian Stackelberg Incentive Compatibility

1. **Bayesian Stackelberg Incentive Compatibility for the Leader.** An SCF $f(.)$ is said to be Bayesian Stackelberg incentive compatible (BaSIC) for the leader (or truthfully implementable in BS equilibrium for the leader) if the direct revelation mechanism $\mathcal{D}_{\text{SLRF}} = ((\Theta_i)_{i \in N}, f(.))$ has a *BS equilibrium* $s^* = (s^*_l, t^*_{-l})$ in which $s^*_l(\theta_l) = \theta_l, \forall \theta_l \in \Theta_l$. That is, truth telling is a BS equilibrium strategy for the leader in the game induced by $\mathcal{D}_{\text{SLRF}}$.
2. **Bayesian Stackelberg Incentive Compatibility for the Followers.** An SCF $f(.)$ is said to be Bayesian Stackelberg incentive compatible (BaSIC) for the followers (or truthfully implementable in BS equilibrium for the followers) if the direct revelation mechanism $\mathcal{D}_{\text{SLRF}} = ((\Theta_i)_{i \in N}, f(.))$ has a *BS equilibrium* $s^* = (s^*_l, t^*_{-l})$ in which $t^*_{-l}(\theta_l) = ((s^*_j)_{j \in N-l})\ \forall \theta_l \in \Theta_l$, where $s^*_j(\theta_j) = \theta_j\ \forall \theta_j \in \Theta_j\ \forall j \in N - \{l\}$. That is, truth telling is a BS equilibrium strategy for the followers in the game induced by $\mathcal{D}_{\text{SLRF}}$.
3. **Bayesian Stackelberg Incentive Compatibility.** An SCF $f(.)$ is said to be Bayesian Stackelberg incentive compatible (BaSIC) if it is BaSIC for both the leader and the followers.

5 Incentive Compatibility of Reserve Price Procurement Auctions

In this section, we state two key results pertaining to the Bayesian Stackelberg incentive compatibility of the social choice functions for first-price and second-price procurement auctions with reserve prices that were defined earlier in Section 3. Due to paucity of space, we are unable to include the proofs for these results. We urge the interested reader to refer to our recent technical report [3].

Theorem 1. Under the assumptions **A1 - A4**, the social choice function for the first-price procurement auction with reserve prices is BaSIC for the buyer but is not BaSIC for the sellers. The BS equilibrium of the BS game induced by this function among the sellers and the buyer is given by $s^* = (s^*_b, t^*_{-b})$ where

$$s^*_b(\theta_b) = \theta_b\ \forall \theta_b \in \Theta_b = [\,\underline{\theta}\,, \overline{\theta}\,]$$
$$t^*_{-b}(\hat{\theta}_b) = (s^*(.), \ldots, s^*(.))\ \ \forall \hat{\theta}_b \in \Theta_b = [\,\underline{\theta}\,, \overline{\theta}\,]$$

That is, in the BS equilibrium, the buyer announces his true valuation itself as the reserve price, and for any announced reserve price $\hat{\theta}_b$, the sellers bid as suggested by the (symmetric) BN equilibrium strategy profile $(s^*(.), \ldots, s^*(.))$, where

$$s^*(\theta_i) = \begin{cases} \theta_i & : \ \theta_i \in [\hat{\theta}_b,\ \overline{\theta}\,] \\ \theta_i + \frac{1}{[1-F(\theta_i)]^{n-1}} \int_{\theta_i}^{\hat{\theta}_b} [1 - F(x)]^{n-1} dx & : \ \theta_i \in [\,\underline{\theta}\,, \hat{\theta}_b\,] \end{cases}$$

Proof: The proof of this result is based on a systematic analysis of different possible scenarios for bidding by the sellers after the reserve price is made known to the sellers. The analysis leads to quite interesting insights. For details of the proof, refer to [3].

Corollary. If $\hat{\theta_b} = \overline{\theta}$, then F-PAR will be the same as the traditional first-price procurement auction with no reserve price.

Theorem 2. Under the assumptions **A1 - A4**, the social choice function for the second-price procurement auction with reserve prices is BaSIC for the followers but is not BaSIC for the leader. The BS equilibrium of the BS game induced by this function among the sellers and the buyer is given by $s^* = (s_b^*, t_{-b}^*)$ where

1. $s_b^*(.)$ is the solution of the differential equation $F(s_b^*(\theta_b)) = (\theta_b - s_b^*(\theta_b))$ $f(s_b^*(\theta_b))$ with boundary condition $s_b^*(\underline{\theta}) = \underline{\theta}$
2. $t_{-b}^*(\hat{\theta_b}) = (s^*(.), \ldots, s^*(.))\ \forall \hat{\theta_b} \in \Theta_b = [\ \underline{\theta}\ , \overline{\theta}\ ]$ where $s^*(\theta_i) = \theta_i \quad \forall \theta_i \in \Theta_i = [\ \underline{\theta}\ , \overline{\theta}\ ]$

That is, the buyer always announces $s_b^*(\theta_b)$ as the reserve price if his true valuation is θ_b. For any reserve price $\hat{\theta_b}$ announced by the buyer, the sellers always bid their true valuation.

Proof: The proof of this result is also based on a systematic analysis of different possible scenarios for bidding by the sellers after the reserve price is made known to the sellers. The analysis leads to quite interesting insights. For details of the proof, refer to [3].

Corollaries and Insights

1. The optimal reserve price strategy of the buyer in S-PAR when all the sellers draw their types independently from uniform distribution over the set $[0, 1]$ is given by $s_b^*(\theta_b) = \theta_b/2$.
2. Announcing $s_b^*(\theta_b)$ as the reserve price is a better strategy for the buyer in S-PAR than announcing true type as the reserve price.
3. Announcing the true type as the reserve price is a better strategy for the buyer than always fixing $\overline{\theta}$ as the reserve price. Another interpretation of this result is that buyer will be better off in S-PAR if he fixes his true type as reserve price as compared to having no reserve price.
4. If the buyer announces his true type as the reserve price then for any given type θ_b, his expected payoff and the expected revenue paid by him is the same as in F-PAR. The classical Revenue Equivalence theorem [16, 12, 6, 5, 10, 13] can be derived as a special case of the above result.

6 Summary

In this paper we have taken the first steps in extending the classical mechanism design theory to Stackelberg problems in general and SLRF problems in particular. These problems are natural in areas such as distributed computing, grid computing, network routing, ad hoc networks, electronic commerce, and distributed artificial intelligence. To illustrate the approach, we have taken two examples from the domain of electronic commerce - first-price and second-price procurement auctions with reserve prices and investigated the Bayesian Stackelberg incentive compatibility of these two mechanisms. To the best of our knowledge, this is the first attempt in the direction of designing incentive compatible mechanisms for Stackelberg problems. This opens up an avenue for solving many important problems, for example:

– designing fair pricing schemes for ad hoc and wireless networks
– developing fair routing algorithms in wireless ad hoc networks
– designing efficient scheduling policies in the grid computing environment.

References

1. Tamer Başar and Geert Jan Olsder. *Dynamic Noncooperative Game Theory*. Academic Press, London, 1982.
2. Andreu Mas Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.
3. Dinesh Garg and Y. Narahari. A theory of mechanism design for stackelberg problems with selfish agents. Technical report, Electronic Enterprises Laboratory, Department of Computer Science and Automation, Indian Institute of Science, Bangalore, May 2005.
4. Jerry R. Green and Jean-Jacques Laffont. *Incentives in Public Decision Making*. North-Holland Publishing Company, Amsterdam, 1979.
5. Milton Harris and Artur Raviv. Allocation mechanisms and the design of auctions. *Econometrica*, 49(6):1477–99, Nov. 1981.
6. Jr. Holt, Charles A. Uncertainity and the bidding for incentive contracts. *Amer. Econ. Rev.*, 69(4):697–705, Sept. 1979.
7. Jean-Jacques Laffont. *Fundamentals of Public Economics*. The MIT Press, Cambridge, 1988.
8. Preston R. McAfee and John McMillan. Auctions and bidding. *Journal of Economic Literature*, 25(2):699–738, June 1987.
9. Paul R. Milgrom and Robert R. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5):1089–1122, September 1982.
10. Roger B. Myerson. Optimal auction design. *Math. Operations Res.*, 6(1):58–73, Feb. 1981.
11. Roger B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, Massachusetts, 1997.
12. Armando Ortega-Reichert. *Models for Competitive Bidding under Uncertainity*. PhD thesis, Dept. of Operations Research, Stanford University, January 1968.
13. John G. Riley and William F. Samuelson. Optimal auctions. *Amer. Econ. Rev.*, 71(3):381–92, June 1981.
14. T. Roughgarden. Stackelberg scheduling strategies. *SIAM Journal on Computing*, 33(2):332–350, 2004.
15. Von H. Stackelberg. *Marketform und Gleichgewicht*. Springer, Vienna, 1934. An English translation appeared in 1952 entitled "The Theory of Market Economy", Published by Oxford University Press, Oxford.
16. William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, March 1961.
17. Elmar Wolfstetter. *Topics in Microeconomics*. Cambridge University Press, 1999.