Profitable Bayesian implementation

Wu, Haoyang

Wan-Dou-Miao Research Lab

7 January 2019

Profitable Bayesian implementation

Haoyang Wu [*]

Wan-Dou-Miao Research Lab, Shanghai, China.

Abstract

In mechanism design theory, a designer would like to implement a desired social choice function which specifies her favorite outcome for each possible profile of all agents' types. Since agents' types are modelled as their private information, what the designer can do is to construct a mechanism and choose an outcome after observing a specific profile of agents' strategies. Traditionally, the designer has no way to adjust agents' types and hence may be in a dilemma in the sense that even if she is not satisfied with some outcome with low profit, she has to announce it because she must obey the mechanism designed by herself. In this paper, we generalize the mechanism design theory to a case where the designer can adjust the type distribution of agents, and propose a novel notion, *i.e.*, profitable Bayesian implementation. After defining a series of notions such as adjusted types, optimal adjustment cost and profitable Bayesian implementability, we propose that the revelation principle does not hold in this generalized case. Finally, we construct an auction example to show that the designer can obtain an expected profit greater than the maximum profit that she can obtain in the traditional optimal auction, at the same time each agent's *ex ante* expected profit is also increased.

Key words: Mechanism design; Optimal auction; Bayesian Nash implementation.

---

1 Introduction

In the framework of mechanism design theory [1–3], there are one designer and some agents.[1] The designer would like to implement a desired social choice function which specifies her favorite outcome for each possible profile of agents' types. However, agents' types are modelled as their private properties and unknown to the designer. In order to implement the social choice function,

---

[1] The designer is denoted as "She", and the agent is denoted as "He".

the designer constructs a mechanism which specifies each agent's strategy set (*i.e.*, the allowed actions of each agent) and an outcome function (*i.e.*, a rule for how agents' actions get turned into a social choice).

Traditionally, in a mechanism the designer has no way to adjust agents' types and hence may be in a dilemma in the sense that even if some profile of agents' strategies leads to an outcome with low profit, she has to announce it because she must obey the mechanism designed by herself. The designer may improve her situation by holding a charity auction. Engers and McManus [4] proposed that agents' bids in a first-price charity auction are greater than those in a standard (non-charity) auction [5] because of the charitable benefit that winners receive from their own payments. Besides the charity auction, there may exist another way for the designer to escape from the dilemma.

For example, suppose the designer is an auctioneer who sells a good, and each agent is a bidder whose initial valuation to the good (*i.e.*, private type) is low. In order to obtain more profit, the designer announces that she will rent a luxurious hotel to hold this auction. The gorgeousness of the hotel is an open signal to all agents that may induce each agent to adjust his valuation to the good. Without loss of generality, we assume that each agent's valuation and bid both increase concavely with the rent cost spent by the designer, and the designer's utility is a linear function of the winner agent's bid. From the viewpoint of the designer, as long as her marginal utility is greater than her marginal cost, it is worthwhile for her to continue investing on the rent cost. Obviously, the designer will obtain the maximum profit when her marginal utility is equal to her marginal cost. Thus, if agents' types are adjustable, the designer may *actively* escape from the above-mentioned dilemma and yield an outcome better than what would happened without doing so.

In this paper, we generalize the mechanism design theory to a case where the designer can adjust the type distribution of agents, and propose a novel notion, *i.e.*, profitable Bayesian implementation. In Section 2, we define a series of notions such as adjusted types, optimal adjustment cost, profitable Bayesian implementability and so on. Then we propose that the revelation principle does not hold in this generalized case. In Section 3, we construct an example to show that by choosing the optimal adjustment cost, the designer can obtain an expected profit greater than what she can obtain at most in the traditional optimal auction model.

2   Theoretical analysis

Following Section 23.B of MWG's textbook [1], we consider a setting with one designer and I agents, indexed by $i = 1, \cdots, I$. Let X be a set of possible

alternatives.

Assumption 1: Each agent i is assumed to observe a private parameter (*i.e.*, type θ_i) which determines his preference over alternatives in X. Let Θ_i be the set of agent i's all possible types. Let $\Theta = \Theta_1 \times \cdots \times \Theta_I$, $\theta = (\theta_1, \cdots, \theta_I) \in \Theta$.

Definition 1: For any $x \in X$, each *agent i's utility* is denoted as $u_i(x, \theta_i) \in R$, where $\theta_i \in \Theta_i$, and the *designer's utility* is denoted as $u_d(x) \in R$.

Definition 23.B.1 [1]: A *social choice function* (SCF) is a function $f : \Theta \to X$ that, for each possible profile of the agents' types $\theta \in \Theta$, assigns a collective choice $f(\theta) \in X$.

Definition 23.B.3 [1]: A *mechanism* $\Gamma = (S_1, \cdots, S_I, g(\cdot))$ is a collection of I strategy sets $S_1, \cdots, S_I$, and an outcome function $g : S_1 \times \cdots \times S_I \to X$. A *strategy* of each agent i in Γ is a function $s_i(\cdot) : \Theta_i \to S_i$. Let $s(\cdot) = (s_1(\cdot), \cdots, s_I(\cdot))$.

Assumption 2: Assume that in a mechanism, each agent i can play his strategy s_i without any cost. Hence agent i's profit from the outcome x is just his utility $u_i(x, \theta_i)$. [2]

Assumption 3: Assume that in a mechanism, the designer constructs an outcome function and announces a cost $c \geq 0$ which she will spend to perform the outcome function. The cost c is observable to all agents and acts as an open signal. The outcome function is denoted as $g^c(\cdot) : S_1 \times \cdots \times S_I \to X$, and the mechanism is denoted as $\Gamma^c = (S_1, \cdots, S_I, g^c(\cdot))$. After learning the cost c, each agent i adjusts his private type from the initial value $\theta_i^0 \in \Theta_i$ to a new value $\theta_i^c \in \Theta_i$, [3] and then plays strategy $s_i(\theta_i^c)$. At last, the designer announces $g^c(s_1(\theta_1^c), \cdots, s_I(\theta_I^c))$ as the outcome. The cost c is also denoted as *adjustment cost*. Thus, although the designer does not know agents' private types exactly, she can adjust the type distribution of agents.

Definition 2: For any adjustment cost $c \geq 0$ and each agent i's initial type

---

[2] For example, suppose that each agent is a bidder in an auction, then each agent can be considered to submit his bid to the auctioneer without any cost.

[3] In Ref [5] (Page 60, Line 12), Myerson proposed that "if there are quality uncertainties, then bidder i might tend to revise his valuation of the object after learning about other bidders' value estimates." Similarly, here it is reasonable to assume that each agent i can adjust his private type after observing the cost signal sent by the designer. Section 3 gives an example, where the designer spends some cost to rent a hotel to hold an auction. The gorgeousness of the hotel is just the signal that the designer sends to agents in order to show how precious the sold good is, although the designer may sell a poor good but deliberately rent a luxurious hotel to deceive agents. After observing the signal, each agent adjusts his private type (*i.e.*, valuation to the good).

$\theta_i^0 \in \Theta_i$, by assumption 3 each agent i's preference over the alternatives in X is determined by his *adjusted type* $\theta_i^c \in \Theta_i$. For each $i = 1, \cdots, I$, let

$$\begin{aligned}
\theta^0 &= (\theta_1^0, \cdots, \theta_I^0) \in \Theta, \\
\theta_{-i}^0 &= (\theta_1^0, \cdots, \theta_{i-1}^0, \theta_{i+1}^0, \cdots, \theta_I^0), \\
\theta^c &= (\theta_1^c, \cdots, \theta_I^c) \in \Theta, \\
\theta_{-i}^c &= (\theta_1^c, \cdots, \theta_{i-1}^c, \theta_{i+1}^c, \cdots, \theta_I^c).
\end{aligned}$$

A *type adjustment function* is denoted as $\mu(\theta, c) : \Theta \times R^+ \to \Theta$, in which $\mu(\theta, 0) = \theta$ for any $\theta \in \Theta$, *i.e.* zero adjustment cost means no type adjustment. Let $\theta^c = \mu(\theta^0, c)$. Let $\phi^0(\theta^0) = (\phi_1^0(\theta_1^0), \cdots, \phi_I^0(\theta_I^0))$ be the probability density function of initial type profile $\theta^0 \in \Theta$, and $\phi^c(\theta^c) = (\phi_1^c(\theta_1^c), \cdots, \phi_I^c(\theta_I^c))$ be the probability density function of adjusted type profile $\theta^c \in \Theta$. For each $i = 1, \cdots, I$, let

$$\begin{aligned}
\phi_{-i}^0(\theta_{-i}^0) &= (\phi_1^0(\theta_1^0), \cdots, \phi_{i-1}^0(\theta_{i-1}^0), \phi_{i+1}^0(\theta_{i+1}^0), \cdots, \phi_I^0(\theta_I^0)), \\
\phi_{-i}^c(\theta_{-i}^c) &= (\phi_1^c(\theta_1^c), \cdots, \phi_{i-1}^c(\theta_{i-1}^c), \phi_{i+1}^c(\theta_{i+1}^c), \cdots, \phi_I^c(\theta_I^c)).
\end{aligned}$$

Assumption 4: For any $\theta \in \Theta$ and adjustment cost $c \geq 0$, the designer is assumed to know type adjustment function $\mu(\theta, c)$, the initial type distribution $\phi^0(\cdot)$ and the adjusted type distribution $\phi^c(\cdot)$.

Definition 3: For any $c \geq 0$, given an SCF f and $\phi^0(\cdot)$, the *designer's expected utility* is denoted as $\bar{u}_d(c) = E_{\theta^c} u_d(f(\theta^c))$, and her *initial expected utility* is denoted as $\bar{u}_d(0) = E_{\theta^0} u_d(f(\theta^0))$.

Definition 4: For any $c \geq 0$, given an SCF f and $\phi^0(\cdot)$, the *designer's expected profit* is denoted as $\bar{p}_d(c) = \bar{u}_d(c) - c$, and her *initial expected profit* is denoted as $\bar{p}_d(0) = \bar{u}_d(0)$.

Assumption 5: $\bar{u}_d(c)$ is assumed to be a concave function that satisfies the following inequalities,

$$\frac{\partial \bar{u}_d(c)}{\partial c} > 0, \quad \frac{\partial^2 \bar{u}_d(c)}{\partial c^2} < 0, \quad \text{for any } c \geq 0. \, [4]$$

Proposition 1: If there exists an adjustment cost $c^* \geq 0$ such that

$$\left. \frac{\partial \bar{u}_d(c)}{\partial c} \right|_{c=c^*} = 1, \quad i.e. \quad \left. \frac{\partial \bar{p}_d(c)}{\partial c} \right|_{c=c^*} = 0,$$

---

[4] See the example given in Section 3. When each agent i's adjusted type is a square root function of the designer's cost as specified by Eq (5) and the social choice function is specified by Eq (6), then the inequalities in Assumption 5 holds.

then the designer will obtain the maximum expected profit $\bar{p}_d(c^*)$ at $c = c^*$. c^* is denoted as the *optimal adjustment cost*. By Assumption 5 there holds

$$\left. \frac{\partial \bar{u}_d(c)}{\partial c} \right|_{c=0} \geq 1, \quad i.e. \quad \left. \frac{\partial \bar{p}_d(c)}{\partial c} \right|_{c=0} \geq 0.$$

Definition 23.D.1 [1]: The strategy profile $s^*(\cdot) = (s_1^*(\cdot), \cdots, s_I^*(\cdot))$ is a *Bayesian Nash equilibrium* of mechanism $\Gamma = (S_1, \cdots, S_I, g(\cdot))$ if, for all agent i and all $\theta_i \in \Theta_i$,

$$E_{\theta_{-i}}[u_i(g(s_i^*(\theta_i), s_{-i}^*(\theta_{-i})), \theta_i)|\theta_i] \geq E_{\theta_{-i}}[u_i(g(\hat{s}_i, s_{-i}^*(\theta_{-i})), \theta_i)|\theta_i], \quad (1)$$

for all $\hat{s}_i \in S_i$.

Definition 23.D.2 [1]: The mechanism $\Gamma = (S_1, \cdots, S_I, g(\cdot))$ *implements the social choice function $f(\cdot)$ in Bayesian Nash equilibrium* if there is a Bayesian Nash equilibrium of Γ, $s^*(\cdot) = (s_1^*(\cdot), \cdots, s_I^*(\cdot))$, such that $g(s^*(\theta)) = f(\theta)$ for all $\theta \in \Theta$.

Definition 5: Given an SCF f and $\phi^0(\cdot)$, f is *profitable Bayesian implementable* if the following conditions are satisfied:
1) The optimal adjustment cost $c^* > 0$, which means that the distribution of agents' private types is adjusted from $\phi^0(\cdot)$ to $\phi^{c^*}(\cdot)$.
2) There exist a mechanism $\Gamma^{c^*} = (S_1, \cdots, S_I, g^{c^*}(\cdot))$ that implements f in Bayesian Nash equilibrium. That is, there exists a strategy profile $s^*(\cdot) = (s_1^*(\cdot), \cdots, s_I^*(\cdot))$ such that:
(i) For all agent i and all $\theta_i^{c^*} \in \Theta_i$,

$$E_{\theta_{-i}^{c^*}}[u_i(g^{c^*}(s_i^*(\theta_i^{c^*}), s_{-i}^*(\theta_{-i}^{c^*})), \theta_i^{c^*})|\theta_i^{c^*}] \geq E_{\theta_{-i}^{c^*}}[u_i(g^{c^*}(\hat{s}_i, s_{-i}^*(\theta_{-i}^{c^*})), \theta_i^{c^*})|\theta_i^{c^*}] \quad (2)$$

for all $\hat{s}_i \in S_i$. [5]
(ii) $g^{c^*}(s^*(\theta)) = f(\theta)$ for all $\theta \in \Theta$.

Proposition 2: Given an SCF f and $\phi^0(\cdot)$, if f is profitable Bayesian implementable, then the designer's expected profit at the optimal adjustment cost is greater than her initial expected profit.
Proof: Given that f is profitable Bayesian implementable, then the optimal adjustment cost $c^* > 0$. By Proposition 1, $\bar{p}_d(c^*) > \bar{p}_d(0)$. □

Definition 23.D.3 [1]: A social choice function $f(\cdot)$ is *truthfully implementable in Bayesian Nash equilibrium* (or *Bayesian incentive compatible*) if $s_i^*(\theta_i) = \theta_i$

---

[5] Note that in formula (2), the probability density function of type profile $\theta_{-i}^{c^*} = (\theta_1^{c^*}, \cdots, \theta_{i-1}^{c^*}, \theta_{i+1}^{c^*}, \cdots, \theta_I^{c^*})$ is $\phi_{-i}^{c^*}(\cdot)$. As a comparison, in the notion of Bayesian Nash equilibrium (see Definition 23.D.1), there is no type adjustment. Thus, the probability density function of type profile $\theta_{-i} = (\theta_1, \cdots, \theta_{i-1}, \theta_{i+1}, \cdots, \theta_I)$ in formula (1) is just $\phi_{-i}^0(\cdot)$.

for all $\theta_i \in \Theta_i$ $(i = 1, \cdots, I)$ is a Bayesian Nash equilibrium of the direct mechanism $\Gamma = (S_1, \cdots, S_I, g(\cdot))$, in which $S_i = \Theta_i$, $g = f$. That is, if for all $i = 1, \cdots, I$ and all $\theta_i \in \Theta_i$,

$$E_{\theta_{-i}}[u_i(f(\theta_i, \theta_{-i}), \theta_i)|\theta_i] \geq E_{\theta_{-i}}[u_i(f(\hat{\theta}_i, \theta_{-i}), \theta_i)|\theta_i], \qquad (3)$$

for all $\hat{\theta}_i \in \Theta_i$.

Proposition 23.D.1 [1]: (*The Revelation Principle for Bayesian Nash Equilibrium*) Suppose that there exists a mechanism $\Gamma = (S_1, \cdots, S_I, g(\cdot))$ that implements the social choice function $f(\cdot)$ in Bayesian Nash equilibrium. Then $f(\cdot)$ is truthfully implementable in Bayesian Nash equilibrium.

Proposition 3: Given an SCF f and $\phi^0(\cdot)$, if f is profitable Bayesian implementable, then it cannot be inferred that f is truthfully implementable in Bayesian Nash equilibrium. That is, the revelation principle does not hold in this generalized case.

Proof: Given that f is profitable Bayesian implementable, then the optimal adjustment cost $c^* > 0$, and there exist a mechanism $\Gamma^{c^*} = (S_1, \cdots, S_I, g^{c^*}(\cdot))$ that implements f in Bayesian Nash equilibrium. Thus, formula (2) is related to the type distribution $\phi^{c^*}(\cdot)$ (see Footnote 5).
As a comparison, in the notion of Bayesian incentive compatibility, there is no type adjustment in the direct mechanism by Definition 23.D.3 . Thus, formula (3) is related to the type distribution $\phi^0(\cdot)$.
Since $\theta^{c^*} = \mu(\theta^0, c^*)$ and $\phi^0(\cdot)$ is not equal to $\phi^{c^*}(\cdot)$, formula (3) cannot be inferred from formula (2). Therefore, given that f is profitable Bayesian implementable, it cannot be inferred that f is truthfully implementable in Bayesian Nash equilibrium. Consequently, the revelation principle does not hold in this generalized case. $\square$

Proposition 4: If the designer's expected utility $\bar{u}_d(c)$ and expected profit $\bar{p}_d(c)$ satisfy the following condition,

$$\left.\frac{\partial \bar{u}_d(c)}{\partial c}\right|_{c=0} < 1, \quad i.e. \quad \left.\frac{\partial \bar{p}_d(c)}{\partial c}\right|_{c=0} < 0, \qquad (4)$$

then the designer will obtain the maximum expected profit at $c = 0$. Put differently, in this case the designer cannot obtain any more profit by spending any cost to adjust agents' types.

3   Example

Following the auction model in MWG's book (Page 863, [1]), suppose that there are one designer and two agents. Let the designer be an auctioneer who wants to sell a good, and each agent be a bidder whose valuation to the good is $\theta_i \geq 0$, $i.e.$, $\Theta_i = R^+$. We consider a first-price-sealed-bid auction setting: Each agent i is allowed to submit a sealed bid $b_i \geq 0$. The bids are then opened, and the agent with the higher bid gets the good, and must pay money equal to his bid to the auctioneer. The auction is not an on-line format and the designer holds the auction in a hotel.

Suppose that:
1) Each agent i's initial valuation ($i.e.$, his initial type) θ_i^0 is drawn independently from the uniform distribution on $[0,1]$. The distribution is known by the designer but the exact value of each θ_i^0 is agent i's private information.
2) The designer spends cost $c \geq 0$ to rent a hotel.
3) The gorgeousness of hotel is characterized by the rent cost. Each agent i's valuation to the good is adjusted by the gorgeousness of the hotel: The more expensive the rent cost is, the greater the bidder's private valuation to the good is.
4) Let $\beta > 0$ be a coefficient, each agent i's valuation to the good ($i.e.$, his adjusted type θ_i^c) is a square root function of the rent cost c,

$$\theta_i^c = (1 + \beta\sqrt{c})\theta_i^0. \tag{5}$$

Thus,

$$\frac{\partial \theta_i^c}{\partial c} = \frac{\beta\theta_i^0}{2\sqrt{c}}, \quad \frac{\partial^2 \theta_i^c}{\partial c^2} = -\frac{\beta\theta_i^0}{4}c^{-3/2}.$$

That is, for any $c \geq 0$, the following formulas hold:

$$\frac{\partial \theta_i^c}{\partial c}\bigg|_{c=0} = +\infty, \ \frac{\partial \theta_i^c}{\partial c} > 0, \ \frac{\partial^2 \theta_i^c}{\partial c^2} < 0.$$

Let $\theta = (\theta_1, \theta_2)$, consider the social choice function

$$f(\theta) = (y_1(\theta), y_2(\theta), y_d(\theta), t_1(\theta), t_2(\theta), t_d(\theta)), \tag{6}$$

in which

$$\begin{aligned}
y_1(\theta) &= 1, \quad \text{if } \theta_1 \geq \theta_2; \quad = 0 \text{ if } \theta_1 < \theta_2 \\
y_2(\theta) &= 1, \quad \text{if } \theta_1 < \theta_2; \quad = 0 \text{ if } \theta_1 \geq \theta_2 \\
y_d(\theta) &= 0, \quad \text{for all } \theta \in \Theta \\
t_1(\theta) &= -\theta_1 y_1(\theta)/2 \\
t_2(\theta) &= -\theta_2 y_2(\theta)/2 \\
t_d(\theta) &= [\theta_1 y_1(\theta) + \theta_2 y_2(\theta)]/2.
\end{aligned}$$

The subscript d stands for the designer, and the subscript $1, 2$ stands for the agent 1 and agent 2 respectively. $y_i = 1$ means that agent $i = 1, 2$ gets the good. t_i denotes agent i's payment to the designer. t_d denotes the sum of two agents' payment to the designer. Obviously, for any $c \geq 0$, $y_1(\theta^c) = y_1(\theta^0)$, $y_2(\theta^c) = y_2(\theta^0)$.

Now we investigate whether this social choice function is Bayesian implementable. We will look for an equilibrium in which each agent i's strategy $b_i(\cdot)$ takes the form $b_i(\theta_i^c) = \alpha_i \theta_i^c = \alpha_i(1 + \beta\sqrt{c})\theta_i^0$ for $\alpha_i \in [0, 1]$.

Suppose that agent 2's strategy has this form, and consider agent 1's problem. For each possible θ_1^c, agent 1 wants to solve the following problem:

$$\max_{b_1 \geq 0} (\theta_1^c - b_1)\mathrm{Prob}(b_2(\theta_2^c) \leq b_1).$$

Because agent 2's highest possible bid is $\alpha_2(1 + \beta\sqrt{c})$ when $\theta_2^0 = 1$, it is evident that agent 1's bid b_1 should never more than $\alpha_2(1 + \beta\sqrt{c})$. Note that θ_2^0 is uniformly distributed on $[0, 1]$, and $b_2(\theta_2^c) = \alpha_2(1 + \beta\sqrt{c})\theta_2^0 \leq b_1$ if and only if $\theta_2^0 \leq b_1/[\alpha_2(1 + \beta\sqrt{c})]$. Hence we can write agent 1's problem as:

$$\max_{0 \leq b_1 \leq \alpha_2(1+\beta\sqrt{c})} \frac{(\theta_1^c - b_1)b_1}{\alpha_2(1 + \beta\sqrt{c})}$$

The solution to this problem is

$$b_1^*(\theta_1^c) = \begin{cases} \theta_1^c/2, & \text{if } \theta_1^0/2 \leq \alpha_2 \\ \alpha_2(1 + \beta\sqrt{c}), & \text{if } \theta_1^0/2 > \alpha_2 \end{cases}.$$

Similarly,

$$b_2^*(\theta_2^c) = \begin{cases} \theta_2^c/2, & \text{if } \theta_2^0/2 \leq \alpha_1 \\ \alpha_1(1 + \beta\sqrt{c}), & \text{if } \theta_2^0/2 > \alpha_1 \end{cases}.$$

Letting $\alpha_1 = \alpha_2 = 1/2$, we see that the strategies $b_i^*(\theta_i^c) = \theta_i^c/2$ for $i = 1, 2$ constitute a Bayesian Nash equilibrium for this mechanism. Thus, there is a Bayesian Nash equilibrium of this first-price-sealed-bid auction mechanism that indirectly yields the outcomes specified by the social choice function $f(\theta)$, and hence $f(\theta)$ is Bayesian Nash implementable.

Let us consider the designer's expected profit:

$$\bar{p}_d(c) = (1 + \beta\sqrt{c})E[\theta_1^0 y_1(\theta^0) + \theta_2^0 y_2(\theta^0)]/2 - c.$$

The designer's problem is to choose an optimal adjustment cost $c \geq 0$ to maximize her expected profit, *i.e.*,

$$\max_{c \geq 0} (1 + \beta\sqrt{c})E[\theta_1^0 y_1(\theta^0) + \theta_2^0 y_2(\theta^0)]/2 - c$$

According to Appendix A, the designer's initial expected profit is $\bar{p}_d(0) = E[\theta_1^0 y_1(\theta^0) + \theta_2^0 y_2(\theta^0)]/2 = 1/3$. Thus, the designer's problem is reformulated as:

$$\max_{c \geq 0}(1 + \beta\sqrt{c})/3 - c$$

It can be easily derived that the optimal adjustment cost $c^* = \beta^2/36$. By Definition 5, $f(\theta)$ is profitable Bayesian implementable. The maximum expected profit of the designer is:

$$\bar{p}_d(c^*) = (1 + \beta\sqrt{c^*})/3 - c^* = \frac{1}{3}(1 + \frac{\beta^2}{12}).$$

Obviously, when $\beta > \sqrt{3}$, there exists $\bar{p}_d(c^*) > 5/12$. Note that the designer's maximum expected profit in the traditional optimal auction with two bidders is $5/12$ (see Page 23, the ninth line from the bottom, Ref [6]). Therefore, if $\beta > \sqrt{3}$, then by choosing the optimal adjustment cost $c^* = \beta^2/36$, *the designer can obtain an expected profit greater than the maximum expected profit given by the traditional optimal auction.*

Now we consider the *ex ante* expected profit of each agent i when agents' types are adjustable and the designer chooses the optimal adjustment cost. Since two agents are symmetric, each of them has the same probability $1/2$ to win the good, which leads to the expected profit as follows:

$$E[\theta_i^{c^*} - b_i^*(\theta_i^{c^*})] = E[\theta_i^{c^*}/2] = (1 + \beta\sqrt{c^*})E[\theta_i^0]/2 = (1 + \frac{\beta^2}{6})/4.$$

Therefore, the *ex ante* expected profit of each agent when agents' types are adjustable is $1/8 + \beta^2/48$.

As a comparison, we consider the *ex ante* expected profit of each agent in the traditional optimal auction model. By Ref [6] (Page 22), the *ex ante* expected payment of a bidder is

$$r(1 - F(r))G(r) + \int_r^\omega y(1 - F(y))g(y)dy,$$

where $r > 0$ is the reserve price, $[r, \omega]$ is the interval of each agent i's valuation X_i. Each X_i is independently and identically distributed according to the increasing distribution function F. Fix a bidder, G denotes the distribution function of the highest valuation among the rest remaining bidders. For the case of two agents with valuation range $[r, 1]$ and uniform distribution, there exist

$$F(r) = r, \ G(r) = r, \ \omega = 1, \ F(y) = y, \ g(y) = 1, \text{for any } y \in [r, 1].$$

By Ref [6] (Page 23), given that each of two agent's valuation to the good is uniformly distributed on interval $[0, 1]$, the optimal reserve price $r^* = 1/2$.

Therefore, the *ex ante* expected payment of each agent in the optimal auction with reserve price $r^* = 1/2$ is

$$r^*(1-r^*)r^* + \int_{r^*}^1 y(1-y)dy$$
$$= \frac{1}{8} + \int_{\frac{1}{2}}^1 y(1-y)dy = \frac{5}{24}.$$

The *ex ante* expected profit of each agent in the traditional optimal auction is his expected valuation minus his *ex ante* expected payment, *i.e.*,

$$\int_{\frac{1}{2}}^1 ydy - \frac{5}{24} = \frac{3}{8} - \frac{5}{24} = \frac{1}{6}.$$

It can be seen that if $\beta > \sqrt{2}$, then $1/8 + \beta^2/48 > 1/6$.

To sum up, when agents' types are adjustable and the designer chooses the optimal adjustment cost, if $\beta > \sqrt{3}$, then the expected profit of the designer and the *ex ante* expected profit of each agent are all greater than their counterparts in the traditional optimal auction with reserve price.

4   Conclusions

Traditionally, agents' types are considered as private and endogenous values, which means that the designer has no way to know and adjust each agent's type. Thus, although the designer constructs a mechanism in order to implement her favorite social choice function, she behaves like *a passive observer in a dilemma* after receiving a profile of agents' strategies: *i.e.*, she must obey the mechanism and announce the outcome specified by the outcome function, no matter whether she is satisfied with the outcome or not.

In this paper, we generalize the mechanism design theory to a case where the designer can adjust the type distribution of agents, and propose a novel notion, *i.e.*, profitable Bayesian implementation. In the generalized case, the designer behaves like *an active modulator* and escape from the dilemma by choosing the adjustment cost and adjusting agents' type distribution. There are two novelties related to profitable Bayesian implementation:
1) For a profitable Bayesian implementable social choice function, it cannot be inferred that it is truthfully implementable in Bayesian Nash equilibrium. That is, the revelation principle does not hold in this generalized case.
2) By choosing the optimal adjustment cost, the designer can obtain an expected profit greater than the maximum expected profit yielded by the traditional optimal auction. At the same time, each agent's *ex ante* expected profit is also increased.

Appendix

As specified in Section 3, θ_1^0 and θ_2^0 are drawn independently from the uniform distribution on $[0,1]$. Let Z be a random variable defined as $Z = \theta_1^0 y_1(\theta^0) + \theta_2^0 y_2(\theta^0)$.

$$f_{\theta_1^0}(z) = \begin{cases} 0, & z < 0 \\ 1, & z \in [0,1] \\ 0, & z > 1 \end{cases}.$$

$$F_{\theta_1^0}(z) = Prob\{\theta_1^0 \leq z\} = \begin{cases} 0, & z < 0 \\ z, & z \in [0,1] \\ 1, & z > 1 \end{cases}.$$

$$F_Z(z) = [F_{\theta_1^0}(z)]^2 = \begin{cases} 0, & z < 0 \\ z^2, & z \in [0,1] \\ 1, & z > 1 \end{cases}.$$

Therefore,

$$f_Z(z) = \begin{cases} 0, & z < 0 \\ 2z, & z \in [0,1] \\ 0, & z > 1 \end{cases}.$$

As a result,

$$E(Z) = \int_0^1 z \cdot 2z dz = \int_0^1 2z^2 dz = 2/3.$$

Therefore, $E[\theta_1^0 y_1(\theta^0) + \theta_2^0 y_2(\theta^0)]/2 = 1/3$. According to Eq (6), the designer's initial expected profit and utility are $\bar{p}_d(0) = \bar{u}_d(0) = 1/3$.

Acknowledgments

References

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

[2] Y. Narahari et al, Game Theoretic Problems in Network Economics and Mechanism Design Solutions, Springer, 2009.

[3] R. Serrano, The Theory of Implementation of Social Choice Function, SIAM Review, vol.46, No.3, 377-414, 2004.

[4] M. Engers and B. McManus, Charity Auctions, International Economic Review, vol.48, No.3, 953-994, 2007.

[5] R. Myerson, Optimal Auction Design, Mathematics of Operations Research, vol.6, No.1, 58-73, 1981.

[6] V. Krishna, Auction Theory (Second Edition), Academic Press, 2010.