Optimal Auctions for Partially Rational Bidders*

Zihe Wang and **Pingzhong Tang**
Institute for Interdisciplinary Information Sciences,
Tsinghua University, Beijing, China
wangzihe11@mails.tsinghua.edu.cn, kenshin@tsinghua.edu.cn

Abstract

We investigate the problem of revenue optimal mechanism design [Myerson, 1981] under the context of the *partial rationality model*, where buyers randomize between two modes: rational and irrational. When a buyer is irrational (can be thought of as lazy), he acts according to certain fixed strategies, such as bidding his true valuation. The seller cannot observe the buyer's valuation, or his rationality mode, but treat them as random variables from known distributions. The seller's goal is to design a single-shot auction that maximizes her expected revenue.

A minor generalization as it may seem, our findings are in sharp contrast to Myerson's theory on the standard rational bidder case. In particular, we show that, even for the simplest setting with one buyer, direct value revelation loses generality. However, we do show that, in terms of revenue, the optimal value-revelation and type-revelation mechanisms are equivalent. In addition, the posted-price mechanism is no longer optimal. In fact, the more complicated the mechanism, the higher the revenue. For the case where there are multiple bidders with IID uniform valuations, we show that when the irrational buyers are truthful, first price auction yields more revenue than second price auction.

1 Introduction

Rationality assumption, which assumes that agents are utility maximizers, is perhaps the most important assumption in game theory [Von Neumann and Morgenstern, 1947]. However, this assumption has been questioned repeatedly in many practical scenarios. Violation of this assumption invalidates fundamental theories and raises research challenges.

Consider, for example, the literature of bounded rationality [Rubinstein, 1986; Osborne and Rubinstein, 1994], where agents' strategies are computationally constrained. Rubinstein [1986] shows that, in repeated games, new Nash equilbria can emerge by restricting agents' strategies to *finite automata*. A more realistic example is the drainage tract auctions [Milgrom, 2004, Chapter 5.3] where buyers compete for tracts to extract oil. Evidences show that, in such auctions, bidders generally do not know their utility functions, and sometimes have to learn the functions via certain costly research [Larson and Sandholm, 2000; 2004]. The same phenomenon has been studied under the context of revenue optimal auction design [Celis *et al.*, 2012; Tang and Sandholm, 2012]. Wright and Leyton-Brown [2010; 2012; 2014] also provide evidences that, rather than being rational, agents' behaviors follow certain fixed partterns.

Along this line, in this paper, we investigate how the rationality assumption affects the theory of optimal auction design [Myerson, 1981], one of the most prevailing theories in economic theory. In our setting, we have a seller who has a single item for sale. Each buyer is partially rational in the sense that he is rational with some probability p and irrational with probability $1-p$. When he is irrational, he adopts some fixed strategy (e.g., he truthfully reports his valuation, regardless of the mechanisms). Our goal is to design a revenue optimal mechanism for this type of buyers.

The rationality model above is motivated from a number of real-world observations. Consider an advertiser that participates in an online advertisement auction, at any given time, only part of the agents are active and change their bids according to current profile, while the remainder simply repeat their previous bids. Such repetition can be regarded as a type of partially irrational behavior in our model. Another example is agents tend to be truthful when the underlying mechanism is not. Examples are when an agent purchases a product (say, mobile phone, jewelry) under the recommendation of a shop assistant, who strategically displays a list of items after knowing the agent's type (budget, etc.). The items recommended may not be the most suitable for the agent, but the ones with the largest profit for the shop. Revealing your real type (as most people will do) may be sub-optimal in this case and can be regarded as an irrational behavior in our model.

In online ad-auctions [Pai and Vohra, 2014], advertisers are known to sometimes bid their entire budgets (valuation) to ensure an ad slot (It will be difficult for the bidder in charge to explain to their superior why he fails to obtain any slot while

*This work was supported by the National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301, the Natural Science Foundation of China Grant 61033001, 61361136003, 61303077, a Tsinghua Initiative Scientific Research Grant and a China Youth 1000-talent program.

still has some budget left). Lee and Malmendier[2011] conduct experiments on sets of bidders that are simultaneously available to two auctions: a posted price as well as an ascending auction. They show that, 42 percent of the time, the ascending auction ends up with a price higher than the posted price (which is also available to them)!

In this paper, we investigate the optimal mechanism design theory under this rationality model. Our conclusion is that, Myerson's theory [1981] does not apply to this partial rationality model, even for the simplest case with one buyer. In particular, our findings are as follows.

- It loses generality to consider only direct value revelation. In the 1-buyer setting, when restricted to value revelation, the optimal mechanism is no longer a posted price mechanism, or even truthful.
- In fact, the more choices (menu items [Hart and Nisan, 2013; Wang and Tang, 2014], to be rigorously defined later) the seller provides to the buyer, the higher her revenue. We show, however, that the optimal revenue can be approximated via simple mechanisms (k-piecewise-linear mechanisms).
- The optimal mechanism that can solicit both rationality modes and bids yields the same revenue as the optimal mechanism that can only solicit bids.
- For the general case where there are several buyers with IID uniform valuations and irrational behavior under consideration is to truthfully report valuation, we compute the Bayes Nash equilibrium for first price auction and show that it yields more revenue than the second price auction.

Our techniques for obtaining these results make use of novel geometrical arguments in optimal mechanism design [Hartline, 2013; Manelli and Vincent, 2007].

2 Preliminaries

In our basic setting, there is a single bidder that is randomized between rationality and irrationality. We then extend to the case where there are a set of such bidders.

2.1 Partially rational bidders

In the basic setting, a seller has one indivisible good for sale. Her valuation towards the item is normalized to 0. There is a single bidder that is interested in the good. His valuation (aka. type) x of the item is drawn from a distribution F with density function $f(x)$ positive everywhere on $[0, 1]$. With probability p, the bidder is rational, in the sense that he is an expected utility maximizer, in the game designed by the seller; with probability $1 - p$, the bidder is irrational, in the sense that he will adopt a fixed strategy $c(x)$. We assume that $c(x)$ is differentiable and an *onto* on $[0, 1]$. We sometimes call the bidder in his rational mode *the rational*. Same for the irrational case.

The seller's objective is to maximize her expected revenue, over the randomness of the bidder's valuations and rationality mode. She does so by designing a *mechanism*. We will be formal with the definition of a mechanism shortly.

2.2 Type revelation versus value revelation

In principle, we consider here a mechanism design problem with *two-dimensional* types, where the first dimension is one's valuation and the second is one's rationality mode. If it were feasible for the bidders to reveal their complete types, by revelation principle, we can first solicit both valuation and rationality mode from the bidders and then price and allocate the item conditioning on agents' modes, subject to a two-dimensional incentive compatibility constraint (e.g., an agent in its rational mode can misreport either her valuation or her rationality mode).

However, in reality, it is impractical to implement full type revelation. As stated in the introduction, rationality mode may emerge implicitly in agents' behaviors (such as laziness) and it is sometimes hard or uncomfortable to evaluate even by the agent itself. To our best knowledge, ask the bidders for their rationality modes has never been seen in any real-world auction design scenario. In this paper, we follow the tradition of auction design and focus on the mechanisms that can only solicit bids from bidders.

In our model, it loses generality to restrict attention to direct value-revelation mechanisms where it is incentive compatible for rational agents to reveal their valuations. The standard revelation principle claims that, for any objective the seller wants to implement, it is without loss of generality to restrict the seller's design space to one-shot truthful type-revelation mechanisms. If restricted to value revelation settings, the principle no longer holds. Recall its proof [Shoham and Leyton-Brown, 2009, Page 266] shown in Fig 1.

For any indirect mechanism M, assume player's strategy is s. The proof constructs a direct mechanism $M(s(\cdot))$ that first applies the strategy s and then calls the indirect mechanism as a black box. The new mechanism $M(s(\cdot))$ satisfies the key property that a manipulation in the direct mechanism will cause a manipulation in the indirect mechanism.

The same proof does not go through in the value-revelation setting, shown in Fig 2. When we construct new mechanism in order to make the rational bidder truthful, i.e. creating $M(\rho(\cdot))$. the payment from the irrational mode changes. The key property does not hold anymore: the indirect mechanism cannot replicate a manipulation of the direct mechanism because he is not in control of his strategy when he is irrational!

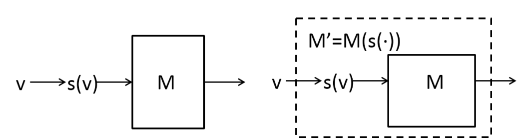

Figure 1: Standard proof

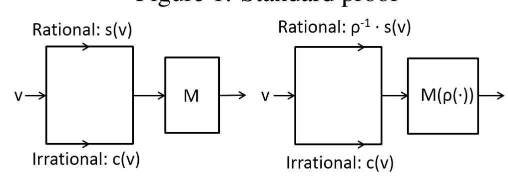

Figure 2: The standard proof fails

Observation 1 *Restricting attention to truthful value-revelation mechanisms (i.e., truthful auctions in the standard sense) loses generality.*

The observation above is further confirmed by Theorem 3, where we demonstrate that the rational bidder misreports his valuation in the optimal mechanism.

Finally, it is important to note that the optimal mechanism that can solicit both rationality modes and bids yields the same revenue as the optimal mechanism that can only solicit bids: the difference only lies in the representation.

2.3 Mechanisms

The aforementioned sub-optimality of truthful value-revelation complicates revenue optimal mechanism design. As mentioned, Theorem 3 guarantees that the revenue in the optimal mechanism that allows type revelation, is the same as that in the optimal mechanism which only allows value revelation. Thus, we can focus on the one-shot value revelation mechanism: the seller solicits a bid in $[0, 1]$ from the bidder, and then allocates the item with some probability q at price t. The bidder's (quasi-linear) utility is given by $qx - t$. From now on, we use "optimal mechanism" as a shorthand for "optimal one-shot mechanism" defined above.

It is useful to think of such a mechanism as a set of *menu items*[Hart and Nisan, 2013] $\{(b, q, t), ...\}$, where a menu item (b, q, t) means the item will be sold with probability q at price t, if the bidder bids b. Therefore, with probability p, the bidder chooses a menu item that maximizes his utility; with probability $1 - p$, he chooses a menu item according to a fixed strategy.

We impose an individual rationality assumption on mechanisms, so that the bidder is not worse off by participation.

Definition 1 *A mechanism is* ex-post *individually rational if a bidder gets non-negative utility at any realization of his type and rationality mode.*

The following lemma states the equivalence between an irrational strategy $c(x)$ given distribution F and a truthful strategy given some different distribution G. It serves as a technical tool for us to characterize the optimal mechanism.

Lemma 1 *For any irrational strategy $c(x)$, there is an equivalent (in turns of bidder's utility) irrational strategy that truthfully reports one's valuation from a different distribution.*

Proof: In mechanism M, let $s(x)$ denote the rational bidder's strategy in Bayes Nash equilibria when irrational behavior is $c(x)$. Let G denote the distribution of irrational bids $c(x), x \sim F$.

Consider a scenario where there are two bidders, with probability p the rational one participates, in other time the irrational one participates, the rational's valuation is distributed according F; the irrational's valuation is distributed according G and the irrational behavior is to truthfully report valuation. In this scenario we use same mechanism M, let the rational bidder's strategy be $s(x)$, it is easy to check that this still forms a Bayes Nash equilibrium for the rational, since bids distribution does not change, the winning probability for same bid does not change. The Rational bidder's utility is still maximized by strategy s. ■

From the seller's view, she cares about the irrational *bids* distribution rather than the *strategies* used by the bidder. When the distribution is fixed, the payment from buyer in the irrational mode is fixed. According to Lemma 1, we can assume from now on the irrational behavior is to truthfully report the value: the value distribution in the rational case is F, and the value distribution in the irrational case is G.

3 Representing a mechanism by the bidder's utility function

In this part, we show that constructing an optimal mechanism is equivalent to designing a convex monotone utility function. For any optimal mechanism, it can be constructed by the rational bidder's utility function. The idea is we first fix the utility function of the rational, then construct the mechanism which brings seller the largest profit.

At the beginning, we prove the rational bidder's utility function is convex and monotone. Given buyer's value x, the utility by choosing menu items (b, q, t) is $xq - t$. So the rational buyer's utility is $u(x) = \sup_{(b,q,t)}\{xq - t\}$, which is the supremum of a set of linear functions of x. Thus, u is convex.

Let $q(x)$ denote the allocation probability when the rational bidder achieves $u(x)$.

$$\begin{aligned} & u(x') - u(x) - q(x)(x' - x) \\ = \; & x'q(x') - t(x') - xq(x) + t(x) - q(x)(x' - x) \\ = \; & x'q(x') - t(x') - (x'q(x) - t(x)) \geq 0 \end{aligned}$$

The last inequality is because $(q(x'), t(x'))$ is the best choice for the rational bidder with value x'. Substitute x' twice by $x^- = x - \epsilon$ and $x^+ = x + \epsilon$ respectively, for any arbitrarily small positive ϵ, we have $u'(x^-) \leq q(x) \leq u'(x^+)$. This imply u is differentiable almost everywhere and $u'(x) = q(x)$. So the rational buyer chooses the menu with allocation probability $u'(x)$ and payment $u'(x)x - u(x)$. Since allocation probability is nonnegative, so u is weakly monotone.

Let $R_{u,i}$ denote the revenue from the irrational bidder based on the mechanism based on u. Let $R_{u,r}$ denote the revenue from the rational bidder based on the mechanism based on u. The revenue getting from the rational buyer is

$$R_{u,r} = \int_0^1 (u'(x)x - u(x))f(x)dx$$

When we fix the utility function for the rational agent, this fixes the price of every menu item chosen by some type of the rational agent. We are flexible to set prices for menu items that are not chosen as long as (a) we only offer items we can deliver(feasibility constraint) (b) don't charge more than an agent's value(IR constraint) (c) don't make some rational agent want to choose this item (IC constraint). When bidder is irrational with private value b, the price given by the following lemma is the largest subject to these 3 constraints.

Lemma 2 *If $b < 1 - u(1)/u'(1)$, the largest payment is the intercept between utility axis and the line that goes through point $(b, 0)$ and touches the utility line. If $b \geq 1 - u(1)/u'(1)$, the largest payement is the b.*

Proof: Suppose in mechanism M, we assign the irrational

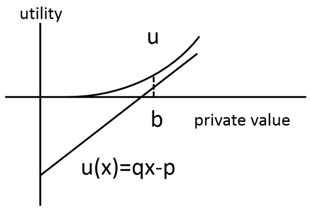

Figure 3: menu chosen by the irrational bidder

bidder the menu (q, p), where p is the price and q is the allocation probability of the item. Look at Fig 3, straight line $u(x) = qx - p$ denotes the buyer's utility function of choosing the menu (q, p). If this line is above the utility curve on some interval, say (x_1, x_2), then the utility of the rational buyer with private value between $[x_1, x_2]$ increases by choosing menu (q, p). So the straight line is below the utility curve. Furthermore, to guarantee the irrational buyer not losing money, this straight line is above the point $(b, 0)$, i.e., the utility of the irrational buyer is nonnegative. Remember that the allocation probability $q \in [0, 1]$. With these three constraints, what's the possible largest payment charging the irrational buyer with value b?

Look at Fig 4. There are 2 cases. Case 1: $b = b_1 <$

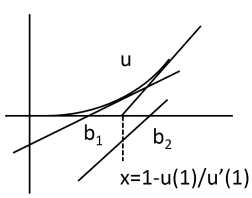

Figure 4: 2 cases

$1 - u(1)/u'(1)$, the largest payment is the intercept between u axis and the line that goes through point $(b_1, 0)$ and touches the utility line. Case 2: $b = b_2 \geq 1 - u(1)/u'(1)$, the largest payement is the b_2, i.e. when line $y = (x - 1)u'(1) + u(1)$ doesn't intersect with utility line. ■

Given any convex and monotone function u with $u(0) = 0$ and $u' \leq 1$, we create a mechanism. For every bid $b \in [0, 1]$, the payment and allocation function is defined below:

$b < 1 - \frac{u(1)}{u'(1)}$	allocation probability = q, payment = p (q is gradient of the touching line, p is the absolute y-value of the cross of y-axis and touching line)
$b \in [1 - \frac{u(1)}{u'(1)}, 1 - u(1)]$	allocation probability = $\frac{u(1)}{1-b}$, payment = $\frac{u(1)b}{1-b}$
$b \in (1 - u(1)]$	allocation probability=1, payment=b

In this mechanism, the rational bidder with private value v achieve the maximal utility by bidding $v - \frac{u(v)}{u'(v)}$. By Lemma 2, we get the maximal irrational payment. Since the revenue from the rational bidder is fixed, this is the optimal mechanism given the rational bidder's utility function.

4 Characterization of the optimal utility function

This section studies the structure of the optimal utility function. We find an important property that when the utility function has a breaking point, we can increase revenue on irrational bidder compared to the loss from the rational bidder by slightly modifying the utility function. The technique is to first focus on a small range around the breaking point where $g(x)$ and $f(x)$ does not change much. Then we compare the gain and loss carefully.

Theorem 1 *When 0<p<1, the optimal utility function belongs to C^1(whose derivative is continuous). Thus posted-price auctions are not optimal(the derivative is not continuous).*

The result also implies that there are infinite many menu pieces in the optimal mechanism.

Proof: Since $u' \in [0, 1]$ and u' is weakly increasing, if u' is continuous then u belongs to C^1. To complete the proof, we will prove that u' is continuous. If u' is not continuous, it must have a breaking point.

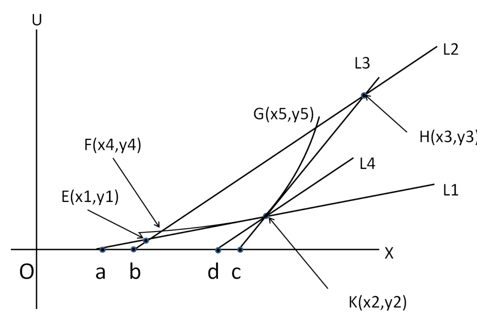

Figure 5: The utility function is curve FKG, where $K(x_2, y_2)$ is a breaking point. L_1 and L_3 are the two touching lines on both sides of point K. L_1 and L_3 has gradient $k_1 = \frac{y_2}{x_2 - a}$, and $k_3 = \frac{y_2}{x_2 - c}$ respectively. We pick point $E(x_1 = x_2 - z, y_1 = y_2 - k_1 z)$ on L_1 and $H(x_3 = x_2 + z, y_3 = y_2 + k_3 z)$ on L_3. Let the straight line goes through these two points be L_2, it has gradient $k_2 = \frac{k_1 + k_3}{2}$ and crosses x axis at point $(b = \frac{k_3 - k_1}{k_3 + k_1} x_2 + \frac{2k_1}{k_1 + k_3} a + \frac{k_1 - k_3}{k_1 + k_3} z, 0)$. $F(x_4, y_4)$ and $G(x_5, y_5)$ are two intersections between L_2 and utility curve. L_4 is parallel to L_2.

Look at Fig.5, we define function u_N as follows. We will prove u_N generates more revenue than u.

$$u_N(x) = \begin{cases} u(x) & 0 \leq x \leq x_4 \\ y_1 + k_2 * (x - x_1) & x_4 < x \leq x_5 \\ u(x) & x_5 < x \leq 1 \end{cases}$$

First we consider the gain from the irrational bidder. We get at least the same profit on every irrational bidder. We only need to consider the gain on interval $[d, c]$.

$$\begin{aligned} & R_{u_N,i} - R_{u,i} \\ = & \int_d^c \frac{y_2 x z}{x_2 + z - x} [\frac{1}{x_2 - c} - \frac{1}{x_2 - x}] g(x) dx \\ \geq & \frac{y_2 dz}{2(x_2 - d)} \int_d^c [\frac{1}{x_2 - c} - \frac{1}{x_2 - x}] g(x) dx \end{aligned}$$

Where the inequality is true when $z < x_2 - d$.

Since $f(x)$ is continuous, for any $\epsilon > 0$, we can pick small enough z such that $|f(x) - f(x_2)| < \epsilon$, for $x \in (x_2 - z, x_2 + z)$. The loss on the rational bidder changes a little:

$$\begin{aligned}
& R_{u,r} - R_{u_N,r} \\
= & \int_{x_4}^{x_5} [(u'(x)x - u(x)) - (u'_N(x)x - u_N(x))] f(x)dx \\
\leq & \int_{x_2}^{x_5} [u'(x)x - u(x) - u'_N(x)x + u_N(x)](f(x_2) + \epsilon)dx \\
+ & \int_{x_4}^{x_2} [u'(x)x - u(x) - u'_N(x)x + u_N(x)](f(x_2) - \epsilon)dx \\
\leq & f(x_2)[u(x)x|_{x_2}^{x_5} - 2\int_{x_4}^{x_5} u(x)dx - u_N(x)x|_{x_2}^{x_5} \\
& +2\int_{x_4}^{x_5} u_N(x)dx] + \epsilon(x_5 - x_2) \\
\leq & 2f(x_2)S_{FKG} + \epsilon(x_3 - x_2) \leq 2f(x_2)z^2/2 + \epsilon z
\end{aligned}$$

We want to prove $R_{u,r} - R_{u_N,r} < R_{u_N,i} - R_{u,i}$, i.e., $z^2 f(x_2)(k_3 - k_1) + z\epsilon x_2(k_3 - k_1) < \frac{y_2 dz}{2(x_2-d)} \int_d^c [\frac{1}{x_2-c} - \frac{1}{x_2-x}]g(x)dx$. Notice that all terms are constant except ϵ and z. Obviously, we can pick small enough ϵ and z to make sure $R_{u_N,i} - R_{u,i} > R_{u,r} - R_{u_N,r}$. That means u_N outperform u. So u' has no breaking point, thus $u \in C^1$. ■

In the following we study some other properties of the optimal utility function. Main technique is the variation of function. From these three properties, one can see that the optimal utility function is complex.

Theorem 2 *If u is the optimal utility function, then*

1. $u(0) = 0$.
2. *u strictly increases.*
3. $u'(x) = 1, x \geq \inf \arg\max_z [1 - F(z)]z$

Proof: (1)Suppose otherwise, we draw a line through zero point and touches utility function at $(l, u(l))$(Fig 6). Define

$$u_N(x) = \begin{cases} \frac{u(l)}{l}x & 0 \leq x \leq l \\ u(x) & l < x \leq 1 \end{cases}$$

u_N generates strictly more revenue when bidder is rational and same revenue when bidder is irrational.

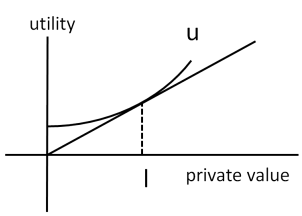

Figure 6: u(0) is 0.

(2) Because u is convex, we only need to prove that $u(x) = 0$ has unique solution in $[0, 1]$. Suppose otherwise, $a = argmax_x\{u(x) = 0\} > 0$. For any $z < a$, define

$$u_N(x) = \begin{cases} 0 & 0 \leq x \leq a - z \\ u(a + z) \cdot \frac{x - a + z}{2z} & a - z < x \leq a + z \\ u(x) & a + z < x \leq 1 \end{cases}$$

We claim u_N generates more revenue than u when z is small enough. We first consider the gain from the irrational bidder.

$$\begin{aligned}
R_{u_N,i} - R_{u,i} & \geq \int_{a-z}^{a} u(a + z)\frac{x}{a + z - x}(g(a) - \epsilon)dx \\
& = u(a + z)(g(a) - \epsilon)((a + z)ln2 - z)
\end{aligned}$$

Before consider the loss from the rational bidder, we need the following inequality for simplicity.

$$\int_a^{a+z} [u'(x)x - u(x)]dx \leq (a + z)u(a + z) - u^2(a + z)$$

The proof uses u is convex and $u(a) = 0$, the proof is omitted. The loss from the rational bidder is:

$$\begin{aligned}
& R_{u,r} - R_{u_N,r} \\
= & \int_a^{a+z} [u'(x)x - u(x)]f(x)dx \\
& - \int_{a-t}^{a+z} u(a + z)\frac{a - z}{2z} f(x)dx \\
\leq & f(a)[2zu(a + z) - u^2(a + z)] \\
& +\epsilon[2au(a + z) - u^2(a + z)]
\end{aligned}$$

We prove the gain is larger than the loss when z is small.

$$\begin{aligned}
& R_{u_N,i} - R_{u,i} \geq R_{u,r} - R_{u_N,r} \\
\Leftarrow & u(a + z)(f(a) - \epsilon)((a + z)ln2 - z) > \\
& f(a)[2zu(a + z) - u^2(a + z)] \\
& +\epsilon[2au(a + z) - u^2(a + z)]
\end{aligned}$$

When ϵ and z approaches zero, the left side is $g(a)aln2 > 0$, the right side is 0. Hence the equality is correct, which means u_N outperform u, contradicts to u is optimal.

(3) Suppose argument is not ture, let $l = \inf \arg\max_z [1 - F(z)]z$ and we have $u'(l) < 1$. Define

$$u'_N(x) = \begin{cases} u'(x) & 0 \leq x \leq l \\ 1 & l < x \leq 1 \end{cases}$$

We prove u_N generates more revenue than u on both irrational bidder and rational bidder.

$$\begin{aligned}
R_{u,r} & = \int_0^1 (u'(x)x - u(x))f(x)dx \\
& = \int_0^l (u'(x)x - u(x))f(x)dx \\
& + \int_l^1 xf(x)u'(x)dx - \int_l^1 u(x)dF(x)
\end{aligned}$$

Since $u(x)$ and $u_N(x)$ differs when $x > l$, so $R_{u,r}$ and $R_{u_N,r}$ differs only at the third term. Because $u_N(x) > u(x)$, $x \in (l, 1]$, we have $R_{u,r} < R_{u_N,r}$.

For the irrational bidder, no matter what the bid is, we can always get more revenue in u_N than in u. In both irrational and rational cases, u_N gives strictly more revenue than u. ■

5 Representing revenue from the irrational bidder via utility function

Given function u, the optimal allocation and payment rule for the irrational bidder is also determined. Let $h(x) = x - \frac{u(x)}{u'(x)}$, by Lemma 2. When $x < 1 - u(1)$, we set the irrational buyer with value $h(x)$ the same menu item as the rational buyer with value x chooses. When $x \geq 1 - u(1)$, we set new menu item: the probability 1 and payment x (Fig 7). Because u is convex, so $h(x)$ is continuous and monotone, so $h^{-1}(x)$ is well defined.

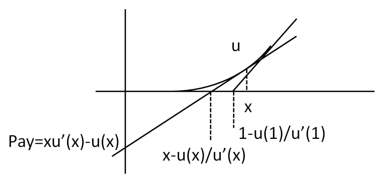

Figure 7: The irrational bidder's payment

So, we have the revenue getting from the irrational buyer:

$$R_{u,i} = \int_0^{1-u(1)} [h^{-1}(x)u'(h^{-1}(x)) - u(h^{-1}(x))]g(x)dx + \int_{1-u(1)}^1 xg(x)dx$$

The representation of total revenue $(1-p)R_{u,i} + pR_{u,r}$ is long and hard to compute the optimal utility function. Take the simplest example: F and G are both uniform distributions and have equal probability of being rational and irrational, i.e. $g = f = 1$ and $p = 0.5$. We also assume u is twice differentiable. Then the total revenue is:

$$(1-p)R_{u,i} + pR_{u,r} \tag{1}$$
$$= u(1) + \frac{u(1)^2}{4} - 2\int_0^1 u(y)dy + \frac{1}{2}\int_0^1 \frac{u(y)}{u'(y)} ydu'(y)$$

Even in this simple setting, we still do not know how to compute the optimal u with the constraints that $u \geq 0$ and $u' \geq 0$ is weakly increasing. We relax the problem and solve the optimization problem with no constraint. After a long computation, we find the optimal solution is in the form $u(t) = c_1 t^{2+\sqrt{3}} + c_2 t^{2-\sqrt{3}}$. The optimal solution has infinite revenue, and does not satisfy our constraints obviously.

For general distribution and probability p, the difficulty is how to optimize with convex and monotone constraints. This is a research agenda in general mechanism design.

We now put forward an important result that states the one-shot value-revelation revenue we have computed so far is equal to the optimal unconstrained revenue where bidders can report the entire type. This result justifies our initial choice of focusing on one-shot value revelation mechanisms.

Theorem 3 *The revenue in the optimal mechanism M_1 which allows type revelation, is the same as that in the optimal mechanism M_2 which only allows value revelation.*

Proof: We start from any type revelation mechanism M that satisfies IR and IC constraints. Apply the following operation to M: we set the payment from the irrational bidder to be the largest payment, among all mechanisms that have the same utility function of the rational bidder as M and satisfy IR and IC constraints. Denote the new mechanism by M'.

In fact, Lemma 2 continues to hold for the irrational bidder, even under the type revelation setting. The only difference in the argument is when we show it is impossible for the line $u(x) = qx - p$ to be above the utility curve, in the value-revelation setting, we created a manipulation (thus a violation of IC) where the rational bidder must choose a different menu, while in the type-revelation setting, we can replicate the same manipulation by letting the rational bidder choosing a different menu (i.e., misreport value) and misreport his rationality mode at the same time.

As a result, in the new mechanism M', the payment from the irrational bidder is the largest payment stated in Lemma 2. Also guaranteed by Lemma 2, Mechanism M' satisfies the IR and IC constraints and its total revenue increases since the revenue from the rational bidder is the same while the revenue from the irrational bidder increases. Hence, the optimal type revelation mechanism M_1 must be among the mechanisms where the payment from the irrational is set to be the largest, fixing the utility function of the rational bidder.

In the optimal type-revelation mechanism M_1, let u denote the utility function of the rational bidder. The total revenue of M_1 is $(1-p)R_{u,i} + pR_{u,r}$, exactly the same as Equation (1), the revenue of M_2. Both mechanisms choose u to be the one which maximizes $(1-p)R_{u,i} + pR_{u,r}$. In other words, the utility function for the rational bidder as well as the optimal revenue are exactly the same in M_1 and M_2. ■

However, M_1 and M_2 differs in their representations. M_2 is defined by the table in Section 3, the rational misreports his valuation in M_2. M_1 conditioning on the irrational mode is the same as M_2, and M_1 conditioning on the rational mode is to allocate the item with probability $u'(b)$ at price $u'(b)b - u(b)$. The rational is truthful in M_1.

6 Approximation via k piecewise linear function

We give an approximation by k piecewise linear function. We show the approximation error can guarantee to be inversely proportional to the number of piece. While the main idea is to pick points on utility curve and this method can maintain the revenue from the irrational bidder. In this case, it is crucial to pick the positions of these breaking points.

Theorem 4 *The difference between revenue of optimal mechanism and optimal k piecewise linear function is at most $1/k$.*

Proof: We want to find $k-1$ points on the optimal utility curve u, say they are $(a_i, u(a_i)), i = 1...k-1$. Let $a_0 = 0$ and $a_k = 1$. We construct the new utility function u_N by connecting $(a_i, u(a_i))$ and $(a_{i+1}, u(a_{i+1}))$ for $i = 0, 1, ..k-1$. The idea is we guarantee the revenue from the irrational bidder increases, and control the loss from the rational bidder to be smaller than $1/k$.

In the following, we only consider the rational case. The payment in the new mechanism u_N when bidder's private value is in $[a_i, a_{i+1}]$, is at least the same as the payment in the old mechanism u when bidder's private value is in $[a_{i-1}, a_i]$. For simplicity, we use $pay(x)$ to denote the payment in the old mechanism, i.e. $pay(x) = xu'(x) - u(x)$. In the old mechanism the revenue is

$$\int_0^1 pay(x)f(x)dx$$

In the new mechanism the revenue is at least

$$\sum_{i=0}^{k-1} pay(a_i)[F(a_{i+1}) - F(a_i)]$$

The difference is at most

$$\sum_{i=0}^{k} (pay(a_i) - pay(a_{i-1}))[F(a_i) - F(a_i - 1)]$$

Pick a_i such that $pay(a_i) = i/k$ for any i. Then the difference is at most $1/k$. ■

Proposition 1 *For uniform distribution, the difference between the optimal mechanism and optimal k piecewise linear function is at most $1/k^2$.*

7 Multiple bidders

In this section, we extend the previous discussion to n bidders. We only consider the uniform distribution of valuations, and $c(x) = x$. In first price auction, we can extract maximum revenue from the irrational bidder.

Theorem 5 *When n bidders have identical uniform valuation distributions and when the irrational bidders are truthful, first price auction generates more revenue than second price auction.*

Proof: First we need to find BNE in the first price auction. Suppose the rational bidder strategy is $s(x)$. The utility of the rational bidder with private value x by reporting $s(t)$ is

$$\begin{aligned} u(x,t) &= (x - s(t))[pt + (1-p)s(t)]^{n-1} \\ \frac{\partial u}{\partial t} &= -s'(t)[pt + (1-p)s(t)]^{n-1} \\ &\quad +(x - s(t))(n-1)[pt + (1-p)s(t)]^{n-2} \cdot \\ &\quad (pt + (1-p)s'(t)) \end{aligned}$$

Because bidding $s(x)$ is a best strategy, we have

$$\begin{aligned} \frac{\partial u}{\partial t}|_{t=x} = 0 \quad \Rightarrow \quad & -s'(x)[px + (1-p)s(x)] \\ & +(x - s(x))(n-1)(p + (1-p)s'(x)) = 0 \end{aligned}$$

By simplification, we find the unique BNE is $s(x) = \frac{n-1}{n}x$, exactly the same as the standard setting. For any bidder, the payment from both irrational and rational aspects, is

$$\begin{aligned} & p\int_0^1 \frac{n-1}{n}x[px + (1-p)\frac{n-1}{n}x]^{n-1} \\ & +(1-p)[\int_{\frac{n-1}{n}}^1 x[p + (1-p)x]^{n-1}dx \\ & +\int_0^{\frac{n-1}{n}} x[p\frac{n}{n-1}x + (1-p)x]^{n-1}dx] \\ = \quad & \frac{n-p}{n+1} + \frac{1}{n+1}\frac{p}{1-p}[(1 - \frac{1-p}{n})^n - p] > \frac{n-p}{n+1} \end{aligned}$$

When we consider "truthful" mechanism, in which the rational bidder's optimal bid is her private valuation. 2nd price auction with reserve 0.5 is the optimal truthful mechanism with revenue

$$\frac{n-1}{n+1} + \frac{n}{2^n}(\frac{1}{2} - \frac{n-1}{2n}) + \frac{1}{2^n}(0 - \frac{n-1}{2n+2}) = \frac{n-1}{n+1} + \frac{1}{(n+1)2^n}$$

So 1st price auction yields approximate $\frac{1-p}{n+1}$ more revenue. ■

References

[Celis *et al.*, 2012] L. Elisa Celis, Anna R. Karlin, Kevin Leyton-Brown, C. Thach Nguyen, and David Robert Martin Thompson. Approximately revenue-maximizing auctions for deliberative agents. In *AAAI*, 2012.

[Hart and Nisan, 2013] Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. In *ACM EC*, pages 565–566, 2013.

[Hartline, 2013] Jason D Hartline. Mechanism design and approximation. *Book draft. October*, 2013.

[Larson and Sandholm, 2000] Kate Larson and Tuomas Sandholm. Costly valuation computation in auctions. In *AAAI*, 2000.

[Larson and Sandholm, 2004] Kate Larson and Tuomas Sandholm. Experiments on deliberation equilibria in auctions. In *AAMAS, 19-23 August 2004, New York, NY, USA*, pages 394–401, 2004.

[Lee and Malmendier, 2011] Young Han Lee and Ulrike Malmendier. The bidder's curse. *American Economic Review*, 101(2):749–787, April 2011.

[Manelli and Vincent, 2007] Alejandro M Manelli and Daniel R Vincent. Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *Journal of Economic Theory*, 137(1):153–185, 2007.

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

[Myerson, 1981] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[Osborne and Rubinstein, 1994] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. MIT Press, 1994.

[Pai and Vohra, 2014] Mallesh M. Pai and Rakesh Vohra. Optimal auctions with financially constrained buyers. *Journal Economic Theory*, 150:383–425, 2014.

[Rubinstein, 1986] Ariel Rubinstein. Finite automata play the repeated prisoner's dilemma. *Journal of Economic Theory*, pages 83–96, 1986.

[Shoham and Leyton-Brown, 2009] Yoav Shoham and Kevin Leyton-Brown. *Multiagent Systems: Algorithmic, Game theoretic and Logical Fundations.* Cambridge Uni. Press, 2009.

[Tang and Sandholm, 2012] Pingzhong Tang and Tuomas Sandholm. Optimal auctions for spiteful bidders. In *AAAI*, 2012.

[Von Neumann and Morgenstern, 1947] J. Von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1947.

[Wang and Tang, 2014] Zihe Wang and Pingzhong Tang. Optimal mechanisms with simple menus. In *ACM EC*, pages 227–240. ACM, 2014.

[Wright and Leyton-Brown, 2010] James R. Wright and Kevin Leyton-Brown. Beyond equilibrium: Predicting human behavior in normal-form games. In *AAAI*, 2010.

[Wright and Leyton-Brown, 2012] James R. Wright and Kevin Leyton-Brown. Behavioral game theoretic models: a bayesian framework for parameter analysis. In *AAMAS*, pages 921–930, 2012.

[Wright and Leyton-Brown, 2014] James R. Wright and Kevin Leyton-Brown. Level-0 meta-models for predicting human behavior in games. In *ACM EC*, pages 857–874, 2014.