

Optimal auctions for spiteful bidders*

Pingzhong Tang and Tuomas Sandholm

Computer Science Department
Carnegie Mellon University

{kenshin, sandholm}@cs.cmu.edu

Abstract

Designing revenue-optimal auctions for various settings is perhaps the most important, yet sometimes most elusive, problem in mechanism design. Spiteful bidders have been intensely studied recently, especially because spite occurs in many applications in multiagent system and electronic commerce. We derive the optimal auction for such bidders (as well as bidders that are altruistic).

It is a generalization of Myerson's (1981) auction. It chooses an allocation that maximizes agents' virtual valuations, but for a generalized definition of virtual valuation. The payment rule is less intuitive. For one, it takes each bidder's own report into consideration when determining his payment. Moreover, bidders pay even if the seller keeps the item; a similar phenomenon has been shown in other settings with negative externalities (Jehiel, Moldovanu, and Stacchetti 1996; Deng and Pekeč 2011). On the other hand, a novel aspect of our auction is that it sometimes subsidizes losers when the item is sold to some other bidder.

We also derive a revenue equivalence theorem for this setting. Using it, we generate a short proof of (a slight generalization of) the previously known result that, in two-bidder settings with independently uniformly drawn valuations, second-price auctions yield greater expected revenue than first-price auctions. Finally, we present a template for comparing the expected revenues of any two auction mechanisms that have the same allocation rule (for the valuations distributions at hand).

Introduction

Auctions are a key class of methods for resource and task allocation in multiagent systems. One of the most important problems in auction theory, and mechanism design at large, is to design revenue-maximizing (aka. *optimal*) auctions. The optimal auction design problem is, for the seller, to design an auction mechanism that maximizes her expected revenue, given the information about bidders' valuation distributions but not the actual values.

There has been a significant amount of research on this topic, most of which has been within the standard single-parameter quasi-linear settings (cf. (Nisan et al. 2007, Chapter 9)). Landmark results include Myerson's (1981) auction

for the standard one-item setting as well as its extension to multiple identical units of an item (Maskin and Riley 1989).

Since those landmark results, several optimal auctions in other special settings have been put forward. To name a few, there is Levin's auction for complements (Levin 1997), multi-item auctions for settings where the bidders' valuations can take on either of only two values (Armstrong 2000; Avery and Hendershott 2000), a 1-item auction for financially constrained bidders (Pai and Vohra 2008), sponsored search auctions (Iyengar and Kumar 2006), multi-item auction with single-minded agents (Ledyard 2007), a procurement auction where both costs and capacities are private information (Iyengar and Kumar 2008), an optimal auction where each bidder's private information (i.e., *type*) is two-dimensional where one dimension describes valuation and the other describes externality (Deng and Pekeč 2011), and an optimal auction where the externality that a bidder experiences is determined by the payments of others (but not by the allocation to others) (Lu 2012).

In this paper, we investigate the optimal auction design problem in the setting with spiteful bidders (and altruistic bidders where the spite factor, defined later, is negative). Loosely speaking, spite is a type of negative externality an agent imposes on others when he wins. Early on, auctions with negative externalities were studied in the context of nuclear weapons (Jehiel, Moldovanu, and Stacchetti 1996)¹. Spite and similar negative externalities have drawn significant interest recently, especially in AI and sponsored search literature (Morgan, Steiglitz, and Reis 2003; Brandt, Sandholm, and Shoham 2007; Zhou and Lukose 2007; Sharma and Sandholm 2010; Krysta et al. 2010; Constantin et al. 2011; Deng and Pekeč 2011; Conitzer and Sandholm 2012). Symmetrical equilibria have been calculated in auction settings with spite for first- and second-price auctions, and revenue comparisons have been derived (Morgan, Steiglitz, and Reis 2003; Brandt, Sandholm, and Shoham 2007; Sharma and Sandholm 2010). A type of vindictive bidding strategy (a strategy that decreases other agents' utilities while maintains one's own utility) has been empirically

*This work is supported by the National Science Foundation under IIS grants 0905390 and 0964579 and by CCF-1101668. Copyright © 2012, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Ukraine had old-fashioned nuclear weapons for sale. The USA and Russia, who were not interested in the weapons themselves, were worried about the danger caused by some other country obtaining the weapons. Consequently, both countries paid significant amounts to Ukraine not to sell any of the weapons.

observed in sponsored search auctions (Zhou and Lukose 2007). Bidding languages that allow one to express certain kinds of externalities have been designed for sponsored search (Parkes and Sandholm 2005; Constantin et al. 2011). Expressive bidding languages and winner determination have been studied for general settings (Krysta et al. 2010; Conitzer and Sandholm 2012; 2011).

The most closely related prior research is the analysis of single-item auctions with spiteful bidders (Morgan, Steiglitz, and Reis 2003; Brandt, Sandholm, and Shoham 2007), where each agent cares about the others' utility (which depends on the others' allocation and payments). We use the same model of spite as that prior work. To our knowledge, the optimal auction in this model was unknown. The prior papers studied first- and second-price auctions, but neither of them is optimal. One aspect that makes our analysis more involved is that to derive an optimal auction, we need to consider the possibility that it might be randomized. (As we will show, randomization turns out not to be necessary.) Furthermore, it is not even clear in this setting whether the losing bidders should pay, get paid, or neither.

Adapting Myerson's approach (Myerson 1981), we give complete characterizations of incentive compatible individually rational auctions, as well as optimal auctions for this setting. The optimal auction, which is a generalization of Myerson's auction to the spiteful setting, takes the following form. It chooses an allocation that maximizes agents' virtual valuations—for a revised definition of virtual valuation—similar to the allocation rule of Myerson's auction. The payment rule, though including Myerson's payment rule as a special case, is less intuitive. For one, it takes each bidder's own report into consideration when determining his payment. Moreover, when the seller keeps the item, bidders usually pay positive amounts for nothing, a phenomenon that rarely occurs in commonly seen auctions. This is similar to the conclusion drawn from a different model of negative externalities that the seller should extract positive rents from losing bidders (Jehiel, Moldovanu, and Stacchetti 1996). Our auction sometimes subsidizes losers when the item is sold to some other bidder—a phenomenon which, to our knowledge, has not been observed in prior work.

We also derive a version of the revenue equivalence theorem for this setting. Using this theorem, we are able to easily compare the revenue of first- and second-price auctions, that is, we obtain a shorter analysis of previously known results (and generalize them slightly).

The setting

We consider a setting where the seller has an indivisible item for sale. Her valuation of the item is normalized to 0. There is a set N of n bidders. Each bidder i has a private valuation t_i for the item. Only agent i knows his valuation. All the others, including the seller, view t_i as a random variable that is drawn from a distribution function $F_i(t_i)$ on a closed interval $[a_i, b_i]$. F_i admits a density function f_i that is positive everywhere on $[a_i, b_i]$. We also use the standard assumption that t_i and t_j are independently distributed, for any $i \neq j$. We denote the joint type by $t = t_1 \times t_2 \times \dots \times t_n$ and the joint type distribution by $F = F_1 \times F_2 \times \dots \times F_n$.

All of our analysis also applies to the setting where the seller has k identical units for sale and each bidder has unit demand. We will point out the differences of the two settings whenever necessary. From now on, for simplicity of presentation, we will word everything in the single-unit setting unless explicitly stated otherwise.

By the revelation principle (Myerson 1981), we can, without loss of generality, focus on the set of direct revelation mechanisms. Each bidder is asked simply to report his valuation in $[a_i, b_i]$. We call the reported valuation a *bid*. Upon receiving a bid from each bidder, the auction determines the probabilities p_i with which each agent i wins the item, and the payments x_i that the bidders have to make.

The setting described in this section so far is the standard *independent private value (IPV)* setting. The key deviation point from the standard setting is the utility functions of the bidders, which will now be spiteful. The utility of agent i is

$$u_i(t) = t_i p_i(t) - x_i(t) - \alpha \sum_{j \neq i} (t_j p_j(t) - x_j(t)).$$

The first two terms define the standard quasi-linear utility. The third term reflects the externality others impose on agent i when they get the item. In other words, an agent's utility is his utility in the standard sense minus the welfare of the other agents times a constant. The constant, α , coined *spite factor*, is symmetric among all the agents. Note that α can be either positive (corresponding to spite) or negative (corresponding to altruism).

Our model is a generalization of prior analyses of auctions with spiteful bidders (Morgan, Steiglitz, and Reis 2003; Brandt, Sandholm, and Shoham 2007) because the prior work studied deterministic auctions (where p_i can only be either 1 or 0) while we allow randomization. This is necessary because we cannot preclude ahead of time the possibility that the optimal auction might need to use randomization.

Due to spite, our setting differs from the standard "single-parameter" environment (cf. (Nisan et al. 2007, Chapter 9)), where the utility is quasi-linear. Our setting also differs from the standard combinatorial auction setting, where externality is absent. Our setting also differs from prior research on externalities in how we model externality. Our externality (specifically, spite) depends on the utilities and payments of others, while Jehiel et. al. (1996) and Deng and Pekec (2011) considered externality to be a number that only depends on the name of the agent who gets the item (no matter how much that agent pays).

When the spite factor $\alpha = 0$ or the number of bidders $n = 1$, the setting reduces to the standard IPV setting. To review Myerson's solution to the optimal auction design problem, we first define the *virtual valuation function* in the standard way as $\tilde{t}_i = t_i - \frac{1 - F_i(t_i)}{f_i(t_i)}$. The allocation rule of the Myerson auction is then to maximize the sum of virtual valuations of all agents (including the seller), that is, the virtual social welfare. The payment rule is for the winning bidder, if any (the seller may have the highest virtual valuation, in which case no bidder wins), to pay the amount of the lowest bid with which he would have won.

If all bidders have the same valuation distribution function ($F_i = F_j$ for all i, j), the virtual valuation function is bidder

independent. Assuming \tilde{t}_i is increasing in t_i (i.e., the standard regularity condition), the bidder with highest virtual valuation is the one with highest actual valuation. Therefore, Myerson's auction reduces to the second-price auction, but with a reserve price (represented by the fact that the seller's virtual valuation might be higher than any bidder's).

Following Myerson's approach, we give a detailed analysis of the seller's problem. In the next section, we define the seller's problem. In the following section, we characterize the incentive compatible individually rational auctions. In the section after that, we derive the seller's revenue formula as a function of only the allocation rule and the utility of each bidder when he happens to have his lowest possible valuation. (As a side effect, we obtain a revenue equivalence theorem for our spiteful bidders.) In the section after that, we present the optimal auction. In the section after that, we provide a more compact way of conducting revenue comparisons of auctions than what was available in prior work.

The seller's problem

Informally, the seller's problem is to maximizing her expected revenue, subject to *incentive compatibility (IC)* and *individual rationality (IR)* constraints in Bayes Nash equilibrium. In addition, we also need the *resource (RS)* constraint, stating that there is only one object for sale.

We first need to write down agent i 's (*ex interim*) utility when i bids truthfully:

$$U_i(t_i) = \int_{t_{-i}} t_i p_i(t) - x_i(t) - \alpha \sum_{j \neq i} (t_j p_j(t) - x_j(t)) dF_{-i}(t_{-i}).$$

His utility, when i bids s_i instead of his true type t_i , is:

$$U'_i(s_i) = \int_{t_{-i}} t_i p_i(s_i, t_{-i}) - x_i(s_i, t_{-i}) - \alpha \sum_{j \neq i} (t_j p_j(s_i, t_{-i}) - x_j(s_i, t_{-i})) dF_{-i}(t_{-i}).$$

Incentive compatibility states that reporting one's true type is no worse for that bidder than reporting any other type, assuming others also report truthfully.

Definition 1 *Incentive compatibility (IC)*

An auction is *incentive compatible in Bayesian Nash equilibrium* if $U_i(t_i) \geq U'_i(s_i)$ for all i, t_i and s_i .

The next constraints we will discuss are individual rationality (IR) constraints (aka participation constraints). As always, IR here states that each agent should be at least as well off by participating as he would be by not participating.

However, unlike in the usual (spiteless) model, an agent in our model will not simply get 0 utility if he does not participate: for example, if $\alpha > 0$, he will get negative utility when the other bidders in aggregate get positive utility (because some other bidder gets the item and/or some other bidder gets paid).² Thus, IR requires the seller to promise each

²This is similar in the nuclear weapons sale example, where a non-participant can face dire consequences if a rival of his gets the weapons—unless the non-participant can somehow make himself immune to the bad consequence, such as by moving to Mars, without changing his utility.

participant expected utility greater than that value. To maximize revenue, the seller needs to construct the worst possible bidder-specific realistic threat with which to threaten each potential non-participant. This will minimize a bidder's utility from non-participation and therefore maximize how much the seller can extract in the auction while still keeping the auction IR.³ The seller accomplishes this by, for each bidder i in turn, threatening to maximize (minimize if $\alpha < 0$) the aggregate value to the other bidders $\int_{t_{-i}} \sum_{j \neq i} (t_j p'_j(t) - x'_j(t)) dF_{-i}(t_{-i})$ (thus damaging i as much as possible according to i 's utility function $U_i(t_i)$ above) if i does not participate. Here, p' and x' are the allocation and payment rule of the "second-stage" mechanism where the players are $N \setminus \{i\}$.

The "second-stage" mechanism would still need to be Bayes Nash IC for the remaining agents. However, it would not have to be IR for them because they already will have agreed to participate up front, and the mechanism will not ask them to reconsider participation before the "second stage". Note that the "second stage" will never actually be reached on the path of play because all agents are incentivized to participate in the first place.

Now, we introduce the following shorthand notation: $g(t_{-i}) = \sum_{j \neq i} (t_j p'_j(t) - x'_j(t))$. With that, we are ready to define the IR constraints.

Definition 2 *Individual rationality (IR)*

An auction is *individually rational in Bayesian Nash equilibrium* if $U_i(t_i) \geq -\alpha \int_{t_{-i}} g(t_{-i}) dF(t_{-i})$ for all i and t_i .

The value of $g(t_{-i})$ depends on the power of the seller. If there are no restrictions and $\alpha > 0$, the seller can threaten to pay the remaining agents infinite amounts, thus setting $g(t_{-i})$ to infinity. In that setting, the optimal revenue is trivially infinite. (Figuroa and Skreta (2009) discuss a spiteless auction setting with two bidders where the seller has two options: 1) to burn the item if even one bidder does not participate (this is also the assumption made by Lu (2012)), or 2) give the item to the participating bidder in that scenario. Clearly neither of these threats is credible because the auctioneer would be better off by re-auctioning.)

Another example would be a seller who has to be budget balanced, that is, she cannot make a positive payment to the agents in aggregate. Now the seller's threat to a potential non-participant i is that if i does not participate, the seller will run a mechanism that tries to maximize the sum of the utilities of the remaining agents $\int_{t_{-i}} \sum_{j \neq i} (t_j p'_j(t) - x'_j(t)) dF_{-i}(t_{-i})$ subject to budget balance.⁴

³Here, the equilibrium payoff is the equilibrium payoff when all other bidders commit to participate. Because (Bayes) Nash equilibrium is used as the solution concept, we only have to guarantee that no agent alone is motivated to not participate. Simultaneous non-participation by multiple bidders need not be considered. That would be a coalitional solution concept, and would give rise to interesting additional questions. For example, it would put additional constraints on how the seller can threaten potential non-participants. She could not threaten bidder A with giving the item to A's arch rival and threaten bidder B with giving the item to B's arch rival if those arch rivals are not the same party.

⁴A different situation arises if agent i has asymmetric spite to-

For the purposes of this paper, we are agnostic about how the seller is constrained. We simply take the function g as given, and assume that it is common knowledge. The seller always has the option of keeping the item, which sets $g(t_{-i}) = 0$ for all i and t_{-i} . So, without loss of generality, we have $g(t_{-i}) \geq 0$ if $\alpha \geq 0$ and $g(t_{-i}) \leq 0$ if $\alpha < 0$.

Finally, the auction has a resource feasibility constraint.

Definition 3 Resource feasibility (RF)

$\sum_i p_i(t) \leq 1$ and $0 \leq p_i(t) \leq 1$ for all t and i .

For the setting where there are k identical unit for sale and each bidder demands only one unit, the first part of the RF constraint is changed to $\sum_i p_i(t) \leq k$.

The seller's problem is to maximize her expected payment, subject to the three kinds of constraints above.

Definition 4 Seller's Problem

$$\max_{p,x} R = \int_{t \in T} \sum_i x_i(t) dF(t),$$

subject to the IC, IR constraints as well as the RF constraint.

Analysis: constraints simplification lemma

We first define the *ex interim* allocation rule, that is, the conditional probability that agent i gets the item, in expectation over all possible other agents' types, given his own type:

$$Q_i(t_i) = \int_{t_{-i}} p_i(t) dF(t_{-i})$$

Lemma 1 *The IC and IR constraints hold if and only if the following conditions are met for all i .*

1. *The ex interim probability that an agent gets the item is weakly increasing in his realized valuation. Formally, $Q_i(t_i)$ is weakly increasing in t_i .*
2. *The ex interim probability is the derivative of the ex interim utility function. Formally, $U_i(t_i) = U_i(a_i) + \int_{a_i}^{t_i} Q_i(s_i) ds_i$.*
3. *A bidder that happens to have his lowest possible type is incentivized to participate. Formally, $U_i(a_i) \geq -\alpha \int_{t_{-i}} g(t_{-i}) dF(t_{-i})$.*

Proof: The proof is similar—but not identical since the utility functions now include spite—to that of Lemma 2 of Myerson (1981). We omit this proof due to limited space. ■

Analysis: objective simplification lemma

In this section, we write the objective function as a function of the allocation rule and the utility of a bidder that happens to have his lowest type. We will later use this formula to derive the optimal auction.

Recall that our spiteful utility functions are as follows:

ward other agents. Say he dislikes A more than he dislikes B. Now the seller can threaten that if i does not participate, she will hurt i infinitely much by having B pay an infinite amount to A. Because the seller can make infinitely painful threats, she can extract an infinite amount of revenue in an IR auction.

$$u_i(t) = t_i p_i(t) - x_i(t) - \alpha \sum_{j \neq i} (t_j p_j(t) - x_j(t)). \quad (1)$$

For both sides of Equation (1), integrate over all possible joint types and sum over all agents. The left hand side is

$$\begin{aligned} \sum_i \int_t u_i(t) dF(t) &= \sum_i \int_{a_i}^{b_i} U_i(t_i) dF(t_i) \\ &= \sum_i \int_{a_i}^{b_i} (U_i(a_i) + \int_{a_i}^{t_i} Q_i(s_i) ds_i) dF(t_i) \\ &= \sum_i U_i(a_i) + \sum_i \int_t (1 - F_i(t_i)) p_i(t) f_{t_{-i}(t_{-i})} dt \end{aligned}$$

The right hand side is

$$\begin{aligned} \sum_i \left(\int_t (t_i p_i(t) - \alpha \sum_{j \neq i} t_j p_j(t)) dF(t) - \int_t (x_i(t) - \alpha \sum_{j \neq i} x_j(t)) dF(t) \right) \\ = \int_t \sum_i (t_i p_i(t) - \alpha \sum_{j \neq i} t_j p_j(t)) - \sum_i (x_i(t) - \alpha \sum_{j \neq i} x_j(t)) dF(t) \\ = \int_t \sum_i (1 - \alpha(n-1)) t_i p_i(t) - \sum_i (1 - \alpha(n-1)) x_i(t) dF(t) \\ = (1 - \alpha(n-1)) \left(\int_t \sum_i (t_i p_i(t) - x_i(t)) dF(t) \right) \\ = (1 - \alpha(n-1)) \left(\int_t \sum_i t_i p_i(t) dF(t) - R \right) \end{aligned}$$

For shorthand we define $C = \frac{1}{1 - \alpha(n-1)}$. In the derivation of the optimal auction, we only need to consider the case $1 - \alpha(n-1) > 0$, which implies $C > 0$. For the case $1 - \alpha(n-1) \leq 0$, the optimal auction is trivial: it is not hard to verify that the auction that charges everyone a huge amount M while keeping the item satisfies IC, IR, and RF.⁵

With C defined, the expected revenue

$$R = \int_t \sum_i \left(t_i - \frac{C(1 - F_i(t_i))}{f_i(t_i)} \right) p_i(t) dF(t) - C \sum_i U_i(a_i)$$

Via the derivation above, the seller's problem is now to maximize the formula above, subject to the resource feasibility constraint RF and the constraints of Lemma 1:

Lemma 2 *The optimal auction problem is equivalent to*

$$\begin{aligned} \max_{p,x} R = \int_t \sum_i \left(t_i - \frac{C(1 - F_i(t_i))}{f_i(t_i)} \right) p_i(t) dF(t) \\ - C \sum_i U_i(a_i) \end{aligned} \quad (2)$$

subject to RF and Constraints 1-3 of Lemma 1.

⁵The IC and RF properties are straightforward. Next we verify IR. A bidder's utility reduces to $-(1 - \alpha(n-1))M$, which is 0 when $1 - \alpha(n-1) = 0$ and very large when $1 - \alpha(n-1) < 0$. The latter case is obviously IR, but the former case warrants a bit more discussion because in the setting with spite, IR is not necessarily satisfied by having utility greater than 0, as discussed above. In the case where the bidder's utility is 0, we have $\alpha = \frac{1}{n-1} > 0$; so IR is satisfied because $-\alpha \int_{t_{-i}} g(t_{-i}) dF(t_{-i}) \leq 0$.

When $\alpha = 0$ or $n = 1$, that is, in the two cases without spite, we have $C = 1$. In those cases, the program above reduces to Myerson's solution of the optimal single item auction. In general, however, we have $C \neq 1$ and the problem is different than Myerson's.

Lemma 2 immediately implies a version of the revenue equivalence theorem for our setting:

Theorem 1 (Revenue equivalence for spiteful bidders)

The seller's expected revenue from any incentive compatible and individually rational auction mechanism is completely determined by the allocation function p and the numbers $U_i(a_i)$, i.e., each bidder's utility if he happens to have his lowest possible valuation. In particular:

- If two IC auctions have the same allocation rule and same $U_i(a_i)$ for all i , then they yield the same expected revenue.
- If two auctions have the same allocation rule and one auction yields greater $-C \sum_i U_i(a_i)$ than the other, then the former yields higher expected revenue.

When bidders' valuations are symmetrically distributed ($F_i = F_j$ for all i, j), the IC constraints can be replaced by the existence of a symmetric, increasing equilibrium and the same proof and conclusion hold. We will use this fact and the second bullet of Theorem 1 later to compare revenue of different auctions, including ones that do not incentivize bidders to report their valuations truthfully.

Analysis: the optimal auction

We are now ready to derive the optimal auction for spiteful bidders. We first set the payment rule that minimizes the utilities of the lowest types (i.e, the second term of Equation (2)) and then choose an allocation rule that maximizes the first term of Equation (2).

The payment rule

The payment rule is more complicated than Myerson's.

Consider Equation (2). The first term does not depend on the payment rule x but only the allocation rule. In the following, we find a payment rule where Constraint 3 in Lemma 1 binds, i.e, for all i , $U_i(a_i) = 0$, which is the best possible value for the second term in Equation (2).

By Constraints 2 and 3 in Lemma 1, we have

$$\begin{aligned} U_i(a_i) &= U_i(t_i) - \int_{a_i}^{t_i} Q_i(s_i) ds_i \\ &= \int_{t_{-i}} (t_i p_i(t) - x_i(t) - \alpha \sum_{j \neq i} (t_j p_j(t) - x_j(t)) \\ &\quad - \int_{a_i}^{t_i} p_i(s_i, t_{-i}) ds_i) dF_{-i}(t_{-i}) \geq -\alpha \int_{t_{-i}} g(t_{-i}) dF(t_{-i}). \end{aligned}$$

Thus, to minimize $U_i(a_i)$, we simply let

$$\begin{aligned} x_i(t) &= t_i p_i(t) - \alpha \sum_{j \neq i} (t_j p_j(t) - x_j(t)) \\ &\quad - \int_{a_i}^{t_i} p_i(s_i, t_{-i}) ds_i + \alpha g(t_{-i}) \end{aligned} \quad (3)$$

Considering Equation (3) for all i , we have n linear equations of x_i 's. Denote $P_i(t) = \int_{a_i}^{t_i} p_i(s_i, t_{-i}) ds_i$ and $D_i(t_{-i}) = \frac{1}{1+\alpha} (\alpha g(t_{-i}) + C \alpha^2 \sum_i g(t_{-i}))$, solving this equation system we obtain

$$x_i(t) = t_i p_i(t) - \frac{1}{1+\alpha} (P_i(t) + C \alpha \sum_i P_i(t)) + D_i(t_{-i}) \quad (4)$$

Again, when $\alpha = 0$ or $n = 1$, i.e., the two cases without spite, and $g(t_{-i}) = 0$, Equation 4 reduces to Myerson's payment rule $x_i(t) = t_i p_i(t) - P_i(t)$.

To sum up, if we choose x according to Equation 4, we have satisfied Constraints 2 and 3 in Lemma 1 and the lowest type of each bidder always gets zero utility. Therefore, when choosing an allocation rule, the remaining constraints to be satisfied are Constraint 1 from Lemma 1 and Constraint RF.

The allocation rule

As is common practice, we consider the regular case. For the irregular case, we can follow Myerson's *ironing technique* to obtain the corresponding solution. In the setting with spite, the regularity condition will be as follows:

Definition 5 (Regularity condition with spite)

A problem is regular if $\tilde{t}_i(t_i)$ is strictly increasing in t_i .

$$\tilde{t}_i(t_i) = t_i - \frac{C(1 - F_i(t_i))}{f_i(t_i)}$$

We call $\tilde{t}_i(t_i)$ the *virtual value* of agent i when his true valuation is t_i . When C is positive, the regularity condition above is implied by a stronger condition called *monotone hazard rate*, which state that the so-called *hazard rate function* $\frac{f_i(t_i)}{1-F_i(t_i)}$ is monotonically increasing in t_i . Commonly seen distributions, such as the normal, uniform and exponential distributions, all satisfy this condition.

If we could ignore Constraint 1, we could simply choose the allocation rule that maximizes the virtual welfare of all the agents (the seller has a virtual valuation of 0). That is,

$$p = \operatorname{argmax}_p \sum_i \max\{\tilde{t}_i(t_i), 0\} p'_i(t). \quad (5)$$

One immediate corollary is that the seller can never benefit from randomization. That is, p_i can be in $\{0, 1\}$ without reducing revenue. Also, in the symmetric setting where all the F_i 's are the same across bidders, the bidders all have the same virtual value function. Therefore, the winner in the symmetric case is the bidder with the largest actual valuation (or no one if each bidder has a negative virtual valuation).⁶

We now show that, Constraint 1 is indeed satisfied, thereby completing the proof that all the constraints are satisfied. If $t_i < t'_i$, then by regularity we have $\tilde{t}_i < \tilde{t}'_i$, which implies that a previous winner (i.e, $p_i(t_i) = 1$) under the report of t_i will still be a winner (i.e, $p_i(t'_i) = 1$). This observation then implies that $Q_i(t_i) \leq Q_i(t'_i)$, for all i , which is exactly Constraint 1.

⁶If there are k identical units for sale and each bidder demands at most one unit, according to Equation 5 an agent gets a unit if and only if he has a positive virtual valuation that is among the k largest ones.

The optimal auction: the explicit payment rule

We finish the derivation of the optimal auction by substituting the allocation rule (derived in the previous subsection) into the version of the payment rule that depended on the allocation rule (Equation 4 derived in the subsection before last). This will yield an explicit payment rule.

There are three cases to consider:

- **Case I: Bidder i wins the item.** Here, $p_i(t) = 1$ and $p_j(t) = 0$ for all $j \neq i$, so Equation 4 becomes

$$x_i(t) = t_i - \frac{1 + C\alpha}{1 + \alpha}(t_i - t_i^0) + D_i(t_{-i}), \quad (6)$$

where t_i^0 is the lowest type that i could have reported in order to win. Here, $x_i(t)$ is a positive number that does depends on i 's own bid, unlike in Myerson's auction!

For the sanity check cases $\alpha = 0$ or $n = 1$ (with $D_i(t_{-i}) = 0$ for both), Equation (6) reduces to $x_i(t) = t_{i_0}$. This is Myerson's payment rule: the winner pays t_i^0 .

- **Case II: Seller keeps the item.** When no bidder wins ($p_i(t) = 0$ for all i), it follows immediately from Equation 4 that

$$x_i(t) = D_i(t_{-i}). \quad (7)$$

Clearly, $D_i(t_{-i})$ is non-negative when $\alpha \geq 0$. ($D_i(t_{-i}) = 0$ if and only if $\alpha = 0$ or $g(t_{-i}) = 0$ for all i and t .) So, bidders usually pay positive amounts. This is similar to the conclusion drawn in different models of negative externalities (Jehiel, Moldovanu, and Stacchetti 1996; Deng and Pekeč 2011) that the seller should charge the bidders even if no bidder gets the item.

- **Case III: Bidder i loses and some other bidder wins.** Here, $p_i(t) = 0$ and $p_j(t) = 1$ for some $j \neq i$, so Equation (4) reduces to

$$x_i(t) = -\frac{C\alpha}{1 + \alpha}(t_j - t_j^0) + D_i(t_{-i}) \quad (8)$$

where t_j^0 is the lowest type that j could have reported in order to win. $D_i(t_{-i}) \geq 0$ when $\alpha \geq 0$. Here $x_i(t)$ can be positive or negative.

An interesting special case is $D_i(t_{-i}) = 0$. Then Equation 8 is negative for all $\alpha > 0$.⁷ In other words, when a bidder i loses to another bidder, i has to be subsidized by the seller. One extreme case is where i has his lowest possible valuation ($t_i = a_i$), so i loses almost surely. In that case he would get a subsidy almost surely.

To summarize, we have the following theorem.

Theorem 2 *The optimal auction uses (5) as the allocation rule. If bidder i wins, his payment is determined by (6). If i loses to another bidder, his payment is determined by (8). If the seller keeps the item, the bidders' payments are determined by (7).*

⁷In very special case it can also be zero, namely $t_j - t_j^0 = 0$ (i.e., the winning bidder j happens to have a valuation that exactly equals his lowest report needed to win).

The payments in the k -unit unit-demand case can also be simplified from Equation (4) by substituting in the explicit allocation rule (agent i gets a unit ($p_i(t) = 1$) if and only if $\tilde{t}_i(t_i)$ is among the k largest positive virtual values). The calculation, which is similar to the 1-unit case, is omitted due to limited space.

Revenue comparison of different auctions

In prior work, symmetric equilibria have been derived for first- and second-price auctions with spite (Morgan, Steiglitz, and Reis 2003; Brandt, Sandholm, and Shoham 2007; Sharma and Sandholm 2010). All three of those papers restrict attention to deterministic mechanisms ($p_i \in \{0, 1\}$), while we relax this to capture randomized ones as well. So our setting includes Morgan et al.'s and Brandt et al.'s settings as special cases. Morgan et al. calculated symmetric equilibria for first- and second-price auctions, and apply them to conclude that when $0 < \alpha < 1$, the second-price auction yields more expected revenue than the first-price auction. Brandt et al. extended those results to a slightly more general setting.⁸

In this section we will show how the theory derived in this paper can be used to prove those results (and generalizations thereof) in a simpler way, and to compare expected revenues of other auction mechanisms under spite. For two-bidder auctions where the bidders' valuations are drawn independently uniformly on $[0, 1]$, Morgan et al. showed that in the first-price auction, the symmetrical equilibrium strategy is $b_{FP}(t_i) = \frac{1+\alpha}{2+\alpha}t_i$, while in the second-price auction, the symmetric equilibrium strategy is $b_{SP}(t_i) = \frac{1+\alpha}{1+2\alpha}t_i + \frac{\alpha}{1+2\alpha}$.

Therefore, the utility of a bidder that has his lowest possible type in the first-price auction is

$$U_i(a_i) = \int_0^1 -\alpha \left(s_j - \frac{1+\alpha}{2+\alpha} s_j \right) ds_j = -\frac{\alpha}{2(2+\alpha)} \quad (9)$$

and in the second-price auction it is

$$U_i(a_i) = \int_0^1 -\alpha \left(s_j - \frac{\alpha}{1+2\alpha} \right) ds_j = -\frac{\alpha}{2(1+2\alpha)} \quad (10)$$

It is straightforward to verify that $U_i(a_i)$ is no less than i 's utility when he does not participate, so both auctions satisfy the IR constraints. The value of Equation (10) is strictly less than that of Equation (9). When $0 < \alpha < 1$, we have $C = \frac{1}{1-\alpha} > 0$, so the term $-C \sum_i U_i(a_i)$ is strictly larger in the second-price auction. Since both auction have a symmetric, increasing equilibrium, they have the same allocation rule in this symmetric setting (i.e., efficient allocation). Therefore, by Theorem 1, the second-price auction yields higher expected revenue than the first-price auction.

With this methodology, we can also generalize the prior results to the following settings of spite.

- When $\alpha = 0$, both auctions yield equal expected revenue.

⁸Sharma and Sandholm (2010) derived equilibria for the setting with asymmetric spite factors.

- When $\alpha > 1$, we have $C < 0$, so Equation (10) is strictly greater than Equation (9). We still have that $-C \sum_i U_i(a_i)$ is strictly greater in the second-price auction. Therefore, we can conclude that the second-price auction yields higher revenue than the first-price auction.
- When $-0.5 < \alpha < 0$ or $\alpha < -2$, the first-price auction yields higher revenue.
- When $-2 \leq \alpha \leq -0.5$, either b_{FP} or b_{SP} is non-increasing in t_i . In this case, neither Morgan et al.'s equilibrium analysis nor our Theorem 1 apply.

To sum up this section, we have the following theorem.

Theorem 3 *Consider two-bidder auctions where the bidders' valuations are drawn independently uniformly on $[0, 1]$. When the spite factor $\alpha \geq 0$, the second-price auction yields no less expected revenue than the first-price auction. When $\alpha < 0$, the first-price auction yields no less expected revenue than the second-price auction, as long as an increasing symmetric equilibrium exists for both auctions.*

Furthermore, for general valuation distributions, revenues of any two auctions with the same allocation rule can be compared as follows using Theorem 1.

1. Compute a symmetric and increasing equilibrium of each auction.
2. Prove that they have the same allocation rule.
3. Compute $-C \sum_i U_i(a_i)$ for both auctions. The larger value implies larger expected revenue.

Conclusions and future research

We considered the optimal auction design problem among spiteful (or altruistic) bidders. In this setting, individual rationality constraints are non-trivial and depend on the seller's ability to hurt non-participants—who have spite toward the participants that will enjoy the allocation.

The optimal auction is a generalization of Myerson's (1981) auction: it chooses an allocation that maximizes agents' virtual valuations, but for a generalized definition of virtual valuation. The payment rule is much less obvious. For one, it takes each bidder's own report into consideration when determining his payment. Moreover, bidders pay even if the seller keeps the item; a similar phenomenon has been shown in other settings with negative externalities (Jehiel, Moldovanu, and Stacchetti 1996; Deng and Pekeč 2011). On the other hand, a novel aspect of our auction is that it sometimes subsidizes losers when the item is sold to some other bidder.

We also derived a revenue equivalence theorem for this setting. Using it, we came up with a short proof of (a slight generalization of) the previously known result that, in two-bidder settings with independently uniformly drawn valuations, second-price auctions yield greater expected revenue than first-price auctions. Finally, we presented a template for comparing the expected revenues of any two auction mechanisms that have the same allocation rule (for the valuations distributions at hand).

There are several interesting directions for future research. For one, it is of both practical and theoretical interests to compare the revenue of the VCG and GSP auctions

for sponsored search in the setting with spite. Furthermore, it would be interesting to study settings where a bidder's externality is determined by others' payments only (i.e., independence of others' valuations), but unlike in Lu (2012), with different threats than auction cancellation.

References

- Armstrong, M. 2000. Optimal multi-object auctions. *Review of Economic Studies* 67(3):455–81.
- Avery, C. N., and Hendershott, T. 2000. Bundling and optimal auctions of multiple products. *Review of Economic Studies* 67(3):483–97.
- Brandt, F.; Sandholm, T.; and Shoham, Y. 2007. Spiteful bidding in sealed-bid auctions. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 1207–1214.
- Conitzer, V., and Sandholm, T. 2011. Expressive markets for donating to charities. *Artificial Intelligence* 175:1251–1271.
- Conitzer, V., and Sandholm, T. 2012. Computing optimal outcomes under an expressive representation of settings with externalities. *Journal of Computer and System Sciences* 78:2–14. Special issue on Knowledge Representation and Reasoning. Early version in AAAI-05.
- Constantin, F.; Rao, M.; Huang, C.-C.; and Parkes, D. C. 2011. On Expressing Value Externalities in Position Auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*.
- Deng, C., and Pekeč, S. 2011. Money for nothing: exploiting negative externalities. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, 361–370.
- Figuerola, N., and Skreta, V. 2009. The role of optimal threats in auction design. *Journal of Economic Theory* 144:884–897.
- Iyengar, G., and Kumar, A. 2006. Characterizing optimal keyword auctions. In *Proceedings of the Workshop on Sponsored Search Auctions*.
- Iyengar, G., and Kumar, A. 2008. Optimal procurement mechanisms for divisible goods with capacitated suppliers. *Review of Economic Design* 12:129–154.
- Jehiel, P.; Moldovanu, B.; and Stacchetti, E. 1996. How (not) to sell nuclear weapons. *American Economic Review* 86(4):814–29.
- Krysta, P.; Michalak, T.; Sandholm, T.; and Wooldridge, M. 2010. Combinatorial auctions with externalities. In *Proceedings of International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.
- Ledyard, J. O. 2007. Optimal combinatoric auctions with single-minded bidders. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, 237–242.
- Levin, J. 1997. An optimal auction for complements. *Games and Economic Behavior* 18(2):176–192.
- Lu, J. 2012. Optimal auctions with asymmetric financial externalities. *Games and Economic Behavior* 74:561–575.
- Maskin, E., and Riley, J. 1989. Optimal multi-unit auctions. In *The Economics of Missing Markets, Information, and Games*. chapter 14, 312–335.
- Morgan, J.; Steiglitz, K.; and Reis, G. 2003. The spite motive and equilibrium behavior in auctions. *The B.E. Journal of Economic Analysis & Policy* 0(1):5.
- Myerson, R. B. 1981. Optimal auction design. *Mathematics of Operations Research* 6(1):58–73.
- Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V. V. 2007. *Algorithmic Game Theory*. Cambridge University Press.
- Pai, M. M., and Vohra, R. 2008. Optimal auctions with financially constrained bidders. Discussion Papers 1471, Northwestern University.
- Parkes, D. C., and Sandholm, T. 2005. Optimize-and-dispatch architecture for expressive ad auctions. In *Proceedings of the Workshop on Sponsored Search Auctions*.
- Sharma, A., and Sandholm, T. 2010. Asymmetric spite in auctions. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 867–873.
- Zhou, Y., and Lukose, R. 2007. Vindictive bidding in keyword auctions. In *Proceedings of the International Conference on Electronic Commerce (ICEC)*, 141–146.