Mechanism Design with Money

Shikha Singh

Contents

1	Game Theory Background			
	1.1 Solution Concepts	3		
2	Single-Item Auction	6		
	2.1 Second-Price (Vickrey) Auction	6		
	2.2 First-Price Auction			
	2.3 Revenue Equivalence	9		
3	Myerson's Lemma for Single-Parameter Domains	10		
	3.1 Proof of Myerson's Lemma			
	3.2 Application of Myerson's Lemma	14		
	3.3 Characterizing Bayes Nash Equilibrium	15		
4	General Valuation: VCG Mechanism			
5	Case Study: Sponsored Search Auctions	20		
	5.1 Nash Equilibrium of GSP	21		
	5.2 Envy-free Outcome	22		
	5.3 Equivalence of GSP and VCG	25		

1 Game Theory Background

At a high level, a game consists of players, the actions available to the players and the utilities the players get from each outcome (a sequence of action played). A game representation depends on whether the game is a *simultaneous-move* game (all players act at once) or a *sequential* game (players act over rounds). For this part of the course, we will only consider simulateneous move games. Such games are represented in **normal form** (with utilities described in a matrix).

Definition 1. A finite, normal-form game (N, u, A), is defined by the following:

- Player set $N = \{1, 2, ..., n\}$
- Action set A_i that are a set of actions available to player i, for each $i \in N$
- $A = A_1 \times \cdots \times A_n$ are the set of **action profiles** of all *n* players and are also referred to as the **outcomes** of the game
- Utility function $u_i: A \to \mathbb{R}$ maps each action profile that can be played to a payoff

In game theory, we make two assumptions:

- Rationality. Each player's goal is to maximize their own utility.
- *Rationality is Common-Knowledge.* Each player knows that everyone else is rational and that everyone else knows that they they are rational and that they know that everyone else is rational and so on, infinitely.

Example. (Prisoner's Dilemma) A classic example of a game where two alleged criminals are questioned in separate rooms and each player has two actions:

- Cooperate (C): stay silent and not admit to anything
- Defect (D): testify against the other person

If both cooperate (C, C), then each serves 1 year in prison for minor offense. If one cooperates and other defects (C, D) or (D, C), then the confessor goes free while other person gets a long prison sentence. If both defect (D, D), then each serve 3 years in prison. Normalizing the utilities, we can represent this game in normal-form with a row-player and column-player as follows.

	C	D	
C	4,4	0,5	
D	5.0	2.2	

A global guarantee of an outcome of a game that is often desirable is called **Pareto-optimality**.

Definition 2. Consider a finite game with outcomes *O*.

- An outcome $o \in O$ **Pareto-dominates** an outcome $o' \in O$, if $u_i(o) \ge u_i(o')$ for all players i and $u_j(o) > u_j(o')$ for a specific player j.
- An outcome *o* is **Pareto-optimal**, if no other outcome Pareto-dominates it.

In Prisoner's Dilemma (Example), the outcome (D, D) is Pareto-dominated by (C, C) and the other three outcomes are Pareto-optimal.

1.1 Solution Concepts

When rational players act to maximize their utility in a game, the process of determining what outcome will likely occur is called *solving the game*. Different "solution concepts" or "equilibrium concepts" are used in game theory to justify likely behavior. In this section, we will review two most popular solution concepts.

Dominant-strategy equilibrium. The strongest possible guarantee on player behavior is given by the existence of dominant-strategy equilibrium (DSE). Intuitively, in a DSE, each player can maximize their utility by playing a dominant action that is agnostic to the actions played by others.

An important notational shorthand in game theory is $A_{-i} = A_1 \times \cdots \times A_{i-1} \times A_{i+1} \cdots \times A_n$ which is the set of all possible action profiles excluding player *i*. Similarly, $a_{-i} = (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n) \in A_{-i}$ represents a specific action profile of all players except *i*.

Definition 3. An action profile a_1^*, \ldots, a_n^* is a **dominant-strategy equilibrium** (DSE) if and only if for all players *i*:

 $u_i(a_i^*, a_{-i}) \ge u_i(a_i', a_{-i})$ for all $a_i' \in A_i$ and for all $a_{-i} \in A_{-i}$

In Prisoner's Dilemma (Example), the unique DSE is (D, D) which is the only non-Pareto optimal outcome, explaining the *dilemma*—what is good for individuals sometimes is not good for the group.

A DSE is a very strong gaurantee but may not always exist. The second-best solution concept a game can admit is called a pure-Nash equilibrium.

Definition 4. An action profile a_1^*, \ldots, a_n^* is a **pure Nash equilibrium** (PNE) if and only if for all players *i*:

$$u_i(a_i^*, a_{-i}) \ge u_i(a_i', a_{-i})$$
 for all $a_i' \in A_i$

The action a_i^* is called the **best response** of player *i* to the action profile a_{-i} of others.

Thus, in a pure Nash equilibrium, each player plays a best response to others and no player has any incentive to deviate unilaterally to improve their utility.

The Nash equilibrium concept has several challenges:

- A pure Nash equilibrium may not always exists in a game.
- Many pure Nash equilibrium might exist (causes the problem of equilibrium selection—which equilibrium should players choose?).
- Computing a Nash equilibrium is a computationally difficult problem.¹

Incomplete-Information Games and Bayes Nash. The definition of Nash equilibrium requires each player to know a_{-i} , the actions other players are playing. The assumption that rationality is common knowledge requires each player to also know the utility structure of others. The games where players have full information of the payoff structure of others are called *completeinformation games*. In an incomplete-information game, each player *i* may have a private type t_i that affects their utility. The concept of Nash equilibrium can be extended to such games by assuming that the types t_i are drawn identically and independently (i.i.d.) from a distribution *G* that is common knowledge.

A strategy s_i of a player *i* now is a function that maps their private type t_i to an action and their goal is to maximize their **expected utility** for a strategy profile $s = (s_1, \ldots, s_n)$, defined as:

$$\mathbb{E}[u_i(s)] = \sum_{t_{-i}} u_i(s|t_{-i}) \cdot \Pr(t_{-i})$$

Definition 5. Consider a Bayesian game where each player *i* has a private types t_i that is drawn i.i.d. from a distribution *G*. Then the strategy profile s_1^*, \ldots, s_n^* is a **pure Bayes** Nash equilibrium if and only if all players *i* maximize their interim expected utility, that

¹It is PPAD-complete, a notion of hardness we will discuss later.

 $\qquad \text{is},$

$$\mathbb{E}[u_i(s_i, s_{-i})] \ge \mathbb{E}[u_i(s'_i, s_{-i})] \text{ for all } s'_i$$

where the expectation is over the private types t_{-i} and each player knows their own type t_i .

2 Single-Item Auction

The first market we consider is a single-item market with n buyers $\{1, \ldots, N\}$. Each buyer has a private valuation $v_i \in \mathbb{R}$ for obtaining the item.

In these markets, an **allocation rule** $x = (x_1, \ldots, x_n)$ defines who gets what and a **payment rule** $p = (p_1, \ldots, p_n)$ defines who pays what. For a single item auction, if buyer j gets the item then $x_j = 1$ and $x_i = p_i = 0$ for all $i \neq j$. The utility of a bidder i is defined as $u_i = x_i \cdot v_i - p_i$. That is, their utility is zero if they do not get allocated (and pay nothing) and otherwise it is their value minus their price.

To define an auction for a single-item auction, we need to determine the winner j should be and what they payment p_j they should be charged. The global optimization objective is to allocate the item so as to maximize **social welfare**, defined as, $\sum_{i=1}^{n} x_i v_i$.

For a single-item setting, this means the goal is to allocate to the highest-valued buyer. However, buyer valuations are private so the auction needs to elicit these from the buyers in the form of *bids*. A bid b_i of a buyer (now called *bidder*) *i* is their alleged value v_i . In a **truthful** bid profile, $b_i = v_i$ for all *i*.

We will consider **sealed-bid** auctions, where bidders submit their bids privately to the seller in the beginning of auction.

2.1 Second-Price (Vickrey) Auction

A second price auction is defined as follows:

- Each bidder *i* submits their private bid b_i , for $i \in \{1, \ldots, N\}$
- The item is allocated to the highest bidder $j = \operatorname{argmax}_{i \in N} b_i$.
- The bidder j is charged the payment equal to the second-highest bid $p_j = \max_{i \neq j} b_i$.

A second-price auction admits truthful bidding as its unique dominant-strategy equilibrium.

Theorem 1. Truthful bidding is a dominant-strategy equilibrium of the second price auction.

Proof. Consider an abritrary bidder *i* with valuation v_i . Fix b_{-i} of other bidders. We show that setting $b_i = v_i$ maximizes bidder *i*'s utility (among all possible bids they can submit) for all b_{-i} .

Let $B = \max_{j \neq i} b_j$ be the maximum bid among b_{-i} . There are two possible outcomes for bidder *i* with valuation v_i .

Case 1. $v_i \ge B$. The maximum possible utility the bidder can obtain in this case is by winning and paying B, that is, $v_i - B$.

Case 2. $v_i < B$. The maximize possible utility the bidder can obtain in this case is 0 because if they win they will have to pay *B* and get negative utility.

In both cases, setting $b_i = v_i$, gives bidder *i* the maximize possible utility and holds for all b_{-i} . Thus, truthful bidding is a DSE of the second-price auction.

An auction or mechanism **dominant-strategyproof** if truthful bidding is a dominant-strategy equilibrium of the resulting game.

As the item is allocated to the highest-valued bidder and truthful bidding is a DSE, the secondprice auction maximizes social welfare at DSE.

2.2 First-Price Auction

A first-price auction is defined as follows:

- Each bidder *i* submits their private bid b_i , for $i \in \{1, \ldots, N\}$
- The item is allocated to the highest bidder $j = \operatorname{argmax}_{i \in N} b_i$.
- The bidder j is charged their bid b_j .

A first-price auction is not dominant strategy proof. To see this, consider the bidder with the highest valuation and assume every one else bids their value. If the bidder bids $b_i = v_i$, then their utility is zero. However, if they bid $b_i = \max_{j \neq i} v_j$, their utility is $v_i - \max_{j \neq i} v_j$ which is strictly positive.

We analyze bidder behavior in a first-price auction in a restricted setting. In particular, assume that each bidders value v_i is drawn i.i.d. from a uniform distribution on the interval [0, 1]. To derive the Bayes-Nash equilibrium strategy, we first make the assumption that each bidder sets their bid proportional to their value, that is, $b_i = \alpha(n) \cdot v_i$. Later, we verify that this assumption is in fact correct.

Deriving BNE of First-Price Auction. Consider bidder 1 and assume that $b_i = \alpha v_i$ for all other bidders 2,..., n. Let $\mathbf{E}(u_1)$ denote the expected utility of bidder 1, assuming all bidder values are drawn i.i.d from the uniform distribution on [0, 1].

=

$$\mathbf{E}(u_1) = (v_1 - b_1) \cdot \Pr(1 \text{ wins with bid } b_1) + 0 \cdot \Pr(1 \text{ loses with bid } b_1)$$
(1)

$$= (v_1 - b_1) \cdot \Pr[b_1 \ge \max_{i=2}^n b_i]$$
(2)

$$= (v_1 - b_1) \cdot \Pr(b_1 \ge b_2 \cap \dots \cap b_1 \ge b_n)$$
(3)

$$= (v_1 - b_1) \cdot \Pr(b_1 \ge \alpha v_2 \cap \dots \cap b_1 \ge \alpha v_n)$$
(4)

$$= (v_1 - b_1) \cdot \Pr\left(v_2 \le \frac{b_1}{\alpha}\right) \cdots \Pr\left(v_n \le \frac{b_1}{\alpha}\right)$$
(5)

$$= (v_1 - b_1) \cdot \left(\frac{b_1}{\alpha}\right)^{n-1} \tag{6}$$

(7)

Step follows from the fact that $b_i = \alpha v_i$ for i = 2, ..., n by our assumption and that the values are drawn independently. Step 5 follows from the fact that $\Pr(x \le k) = k$ when x is drawn from a uniform distribution on [0, 1].

To find bid b_1 that maximizes $\mathbf{E}(u_1)$, we take the derivative wrt to b_1 and set it to zero.

$$\mathbf{E}'(u_1) = v_1 \cdot (n-1)\frac{b^{n-2}}{\alpha^{n-1}} - n \cdot \frac{b^{n-1}}{\alpha^{n-1}} = 0$$
(8)

(9)

Solving the above gives us that $b_1 = \frac{n-1}{n} \cdot v_1$.

Theorem 2. Assume each of the *n* bidders have values drawn i.i.d. from uniform distribution on [0, 1]. Then, the strategy $s_i = \frac{n-1}{n} \cdot v_i$ for each bidder *i* is a symmetric Bayes Nash equilibrium of the sealed-bid first-price auction.

Proof. This proof is analogous to the derivation with the difference that we take an arbitrary bidder j and fix $b_j = \frac{n-1}{n}v_j$ for all bidders $i \neq j$. Considering the expected utility of j as before and differentiating it to verify that the best response is $b_i = \frac{n-1}{n}v_i$ finishes the proof.

We conclude with the following properties of first-price auction.

- The above analysis generalizes to arbitrary i.i.d. distributions (beyond uniform).
- The Bayes Nash equilibrium (BNE) of the first-price auction is unique.

- At the unique BNE, the first-price auction maximizes the social welfare (this is because the highest valued bidder wins the auction).
- At the unique BNE, the first-price auction generates the same revenue as the second-price auction; see Theorem 3.

2.3 Revenue Equivalence

In general, revenue equivalence in auction theory states that if two auctions have the same allocation at equilibrium and same bidder valuation distribution, then they also generate the same revenue.

We show this for the special case of first and second-price auction for the single-item case. We first define the kth order-statistic and use it in the proof.

Definition 6. Let X_1, \ldots, X_n be *n* independent samples drawn identically from the uniform distribution on [a, b]. Let $X_{(k)}$ denote the *k*th highest value among the *n* samples, called the *k*th-order statistic. Then,

$$\mathbb{E}[X_{(k)}] = a + \frac{n - (k - 1)}{n + 1} \cdot (b - a)$$

Theorem 3. If bidder's values are drawn i.i.d. from the uniform distribution on [0, 1], then the expected revenue (at equilibrium) of the first-price auction is equal to the expected revenue of the second-price auction (at equilibrium).

Proof. Relabel the bidder indices such that $v_1 \ge v_2 \ge \dots v_n$. Let $\mathbb{E}[R_1]$ and $\mathbb{E}[R_2]$ be the expected revenue of the first and second price auctions respectively. Then,

$$\mathbb{E}[R_2] = \mathbb{E}[v_2] = \frac{n-1}{n+1}$$

Using Theorem 2 and linearity of expectation, the expected revenue of the first-price auction is

$$\mathbb{E}[R_1] = \mathbb{E}[b_1] = \mathbb{E}\left[\frac{n-1}{n}v_1\right] = \frac{n-1}{n}\mathbb{E}[v_1] = \frac{n-1}{n} \cdot \frac{n}{n+1} = \frac{n-1}{n+1}.$$

3 Myerson's Lemma for Single-Parameter Domains

A generalization of a single-item market is a market with multiple items where each bidder's valuation v_i for their allocation is captured by a single number. Some examples of this setting:

- k identical items, with each bidder's v_i is their value for obtaining a single copy of the item
- each bidder has a value v_i per "unit" of allocation, e.g. v_i -per-click in a sponsored-search auction, and their allocation determines how many units they receive
- 0-1 allocations with constraints: each bidder either "wins" or "loses" their desired subset, or each bidder is either included or not included in a Knapsack auction, etc.

For single-parameter settings, Myerson's lemma provides a tight characterization of allocation and payment rules that enforce dominant-strategyproof behavior.

Definition 7. An allocation rule $x = (x_1, \ldots, x_n)$ for a single-parameter domain is monotonenon-decreasing if for every bidder *i* and bids b_{-i} of other bidders, *i*'s allocation $x_i(z, b_{-i})$ is non-decreasing in *i*'s bid *z*.

Theorem 4. (Myerson) Fix a single-parameter setting. We state the result for the continuous case.

- An allocation rule x can be made dominant strategyproof if and only if x is monotone (non decreasing).
- If x is monotone, there is a unique payment rule p such that (x, p) is dominant strategyproof. This payment rule is given by the following expression for all i

$$p_i(z, b_{-i}) = z \cdot x_i(z, b_{-i}) - \int_0^z x_i(z, b_{-i}) dz$$

where player *i* bids *z*. Keeping b_{-i} fixed, we can simplify:

$$p_i(z) = z \cdot x_i(z) - \int_0^z x_i(z) dz$$

This assumes that $p_i(0) = 0$.

Suppose the allocation function x is piecewise-constant and there are ℓ points at which the

allocation "jumps" before bid z, the payment at bid z given by Theorem 4 becomes:

$$p_i(z) = \sum_{i=1}^{\ell} z_i \cdot [\text{jump in } x_i \text{ at } z_i]$$

3.1 Proof of Myerson's Lemma

We break up the proof of Myerson's lemma in several parts.

Step 1. If x is dominant strategyproof then x must be monotone non-decreasing.

Consider a bidder i with value v that bids v'. For truthful bidding to be the dominant strategy for i, their utility from being truthful should be at least as high as misreporting, that is,

$$v(x(v) - p(v)) \ge v(x(v') - p(v')) \quad \text{for all } v, v' \tag{10}$$

Assuming that strategies are onto, we consider two cases. Assuming wlog that $z_1 < z_2$.

• Case 1. $v = z_1$ and bidder overbids z_2 . Inequality 10 in this case becomes:

$$z_1(x(z_1) - p(z_1)) \ge z_1(x(z_2) - p(z_2)$$
(11)

• Case 2. $v = z_2$ and bidder underbids z_1 . Inequality 10 in this case becomes:

$$z_2(x(z_2) - p(z_2)) \ge z_2(x(z_1) - p(z_1)$$
(12)

Adding both of the above iequalities and rearranding, we get:

$$(z_2 - z_1) \cdot (x(z_2) - x(z_1)) \ge 0$$

Since $z_2 > z_1$, this only holds if $x(z_2) \ge x(z_1)$ and thus x must be monotone non-decreasing.

Step 2. If x is dominant strategyproof and monotone, then there is a unique payment rule given by Theorem 4.

Using Inequalities (11) and (12), we can derive and upper and lower bound on the payment difference between bids z_1 and z_2 , that is,

$$z_2 \cdot (x(z_2) - x(z_1)) \ge p(z_2) - p(z_1) \ge z_1 \cdot (x(z_2) - x(z_1))$$



We can visualize these upper and lower bounds in the picture from Hartline [2] below.

To finish this part of the proof set $z_1 = 0$ and $z_2 = z$ and we derive the payment to be the shaded area above the allocation curve at z, or mathematically:

$$p_i(z) = z \cdot x_i(z) - \int_0^z x_i(z) dz$$

Step 3. If the allocation x is monotone and the payment rule p is as given by the expression in the lemma then, (x, p) is dominant strategyproof.



This proof is entirely by picture from Hartline [2]. The left column shows (shaded) the welfare, payment, and utility of the bidder playing action $b(v = z_2)$. The right column shows (shaded) the same for the bidder playing action $b(v^{\dagger} = z_1)$. The final diagram shows (shaded) the difference between utility for these strategies. Monotonicity implies this difference is non-negative.

3.2 Application of Myerson's Lemma

Myerson's lemma gives a general characterization of all strategy proof mechanisms (x, p) for single-parameter domains. Specific applications of it are described below.

Critical bid for 0-1 **allocations**. When allocations are 0/1, that is, for each bidder $x_i = 0$ or $x_i = 1$, then the Myerson payment simplifies to charging each bidder their **critical bid**, the lowest bid they could have submitted and still won. More formally,

$$p(b_i, \mathbf{b}_{-i}) = \begin{cases} 0 & \text{if } x_i(b_i, \mathbf{b}_{-i}) = 0\\ b_i^*(\mathbf{b}_{-i}) & \text{if } x_i(b_i, \mathbf{b}_{-i}) = 1 \end{cases}$$

where $b_i^*(\mathbf{b}_{-i})$ is the bidder *i*'s *critical bid*, that is, the lowest bid at which *i* gets a non-zero allocation.

Welfare-maximizing allocation rules. Consider an arbitrary single-parameter environment, with feasible allocation $X = (x_1, x_2, \ldots, x_n) \in \mathbb{R}^n$. Given bids $\mathbf{b} = (b_1, \ldots, b_n)$, the welfare-maximizing allocation rule is $\mathbf{x}(\mathbf{b}) = \operatorname{argmax}_{(x_1,\ldots,x_n)\in X}\sum_{i=1}^n b_i x_i$. In Assignment 2 (Problem 1), we prove that this allocation rule is monotone. Thus, such allocation rules can always be paired with the Myerson payment scheme to give a dominant-strategyproof mechanism.

Consider a welfare maximizing allocation rule $\mathbf{x}(\mathbf{b})$ and feasible allocations X that contain only 0-1 vectors—that is, each bidder either wins or loses.

Given feasible allocations containing 0-1 vectors, we can identify each feasible allocation with a "winning set" of bidders S^* (the set of bidders *i* with $x_i = 1$ in that allocation). As proved in Assignment 2 (Problem 2), the critical bid payment of the winning bidders in S^* given by Myerson then becomes equal to their **externality**, which is the difference between:

- (i) the maximum welfare of a feasible allocation that excludes i^2 —that is, the maximum welfare that can be generated if i was not present
- (ii) the welfare $\sum_{j \in S^* \setminus \{i\}} b_j$ generated by the winners (other than *i*) in the chosen outcome S^* —that is, the welfare that is generated (by others) given *i* wins

Sponsored-search auction. In the sponsored search auction problem, there are k slots, the jth slot has a click-through rate (CTR) of α_j (non-increasing in j), and the utility of bidder i in slot j is $\alpha_j(v_i - p_j)$, where v_i is the value-per-click of the bidder and p_j is the price charged per-click in slot j.

 $^{^{2}}$ You should assume that there is at least one such feasible allocation.

The following algorithm, the *Vickrey-Clarke-Groves* (VCG) auction, maximizes welfare and its payment is derived using Myerson's Lemma.

Vickrey-Clarke-Groves (VCG) auction for sponsored search.

- 1. Rank the advertisers from highest to lowest bid-per-click b_i ; assume without loss of generality that $b_1 \ge b_2 \ge \ldots \ge b_n$.
- 2. For i = 1, 2, ..., k, assign the *i*th bidder to the *i*th slot.
- 3. For i = 1, 2, ..., k, charge the *i*th bidder a price-per-click given by Myerson's formula:

$$p_i = \sum_{j=i}^k b_{j+1} \left(\frac{\alpha_j - \alpha_{j+1}}{\alpha_i} \right)$$

3.3 Characterizing Bayes Nash Equilibrium

Myerson's lemma also generalizes to Bayes Nash equilibrium. In particular, it states that a strategy profile s is a Bayes' Nash equilibrium in (x, p) if and only if for all players i,

- the allocation probability x_i is monotone non decreasing, and,
- player i's expected payment is given by

$$p_i(z) = z \cdot x_i(z) - \int_0^z x_i(z) dz$$

Revenue Equivalence. The main takeaway of Myerson's lemma characterizing the BNE in single-parameter settings is that the expected payment only depends on the allocation probability.

Corollary. (Revenue Equivalence) If two mechanisms have the same distribution of agent values and same allocation (at BNE), then they generate the same revenue.

One application of this corollary is that it lets us reason about the equilibrium of auctions that are harder to analyze by equating their expected payments at equilibrium to that of a strategically-simple dominant-strategyproof auction where bidders bid their value.

Steps to Use Revenue Equivalence for BNE. As an example of how revenue equivalence is useful to "guess" the BNE of an auction, let's use it to recreate the BNE for first-price auction.

- Step 1. Guess what the allocation might be in a Bayes-Nash equilibrium (usually a welfare-maximizing allocation). In this case, we guess that the highest-valued bidder wins in a first-price auction at equilibrium.
- Step 2. Write down the expression for the expected payment of a bidder as a function of their value in a strategically-simpler auction (usually the dominant-strategyproof version). In this case, consider bidder *i*, their expected payment in a second-price auction is

$$E[p_i] = E[\text{second highest bid} | v_i \text{ wins}] \cdot \Pr[v_i \text{ wins}] + 0 \cdot \Pr[v_i \text{ loses}]$$
$$= E[\text{second highest value} | v_i \text{ wins}] \cdot \Pr[v_i \text{ wins}]$$
$$= E[X_{(1)} \text{ in } n - 1 \text{ samples from } [0, v_i]] \cdot \Pr[v_i \text{ wins}]$$
$$= \frac{n-1}{n} v_i \cdot \Pr[v_i \text{ wins}]$$

Here the last equality is because the expected value of the second-highest bidder given v_i is the highest bidder is the same as the expected value of the first-order statistic when n-1samples are drawn from a uniform distribution on $[0, v_i]$. Let X_1, \ldots, X'_n denote n' samples drawn i.i.d. from the uniform distribution on [0, b]. Let X_k be the kth largest value among them, then the last step follows from using $E[X_{(k)}] = \frac{n'-(k-1)}{n'+1}b$ with n' = n-1 and k = 1.

Step 3. Write the expression for the expected payment in terms of the strategy of the bidder in the auction you are trying to solve for the BNE. In this case, a bidder *i* pays their bid in a first-price auction if they win and pay zero otherwise. Assuming a bidder *i* shades their bid down by a factor α of their value, their expected payment is

$$E[p_i] = E[b_i \mid v_i \text{ wins}] \cdot \Pr[v_i \text{ wins}] + 0 \cdot \Pr[v_i \text{ loses}]$$
$$= E[\alpha v_i \mid v_i \text{ wins}] \cdot \Pr[v_i \text{ wins}] = \alpha v_i \cdot \Pr[v_i \text{ wins}]$$

In the last step, $E[\alpha v_i] = \alpha v_i$ because of linearity of expectation and because bidder *i* knows their own value (the expected utility is over their uncertainty of other bidders bids and values).

• Step 4. Solve for the BNE strategy by equating the expected payments between the two auctions. In this case, setting the expected payments equal in first and second price auction, assuming that the probability of winning is the same in both, we get $\alpha = \frac{n-1}{n}$, and thus revenue equivalence suggests that $s(v_i) = \frac{n-1}{n}v_i$ should be the symmetric Bayes Nash equilibrium of the first-price auction with n bidders.

• Step 5. Finally, verify that this strategy profile computed in Step 4 is actually a symmetric BNE. We did this step in lecture for the first-price auction.

4 General Valuation: VCG Mechanism

Consider a general mechanism to allocate $M = \{1, ..., m\}$ different items. Then, there are 2^m possible subsets that a bidder can receive and they may have a different valuation for each such allocation. In such a general combinatorial auction, it can be challenging to even ask bidders to submit such valuation, let alone compute on it.

For a more restricted example, but still more general than the single-parameter setting, consider the unit-deman market.

Example. (Unit-Demand Market) Consider a market with m items and n bidders such that each bidder i has a valuation v_{ij} for item j, for $i \in \{1, \ldots, n\}$ and $j \in \{1, \ldots, m\}$.

Our goal is to find a welfare-maximizing, dominant-strategyproof mechanism for such markets. Surprisingly, such a mechanism always exists, even for these general markets.

Vickrey-Clarkes-Grove (VCG) Mechanism The VCG mechanism is defined as follows. Let A be the set of feasible outcomes.

- Collect sealed bids $\mathbf{b} = \mathbf{b_1}, \mathbf{b_2}, \dots, \mathbf{b_n}$, where each $\mathbf{b_i}$ is a vector indexed by |A| describing the bidders valuation for each possible outcome.
- Find a welfare-maximizing allocation $a^*(\mathbf{b}) = \operatorname{argmax}_{a \in A} \sum_{i=1}^{n} \mathbf{b}_i(a)$, assuming that the bids are truthful.
- Charge each bidder their **externality**, which is the welfare loss they impose on others by their presence, that is,

$$p_i(\mathbf{b}) = \max_{\substack{a_{-i} \in A_{-i} \\ \text{without } i}} \mathbf{b}_j(a_{-i}) - \sum_{\substack{j \neq i \\ \text{with } i}} \mathbf{b}_j(a^*)$$

Theorem 5. The VCG mechanism is dominant strategyproof.

Proof. Fix i and bids \mathbf{b}_{-i} . Suppose the chosen outcome is a^* . Then, the utility of i for

outcome a^* is:

$$u_i(a^*) = v_i(a^*) - p_i(\mathbf{b}) = v_i(a^*) + \sum_{j \neq i} \mathbf{b}_j(a^*) - \max_{a_{-i} \in A_{-i}} \sum_{j \neq i} \mathbf{b}_j(a_{-i})$$

Let $A = v_i(a^*) + \sum_{j \neq i} \mathbf{b}_j(a^*)$ and $B = \max_{a_{-i} \in A_{-i}} \sum_{j \neq i} \mathbf{b}_j(a_{-i})$. Bidder *i*'s goal is to set \mathbf{b}_i to maximize A - B. Since *B* does not depend on b_i , this reduces to setting \mathbf{b}_i so as to maximize *A*. This is exactly the allocation that is chosen by welfare-maximizing rule assuming bidders are truthful. Thus, bidder *i* maximizes their utility by setting $\mathbf{b}_i = \mathbf{v}_i$. \Box

Example. Suppose you are organizing a job fair and each firm has a different preference of the booth assignment they receive There are three firms and three possible locations in the room front (F), middle (M), rear (R) They have the following private valuation profiles:

	F	M	R
1	10	2	1
2	100	100	100
3	50	45	40

We can use VCG to compute allocation and payments. Assuming truthful bidding, the welfare-maximizing allocation is (1, F), (2, R), (3, M), that generates total welfare 10+100+45 = 155. Without firm 1, the allocation would be (2, M/R), (3, F) with welfare 100+50 = 150. Thus $p_1 = 150 - (155 - 10) = 5$. Without firm 2, the allocation would be (1, F), (3, M) with welfare 55, thus $p_2 = 55 - (155 - 100) = 0$. Finally, without firm 3, the allocation would be (1, F), (2, M/R) with welfare 110. Thus, $p_3 = 110 - (155 - 45) = 0$.

5 Case Study: Sponsored Search Auctions

As the main application of the theory of markets-with-money, we looked at a case study of sponsored search auctions by Edelman et al. [1].

Recall that in a sponsored search auction problem, there are k slots, the jth slot has a clickthrough rate (CTR) of α_j (non-increasing in j), and the utility of bidder i in slot j is $\alpha_j(v_i - p_j)$, where v_i is the value-per-click of the bidder and p_j is the price charged per-click in slot j.

The unique welfare-maximizing dominant-strategyproof mechanism for this problem is the VCG mechanism described in Section 3.2. The mechanism implemented by Google however, and analyzed by Edelman et al. [1] is the **Generalized Second Price (GSP)** auction.

The generalized-Second-Price (GSP) auction is defined below:

Gneralized Second Price (GSP) Auction.

- 1. Rank the advertisers from highest to lowest bid; assume without loss of generality that $b_1 \ge b_2 \ge \ldots \ge b_n$.
- 2. For i = 1, 2, ..., k, assign the *i*th bidder to the *i*th slot.
- 3. For i = 1, 2, ..., k, charge the *i*th bidder a price of b_{i+1} per click.

Understanding Edelman et al. (2007). Edelman et al. (2007) analyze this GSP auction formally and show that it has a canonical equilibrium that is equivalent to the dominant strate-gyproof outcome of the VCG auction.

Their analysis can be broken down and formalized in several parts.

Lemma 1. (GSP is not Dominant-Strategy Proof) For every $k \ge 2$ and sequence $\alpha_1 \ge \ldots \alpha_k > 0$ of CTRs, the GSP auction is not dominant strategyproof (that is, truthful bidding is not a dominant strategy).

Proof. To show that GSP is not dominant-strategyproof, it is sufficient to show that there exists some valuation profile (v_1, v_2, \ldots, v_n) , bid profile (b_1, \ldots, b_n) and bidder *i*, such that, the utility $u_i(b_i, b_{-i})$ of *i* by bidding $b_i \neq v_i$ is greater than the truthful utility $u_i(v_i, b_{-i})$.

Assume that $v_1 \ge v_2 \ge \ldots \ge v_n$. Consider a bidder $i \ge 3$ and assume that all other bidders are bidding truthful. When *i* is truthful, they are in slot *i* and pay v_{i+1} as payment. On the other hand, if *i* misreports to get slot i+1, their payment is v_{i+2} . We construct a valuation profile where this misreport improves i's utility proving the claim.

$$\alpha_i(v_i - v_{i+1}) < \alpha_{i+1}(v_i - v_{i+2}) \tag{13}$$

Now suppose that $v_{i+1} = v_{i+2} + 1$ and $v_i = \frac{\alpha_{i+1}}{2(\alpha_i - \alpha_{i+1})} + v_{i+2} + 1$. Then, this satisfies Inequality 13. Thus, for any CTRs, there exists a valuation profile and bid profile such that some bidder benefits from misreporting showing that GSP is not dominant-strategyproof.

5.1 Nash Equilibrium of GSP

First, we look at the bid profiles that form a Nash equilibrium of the GSP auction, assuming each player has full information about the bids and values of others.

Fix CTRs for slots and valuers-per-click for bidders. We can assume that k = n by adding dummy slots with zero CTRs (if k < n) or dummy bidders with zero value-per-click (if k > n).

Lemma 2. A bid profile **b** is a Nash equilibrium of GSP if no bidder can increase her utility by unilaterally changing her bid. Verify that this condition translates to the following inequalities, under our standing assumption that $b_1 \ge b_2 \ldots \ge b_n$ for every *i*:

$$\alpha_i(v_i - b_{i+1}) \ge \alpha_j(v_i - b_j) \quad \text{for every higher slot } j < i \tag{14}$$

$$\alpha_i(v_i - b_{i+1}) \ge \alpha_j(v_i - b_{j+1}) \quad \text{for every lower slot } j > i \tag{15}$$

Proof. Consider a bidder *i* with bid b_i in slot *i* and fix all other bids b_{-i} . The utility of bidder *i* for being truthful is $u_i = v_i(\alpha_i - b_{i+1})$.

Suppose bidder *i* now wants to misreport and target slot *j* above *i* (that is, j < i). To achieve this, *i* has to bid $b'_i \ge b_j$ and $b'_i \le b_{j-1}$, and their payment then becomes b_j . The utility *i* gets from this deviation is $u'_i = v_i(\alpha_j - b_j)$.

Thus, for a Nash equilibrium this deviation should not give better utility than truthful bidding:

$$\alpha_i(v_i - b_{i+1}) \ge \alpha_j(v_i - b_j)$$
 for every higher slot $j < i$

Now, suppose bidder *i* wants to misreport and target a slot *j* below *i* (that is, j > i). To achieve this, *i* has to bid b'_i such that $b_{j+1} \leq b'_i \leq b_j$. Thus, bidder *j* currently in b_j moves up to slot j-1 and bidder *i*'s price for slot *j* is b_{j+1} .

The utility *i* gets from this deviation is $u'_i = v_i(\alpha_j - b_{j+1})$. Thus, for a Nash equilibrium this deviation should not give better utility than truthful bidding:

$$\alpha_i(v_i - b_{i+1}) \ge \alpha_j(v_i - b_{j+1})$$
 for every lower slot $j > i$

The problem with Nash equilibrium in GSP is that there are many such bid profiles and some of them are socially inefficient (bids are not value-ordered).

Example. Consider an example with three slots with $\alpha_1 = 0.2$, $\alpha_2 = 0.18$ and $\alpha_3 = 0.1$ and four bidders with valuations: $v_1 = 4$, $v_2 = 10$, $v_3 = 2$ and $v_4 = 1$. The bid profile $b_1 = 4$, $b_2 = 2.1$, $b_3 = 2$ and $b_4 = 1$ is a Nash equilibrium but does not maximize social welfare as the slots are not assigned in value order.

5.2 Envy-free Outcome

To refine Nash equilibrium further for GSP, the authors look at the envy-free outcomes.

Definition 8. A bid profile **b** with $b_1 \ge \ldots \ge b_n$ is **envy-free**^{*a*} if for every bidder *i* and slot $j \ne i$, we have

$$\alpha_i(v_i - b_{i+1}) \ge \alpha_j(v_i - b_{j+1}). \tag{16}$$

^aWhy "envy free"? Setting $p_j = b_{j+1}$ for the current price-per-click of slot j, then these inequalities translate to: "every bidder i is as happy getting her current slot at the current price as she would be getting any other slot at that slot's current price.

Lemma 3. Every envy-free bid profile is a Nash equilibrium.

Thus, each envy-free bid profile is also a Nash equilibrium.

Proof. The envy-free condition is the same as the Nash equilibrium condition if j > i. Now suppose j < i. Since bids are sorted, we know that $b_{j+1} \le b_j$, thus $\alpha_j(v_i - b_{j+1}) \ge \alpha_j(v_i - b_j)$. Combining this with the envy-free condition gives us the inequality for a Nash equilibrium.

There are still many envy-free Nash equilibria. To understand which of these are most likely to

be played, the authors consider the fact that bidders are likely to keep increasing their bids b_i (and thus the price of the bidder in slot i - 1). The possible downside of this appraoch is that the bidder in slot i - 1 can retaliate and underbids such that that bidder i has to pay their own bid as the price for slot i - 1. This motivates the definition of locally-envy free Nash equilibrium, which is the result of these best-response dynamics between bidders until they do not envy the bidder immediately above or below.

Definition 9. A bid profile is **locally envy-free** if the envy-free condition (Definition 8) holds for every pair of adjacent slots—for every i and $j \in \{i - 1, i + 1\}$.

Lemma 4. In a GSP auction with strictly decreasing CTRs, a bid profile is locally envy-free if and only if it is envy free.

Proof. We split this into three parts.

• Part 1. Bids are value ordered.

Consider the locally envy-free condition for bidders i and i + 1 not envying slots i + 1and i respectively, we get:

$$\alpha_i(v_i - b_{i+1}) \ge \alpha_{i+1}(v_i - b_{i+2}) \tag{17}$$

$$\alpha_{i+1}(v_{i+1} - b_{i+2}) \ge \alpha_i(v_{i+1} - b_{i+2}) \tag{18}$$

(19)

Isolating $\alpha_{i+1}b_{i+2}$ in Inequality 18, we get

$$\alpha_{i+1}b_{i+2} \le \alpha_{i+1}v_{i+1} - \alpha_i v_{i+1} + \alpha_i b_{i+2}$$

Substituting this on the RHS of Inequality 17 we get:

$$\alpha_i(v_i - b_{i+1}) \ge \alpha_{i+1}(v_i - b_{i+2}) \tag{20}$$

$$\geq \alpha_{i+1}v_i - (\alpha_{i+1}v_{i+1} - \alpha_i v_{i+1} + \alpha_i b_{i+2})$$
(21)

$$v_i(\alpha_i - \alpha_{i+1}) \ge v_{i+1}(\alpha_i - \alpha_{i+1}) + \alpha_i(b_{i+1} - b_{i+2})$$
(22)

(23)

Since we know that $b_{i+1} \ge b_{i+2}$, we can simplify to:

$$(v_i - v_{i+1})(\alpha_i - \alpha_{i+1}) \ge 0$$

Since $\alpha_i \geq \alpha_{i+1}$ for all *i*, it must be that $v_i \geq v_{i+1}$ for all *i* as well. Thus, we get that the bids are value ordered at a locally-envy free profile.

Part 2. Downward deviations are envy-free. To argue that downward deviations are envy free, we show that bidder 1 in slot 1 does not have any incentive to deviate to any slot j > 2. The same argument generalizes to all bidders.

Using locally envy free condition on bidder 1 and 2 that says they do not envy slots 2 and 3 respectively, we get:

$$\alpha_1(v_1 - b_2) \ge \alpha_2(v_1 - b_3) \tag{24}$$

$$\alpha_2(v_2 - b_3) \ge \alpha_3(v_2 - b_4) \tag{25}$$

(26)

Isolating $\alpha_2 b_3$ in Inequality 25 above we get:

$$\alpha_2 b_3 \le \alpha_2 v_2 - \alpha_3 v_2 + \alpha_3 b_4$$

Substituting this in Inequality 24 we get:

$$\begin{aligned} \alpha_1(v_1 - b_2) &\geq \alpha_2 v_1 - (\alpha_2 v_2 - \alpha_3 v_2 + \alpha_3 b_4) \\ \alpha_1(v_1 - b_2) &\geq \alpha_2(v_1 - v_2) + \alpha_3(v_2 - b_4) \\ &\geq \alpha_3(v_2 - b_4) + \alpha_2(v_1 - v_2) + \alpha_3 v_1 - \alpha_3 v_2 \\ &\geq \alpha_3(v_1 - b_4) + \alpha_2(v_1 - v_2) + \alpha_3(v_2 - v_1) \\ &\geq \alpha_3(v_1 - b_4) + (v_1 - v_2)(\alpha_2 - \alpha_3) \\ &\geq \alpha_3(v_1 - b_4) \end{aligned}$$

Thus, bidder 1 has no incentive to deviate to slot 3. The same argument can be extended to show that bidder 1 has no incentive to deviate to any slot lower than 3 and that in fact no bidder has any incentive to deviate to any non-adjacent lower slot.

Part 3. Upwards deviations are envy-free. To argue that upward deviations are envy free, we show that bidder 3 in slot 3 does not have any incentive to deviate up to slot 1. The same argument generalizes to all upward slots and bidders.

Using locally envy free condition on bidder 3 and 2 that says they do not envy slots 2 and 1 respectively, we get:

$$\alpha_3(v_3 - b_4) \ge \alpha_2(v_3 - b_3) \tag{27}$$

$$\alpha_2(v_2 - b_3) \ge \alpha_1(v_2 - b_2) \tag{28}$$

(29)

Isolating $\alpha_2 b_3$ in Inequality 28 above we get:

$$\alpha_2 b_3 \le \alpha_2 v_2 - \alpha_1 v_2 + \alpha_1 b_2$$

Substituting this in Inequality 27 we get:

$$\begin{aligned} \alpha_3(v_3 - b_4) &\geq \alpha_2 v_3 - (\alpha_2 v_2 - \alpha_1 v_2 + \alpha_1 b_2) \\ &\geq \alpha_2 v_3 + v_2 (\alpha_1 - \alpha_2) - \alpha_1 b_2 \\ &\geq \alpha_1 v_3 - \alpha_1 v_3 + \alpha_2 v_3 + v_2 (\alpha_1 - \alpha_2) - \alpha_1 b_2 \\ &\geq \alpha_1 (v_3 - b_2) - v_3 (\alpha_1 - \alpha_2) + v_2 (\alpha_1 - \alpha_2) \\ &\geq \alpha_1 (v_3 - b_2) + (v_2 - v_3) (\alpha_1 - \alpha_2) \\ &\geq \alpha_1 (v_3 - b_2) \end{aligned}$$

Thus, bidder 3 does not envy slot 1. The same argument generalizes to all bidders not envying any upward non-adjacent slot. $\hfill \Box$

5.3 Equivalence of GSP and VCG

Finally, to show that the best-response dynamics in a GSP auction converge to a VCG outcome, we refine the locally-envy free condition further to define a unique bid profile of the agents.

Definition 10. (Balanced bidding) Let the bids b_1, \ldots, b_n of the bidders satisfy the balanced bidding condition if and only if for each bidder *i*:

$$\alpha_i(v_i - b_{i+1}) = \alpha_{i-1}(v_i - b_i)$$

Notice that the balanced bidding condition defines a **unique** bid profile (up to the indifference of the top bidder). We can now show that at this bid profile, the outcome of the GSP auction is the same as the truthful VCG auction.

Theorem 6. The outcome of the GSP auction in a locally envy-free Nash equilibrium bid profile that satisfies balanced bidding is equal to the truthful outcome of the VCG auction.

Proof. As the bids in a locally envy-free outcome are value ordered, the winners are the same in both GSP and VCG auction.

We now show that the payment of bidder *i* in slot *i* is the same in both auctions using induction. As the base case, condition the last slot i = k, then $p_k[\text{GSP}] = p_k[\text{VCG}] = \alpha_k b_{k+1}$.

Consider a slot j < k, then the payment $p_j[\text{GSP}] = \alpha_i b_{i+1}$. Applying the balanced bidding condition on bidder i + 1, we get

$$\alpha_{i+1}(v_{i+1} - b_{i+2}) = \alpha_i(v_{i+1} - b_{i+1}) \tag{30}$$

$$\alpha_i b_{i+1} = (\alpha_i - \alpha_{i+1}) v_{i+1} + \alpha_{i+1} b_{i+2} \tag{31}$$

$$p_i[\text{GSP}] = (\alpha_i - \alpha_{i+1})v_{i+1} + p_{i+1}[\text{GSP}]$$
 (32)

$$= (\alpha_i - \alpha_{i+1})v_{i+1} + p_{i+1}[VCG] = p_i[VCG]$$
(33)

As the final step in the analysis of GSP vs VCG, we can compare their revenue as follows.

Lemma 5. Balanced bidding equilibrium is the lowest-revenue among locally envy-free bid profiles of the GSP auction. Thus, all the locally envy-free bid profiles have revenue at least as much as the truthful VCG outcome.

Proof. The proof follows a similar induction as in Theorem 6 and is omitted. \Box



References

- Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American* economic review, 97(1):242–259, 2007.
- [2] Jason D Hartline. Mechanism design and approximation. Book draft. October, 122(1), 2013.