

Last time we started talking about mechanism design: how to allocate an item to the person who has the maximum value for it, so that we are “incentive compatible”—no one has an incentive to misreport their valuations (a.k.a., to lie). This Vickery auction could be implemented as a natural ascending-price auction.

Then we generalized this to the VCG mechanism which extends to arbitrary allocation problems, with the same incentive-compatibility property. But VCG was a sealed-bid auction, people had to submit their entire valuation functions to the center, which would then output an allocation and some prices. In this lecture, we consider a special kind of market called a matching market, where we want a matching between buyers and goods. For this market we give a “natural” ascending-price auction that generalizes the single-item Vickery auction.¹ Along the way this gives a simple algorithm for finding max-weight matchings in bipartite graphs.

1 Matching Markets

Think of the setting of assigning dorm rooms to students, houses to buyers, etc. Each person wants one of each item. Each item can be given to at most one person. Formally, consider the setting where there is a set I of n items on sale, and a set B of n buyers who may buy some of them. The variables j will refer to a generic buyer and variables i will refer to items. Each buyer wants to buy at most one item, and has *valuation* $v_j(i) \in \mathbb{Z}$ for being allocated item i .² We assume that $v_j(i) \geq 0$. An example with $n = 3$, items on the left, buyers on the right, buyer a values items A, B, C at 10, 9, 0 respectively.

Items	Buyers
Ⓐ	Ⓐ 10,9,0
Ⓑ	Ⓑ 8,7,2
Ⓒ	Ⓒ 7,2,5

If the price for item i is p_i , the utility that j gets from being assigned item i is

$$u_j(i) := v_j(i) - p_i.$$

Naturally, the buyer seeks to maximize his utility. Moreover, we want to ensure *individual rationality*, that the assignment results in every buyer achieving non-negative utility; hence if item i is assigned to buyer j at price p_i , it must be that $v_j(i) \geq p_i$.


The *social welfare* achieved by an assignment of items to buyers is the sum of *valuations* of the buyers for their allocations. Not the sum of the utilities, the actual valuations. (If you sum utilities, you must include the utility of the sellers in the summation.) In the example above, at prices

¹CLRS and DPV do not cover auctions and mechanism design. But a gentle and very readable introduction to matching markets (with more examples) can be found in Chapter 10 of the Easley-Kleinberg book: *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*.

²In this model, the value of person j getting some room i does not depend on who else is assigned which other rooms. This model may not be accurate in some cases: e.g., you may value your assigned dorm room more or less depending on who has the neighboring rooms. We are not modeling such effects. In today’s lecture, the value of j for some assignment of items to buyers depends only on the item that j gets in this assignment.

$p_A = 6, p_B = 5, p_C = 3$, assigning $A - a, B - b, C - c$ would achieve social welfare $10 + 7 + 5 = 22$. The utility of the players under this assignment would be $10 - 6 = 4, 7 - 5 = 2, 5 - 3 = 2$ respectively.

We want to find an allocation that maximizes the social welfare. Last time we saw the VCG mechanism that finds an allocation (and a pricing scheme) which maximizes the social welfare, and also gives incentive compatibility. That is, no player has any reason to misreport their valuations. And what would VCG do in this setting of a matching market? It would compute the max-weight matching in the bipartite graph where the weight of an edge (i, j) is the valuation $v_j(i)$. This would give us the allocation. Then the payment that buyer j makes is the “externality” that j causes: the difference between the valuations of all the other players when i participates, and when it does not. In the example above, a max-weight matching is $M_{VCG} = \{A - a, B - b, C - c\}$. The corresponding VCG prices are given below.



(What is the price VCG charges to a ? The valuations of the other players under M_{VCG} is $7 + 5 = 12$. Now if a does not bid, then the optimal allocation for the others is $B - b, A - c$ with total valuation 14. Hence the “externality” caused is $14 - 12 = 2$, which is the price p_a VCG sets for a .)

But the VCG mechanism required us to submit our valuations to the center, which then finds the allocation and prices in this centralized way. Today we will give a different, distributed mechanism that acts more like real markets seem to do: it will gradually raise prices until they “stabilize”, and each buyer will buy their favorite item (which maximizes their own utility), and that will be the optimal allocation, maximizing social welfare.

Important note: for all of Section 1, we will assume everyone acts truthfully. It will be just an algorithm design problem. We come back to incentive-compatibility in Section 2.

1.1 Preferred Items and the Preferred Graph

Given prices p_1, p_2, \dots, p_n for the items, the preferred items for buyer j are all the items that maximize his utility (and for which the utility is non-negative). I.e., if


$$u_j^* := \max_{i \in I} (v_j(i) - p_i) \tag{1}$$

is the maximum utility that j can achieve at these price, and if $u_j^* \geq 0$, the preferred items (at these prices) are

$$S_j := \{i \in I \mid v_j(i) - p_i = u_j^*\}. \tag{2}$$

Else if $u_j^* < 0$ then buyer j has no preferred items (and $S_j = \emptyset$).

Now you can create the *preferred graph* H by having a node for each item $i \in I$, one for each buyer $j \in [n]$, and an edge (i, j) if item i is preferred item for j , i.e., $i \in S_j$. (All this is with respect to the current set of prices, of course.) Here are the preferred graphs with respect to two different price settings.



1.2 Market-Clearing Prices


A set of prices p_1, p_2, \dots, p_n , one for each item, is called *market-clearing* if

- There is a matching in the preferred graph H (at these prices) that matches *all* the buyers to items, and
- if an item i is not matched, then its price $p_i = 0$. (This is irrelevant when the number of items equals the number of buyers, since if all buyers are matched, so are the items. But if you had more items, you need this condition.)

This means at these prices, each buyer can come by and pick some item that maximizes her utility³, and the market will “clear”. (Of course, items that are not desirable might be left behind, but they have no value in this market and hence have zero price.) The matching given by the definition of market-clearing prices gives us the corresponding allocation.

In the previous figure on the right $p_A = 2, p_B = 1, p_C = 0$ were market-clearing prices, the graph has a perfect matching.

But $p_A = 1, p_B = 0, p_C = 0$ are not market-clearing prices. Why not? There is no perfect matching in the preferred graph H (on the left). How do you prove the absence of a perfect matching? Look at the set $S = \{a, b, c\}$. The neighborhood has size $|N(S)| = |\{A, B\}| = 2$ whereas this set has size $|S| = 3$. There is no way we can match S to its neighbors. Hall’s theorem⁴ says that a bipartite graph with equal number of vertices on either side has no perfect matching if and only if there is a set $S \subseteq B$ with neighborhood size $|N(S)| < |S|$. The set $N(S)$ is called a *constricted set* or a *Hall set*.⁵ Here is another example of a bipartite graph with $n = 8$ where there is no perfect matching, as shown by a constricted set, the three (blue) neighbors on the left of the set of four red vertices on the right.



Take away message: if there is no perfect matching in a bipartite graph, we can always find a constricted set. This will be useful later.

³Due to ties, they can’t pick an arbitrary item that maximizes their utility; we might have to give them a utility-maximizing item breaking ties carefully

⁴You may have seen it in 15-251; we also saw this in Recitation #6.

⁵A constricted set is a set $C \subseteq A$ such that there exists $S \subseteq B$ with $N(S) = C$, and $|S| > |C|$.

1.3 Market-Clearing Prices and Social Welfare Maximization

Why are we interested in market-clearing prices, apart from them seeming “natural”? And do they always exist? The first answer is short and sweet.

Theorem 1 *If there are market-clearing prices, then the corresponding allocation maximizes social-welfare.*

Proof: Fix any market-clearing prices p_i . Let M be the matching from items to buyers. By the definition, each player is given a preferred item, and hence is (individually) maximizing her utility. So the maximum utility that can be achieved (by any matching) with respect to these prices is precisely

$$\sum_{(i,j) \in M} u_j(i) = \sum_{(i,j) \in M} (v_j(i) - p_i).$$

Now look at the items. Some of them are matched to buyers and their prices are included in the sum. Others are not matched, and by the definition again, their prices must be zero. Hence the above expressions are equal to

$$\sum_{(i,j) \in M} v_j(i) - \sum_i p_i.$$

But the latter sum does not depend on the matching, so in finding the matching maximizing this expression, we have found a matching that maximizes the social welfare. QED. ■

Cool: if we have market-clearing prices, they automatically maximize the value to the buyers.

And why should these market-clearing prices exist? Here’s an algorithm that will terminate with these prices.

1.4 An Algorithm for Market-Clearing Prices


We want to find market-clearing prices. The idea will be a natural one: if a set of k items has more than k people preferring it at the current prices, it seems likely that the prices of these items should rise. And that is almost exactly what we will do.

Start with all prices $p_i = 0$. Our prices will always be integers. Recall that the valuations $v_j(i)$ are all non-negative integers. Now do the following:

1. Check if the current prices are market-clearing. I.e., build the preferred graph, check if there is a matching that matches all the buyers to items.⁶ If so, stop and output prices/matching.
2. If there is no such matching, Hall’s theorem says there exists some set S of buyers such that the set of neighbors of S in the graph H (denoted by $N(S)$) has strictly smaller cardinality. I.e., $|S| > |N(S)|$. (Recall we called $N(S)$ a *constricted set*).
Such a set is “over-subscribed”, the demand for it (i.e., $|S|$) is more than the supply (i.e., $|N(S)|$). So such a set seems like a natural candidate to see price increases. For each item $i \in N(S)$, increment the prices $p_i \leftarrow p_i + 1$.
3. If this causes the minimum price to become non-zero (i.e., it must be 1), subtract 1 from each price so that the minimum price becomes zero,
4. Go back to Step 1.

⁶Since the number of items equals the number of buyers, if there is a matching that no items are left over, and we don’t have to check the zero-price condition for unassigned items.

Here’s a run of the algorithm (with the constricted sets shown in red).



If the algorithm terminates it will output market-clearing prices. So we show that the algorithm terminates. First, a simple lemma that justifies the price-lowering step at the very end.

Fact 2 *Each buyer has degree at least 1 in the preferred graph (at all times).*

Proof: By construction, at least one item stays at price 0. The utility of every buyer for this item remains non-negative, and hence the preferred set for any buyer can never be empty. ■

Theorem 3 *The algorithm terminates in finite time.*

Proof: We use a potential function. Given prices p_1, p_2, \dots, p_n , define the potential to be

$$\Phi = \sum_{\text{items } i} p_i + \sum_{\text{buyers } j} u_j^*. \quad (3)$$

where u_j^* is defined as in (1). Note that all the terms here are non-negative. At the beginning, the prices are zero, and the $u_j^* = \max_{i \in I} v_j(i)$. So the initial potential $\Phi_0 = \sum_j \max_{i \in I} v_j(i)$.

Now, every time we execute Step 2 of the algorithm, the potential decreases. Why? By Fact 2, each buyer has at least one edge out of it, so $|N(S)| \geq 1$. Raising the prices on $|N(S)|$ nodes increases the potential by $|N(S)|$. However, the value of u_j^* decreases for all nodes in S —this decreases the potential by $|S|$. Since $|S| > |N(S)|$, the potential decreases by at least 1. So the algorithm stops in at most Φ_0 steps. ■

BTW, this algorithm is not polynomial-time, if the maximum valuation is V_{\max} this process could take $\Omega(V_{\max})$ time to finish. You can change the algorithm slightly to take “bigger” steps, and get a poly-time algorithm out of it. (That algorithm is called the *Hungarian Algorithm*, due to Kuhn’s reinterpretation of a theorem of Egervary.)

1.5 The Ascending-Price Mechanism*

You may wonder if the price-lowering step was actually required. What if we just wanted to have “ascending prices”, which is more natural. One can get such an algorithm as well, with a little more care. The new algorithm is the following:

1. Check if the current prices are market-clearing. (I.e., build the preferred graph, check if there is a matching that matches all the buyers to items.) If so, stop and output prices/matching.
2. Else, there must be some set S of buyers such that the set of neighbors of S in the graph H (denoted by $N(S)$) have strictly smaller cardinality. I.e., $|S| > |N(S)|$. (Recall we called $N(S)$ a *constricted set*). Such a set is “over-subscribed”, the demand for it (i.e., $|S|$) is more than the supply (i.e., $|N(S)|$). So such a set seems like a natural candidate to see price increases.

There may be multiple constricted sets $N(S)$. Choose a *minimal* constricted set: i.e., $N(S)$ does not contain another constricted set $N(T)$. For each item $i \in N(S)$, increment the prices $p_i \leftarrow p_i + 1$. Go back to Step 1.

The choice of raising prices on a minimal constricted set gives us that at least one zero-price element remains, as the following lemma shows.

Lemma 4 *At most $n - 1$ items will have strictly positive prices.*

Proof: Just for the sake of analysis, maintain a tentative matching M around, where M contains a subset of edges in the preferred graph. We maintain the invariant that if an item has non-zero cost, that item is tentatively matched to some buyer. In math, $p_i > 0 \implies \exists j : (i, j) \in M$. Initially M is empty, which satisfies the invariant.

Suppose you have prices and a tentative matching, and now you raise prices for items in $N(S)$. Tentatively match all the items in $N(S)$ to buyers in S . (If either these items, or the buyers they are matched to, were tentatively matched earlier, drop those tentative matching edges.) Since $|S| > |N(S)|$ and we chose a minimal $N(S)$, we know that this tentative matching can be done. (Why? Think about it — it is Hall's theorem again!) So the items in $N(S)$, they have non-zero prices and are tentatively matched. And the items outside $N(S)$ must be tentatively matched to buyers outside S (since all neighbors of S are in $N(S)$), and we did not drop any such edge. Hence, the invariant is maintained.

For sake of contradiction, if we ever want to raise the price for the n^{th} item, we would get a tentative matching of size n , which would mean we would have had a matching of all buyers to items, and hence would not want to raise prices at this step. Hence, at most $n - 1$ items would ever have their prices raised. ■

Using this Lemma, one can prove Fact 2 and hence Theorem 3 for this ascending-price algorithm.

2 Relationship to Vickery and VCG

The auctions we just saw are a direct generalization of the Vickery auction from the previous lecture. In the Vickery auction there was a single item we were auctioning among n buyers. To make it fit our model here, just introduce $n - 1$ dummy items (numbered 2 through n), and number the real item 1. And if the buyer j had value v_j for the real item, define $v_j(1) = v_j$ and $v_j(i) = 0$ for $i \geq 2$. Assuming that each $v_j \geq 0$, the preferred graph starts off with edges from item 1 to all buyers. It is easy to check that $\{1\}$ is the unique constricted set (make sure you see why!), and we raise its price. Each time we raise its price, the preferred graph changes, but no other element can enter the constricted set. Finally, all except one buyer has a preferred edge to at least one item from $\{2, 3, \dots, n\}$, and there is no more constricted set. This will happen precisely when p_1 reaches the second-highest valuation, and the person matched to the real item will be the one with the highest valuation.

Also, for the case when we have K identical items to sell, the same argument shows that our algorithm above will sell to the K highest bidders, at a price equal to the valuation of the $K + 1^{\text{st}}$ -highest bidder.

Is this just a coincidence? Or could it be the case that the matching and prices found by our algorithm above are the VCG allocation and prices? That would be great, because this would say that the ascending-price mechanism is also incentive compatible, that it is a dominant strategy for the players to play truthfully. Well, clearly the matching we found is the max-weight matching (by Theorem 1), and hence the same as the VCG allocation. But what about the prices? This is a little more tricky, but it is true: *the prices we find by the ascending-price auction are exactly the VCG prices!* We'll not cover this here, see Chapter 15.9 of the Easley-Kleinberg book for more details.

3 Walrasian equilibrium*

The same ideas work for more general markets where the players need not just want one item, they may be interested in buying multiple items. The valuation functions may have different characteristics. For instance, you may want a hotel room and a flight ticket and a car rental (and any strict subset of these might be useless to you). Or you may want multiple computers, but each additional computer gets less valuable to you the more you have (“diminishing returns”). Or you may value set A at \$50, or set B at \$100, but all other sets are worthless to you. One can imagine very general valuation functions.⁷

Formally, suppose you have a set of m items I , a set of n buyers B . Each player $j \in B$ now has valuations $v_j(S)$ for each subset $S \subseteq I$ of items. Let us define some useful shorthand: given prices p_1, p_2, \dots, p_m , define

$$p(S) := \sum_{i \in S} p_i. \quad (4)$$

The utility of a set S for buyer j is $u_j(S) := v_j(S) - p(S)$. We imagine that $v_j(\emptyset) = 0$, and hence $u_j(\emptyset) = 0 - p(\emptyset) = 0$.

Definition 5 (Preferred Set) *At prices p_1, p_2, \dots, p_m , a set S is called preferred for player j if $u_j(S) \geq 0$, and also for all $T \subseteq I$, $u_j(T) \leq u_j(S)$.*

Definition 6 (Walrasian Equilibrium) *The prices p_1, p_2, \dots, p_m and allocations S_1, S_2, \dots, S_n are at Walrasian equilibrium if (a) for each j the set S_j is preferred for player j , and (b) if item i is not allocated (i.e., $i \notin \cup_j S_j$), the price $p_i = 0$.*

This is the natural way to extend market-clearing prices to the more general setting, and a proof similar to Theorem 1 shows that if $\{p_i\}_{i \in I}, \{S_j\}_{j \in B}$ are a Walrasian equilibrium, then the allocation $\{S_j\}$ maximizes social welfare.

3.1 Existence No Longer Guaranteed

Unfortunately, Walrasian equilibria don’t always exist, when the valuation functions are allowed to be so general.

There are 2 items $\{A, B\}$ and 2 buyers $\{a, b\}$. Buyer a values the entire set at 3, and all other sets have value 0. b values every non-empty subset at 2. The welfare maximizing allocation is to give both items to a , with social welfare 3. For this to be a Walrasian equilibrium, what prices should we choose? The prices of each of the items must be at least 2 (for b ’s preferred set to be empty). But then the price of the entire set is at least 4, so a ’s preferred set cannot be $\{A, B\}$.

3.2 An Ascending-Price Mechanism for “Nice” Valuations

The problem in the example above comes from “complementarity”. Items A and B “complement” each other, and the whole $\{A, B\}$ set is worth more to a than the sum of its parts. For example, think of A as being a left shoe and B the right shoe. Or a hotel room and a flight ticket. Or bread and butter.

But if the valuation functions are “nice”, an ascending-price auction find an (approximate) Walrasian equilibrium. What are these “nice” valuations? These are called “substitutes”, and is one

⁷For more details on this section’s material, see the survey by Blumrosen and Nisan from the *Algorithmic Game Theory* book, which has loads of other cool topics. (Link to download entire book, 5Mb.)

way of ensuring the lack of complements. Loosely, it says that if you raise the price of a set of items, your preference or demand for some other item does not go down. (This clearly excludes complements, since if you were to raise the price of left shoes, not only would the demand for left shoes go down, but the demand for right shoes would also go down despite their prices not changing.)

Definition 7 (Substitutes) *Formally, valuation v_j satisfies the substitutes property if the following holds. For any two prices p_1, p_2, \dots, p_m and q_1, q_2, \dots, q_m where $p_i \leq q_i$, if P is a preferred set for buyer j at prices p , and $P^= := \{i \in P \mid p_i = q_i\}$ is the subset of P containing those items whose prices did not change, then there exists a preferred set Q for j at prices q with $P^= \subseteq Q$.*

Start with tentative assignments $T_j = \emptyset$, and prices $p_i = 0$.

1. Look at the preferred set P_j for buyer j under the following prices: each item $i \in T_j$ costs p_i , and each item $i \notin T_j$ costs $p_i + \epsilon$. (Think of this as saying that acquiring an item not in our tentative set has some ϵ extra cost, you have to work for it.)
2. If each such preferred set $P_j = T_j$, stop and return the prices p_i .
3. Else, pick some buyer with $P_j \neq T_j$, who is not satisfied. Set $T_j = P_j$, that is, give him his preferred set.

Now increase the prices of all items newly assigned to j , namely those in $P_j \setminus T_j$, to $p_i + \epsilon$. Finally, remove these items in P_j from other tentative sets: $T_{j'} \leftarrow T_{j'} \setminus P_j$. Go back to Step 1.

Again, a natural ascending-price mechanism. If the valuations satisfy the substitutes property, this gives us an ϵ -approximate Walrasian equilibrium. The precise definition of this approximate equilibrium is not important, what is interesting is how this simple algorithm can give market-clearing prices. (Sadly, now the prices are no longer VCG prices, and players have an incentives to lie. One can get around this by pricing sets of items, but that is a story for another day.)