

Mini-Course on Market Design: Matching, Carpooling, and Congestion Pricing

Michael Ostrovsky
Stanford GSB

IWGTEA 2026
São Paulo, Brazil



Lecture 1: Two-sided matching with transfers

Lecture 2: Vickrey (1969) model of traffic congestion. Congestion pricing, carpooling, and commuter welfare.

Lecture 3: Bringing it all together: self-driving cars, carpooling, and congestion pricing

Lecture 4: Empirics: Congestion Pricing in New York City

Lecture 1: Two-sided matching with transfers

Warmup: Matching without transfers

m men, n women, the usual Gale-Shapley notion of stability

Sotomayor (1996), “A Non-constructive Elementary Proof of the Existence of Stable Marriages,” *Games and Economic Behavior*, 13, 135–137.

Definition: Matching μ is *simple* if for any blocking pair (m, w) , woman w is single.

The empty matching is simple, and the set of simple matchings is finite. Therefore, there must exist a simple matching that is Pareto optimal for men (say, μ^*).

Key observation: matching μ^* must be stable! Why? Suppose it is blocked by some pair (m, w) . By simplicity of μ^* , woman w is single. Of all blocking pairs involving w , pick her favorite man m' , and generate a new matching by matching w with m' (and dropping the old match of m' if necessary). The new matching m' is also simple, and (for men) Pareto dominates matching μ^* .

Matching with Transfers

Starting point: 1-to-1 matching with transfers (1 firm is matched to 1 worker, 1 customer is matched to 1 object, etc.).

Primary references: Shapley-Shubik (1971); Chapter 8 in Roth and Sotomayor (1991).

This is also known as the “assignment game” or “optimal assignment problem”.

In Group P, we have m agents $\{i_1, i_2, \dots, i_m\}$.

In Group Q, we have n items $\{j_1, j_2, \dots, j_n\}$.

The surplus produced by a given match $i-j$ is $\alpha_{ij} \geq 0$. You can think of α_{ij} as the “total value that pair $i-j$ generates.”

This gives us a matrix of size $m \times n$ listing all α_{ij} , summarizing the values of all potential matches.

Definition: An *assignment* is a matrix X of size $m \times n$ whose elements are zeros and ones, such that the sum of elements in each column and in each row is at most one.

Definition: A *feasible outcome* is a triple (u, v, X) , where X is an assignment, $u \in \mathbb{R}_+^m$ is a vector summarizing the payoffs of P -agents, and $v \in \mathbb{R}_+^n$ is a vector summarizing the payoffs of Q -agents, such that $\sum_i u_i + \sum_j v_j \leq \sum_{i,j} \alpha_{ij} X_{ij}$.

(I.e., an assignment tells us who is matched to whom; a feasible outcome also tells us how the agents allocate the total surplus generated by all of them together.)

Definition: assignment X is *optimal* if it maximizes the expression $\sum_{i,j} \alpha_{ij} X_{ij}$.

Definition: feasible outcome (u, v, X) is *stable* if for every P -agent i and Q -agent j ,

$$u_i + v_j \geq \alpha_{ij}.$$

Lemma: if (u, v, X) is stable, then

- $u_i = 0, v_j = 0$ for all unassigned i, j ;
- $u_i + v_j = \alpha_{ij}$ for all matched pairs.

Theorem: if (u, v, X) is stable, then it is optimal.

Proof: Take any optimal assignment Y . We have

$$\sum_{i,j} \alpha_{ij} X_{ij} = \sum_i u_i + \sum_j v_j \geq \sum_{i,j} \alpha_{ij} Y_{ij}.$$

Theorem: if (u, v, X) is stable, then it is in the core.

Proof: Analogous to the previous one, with assignment Y restricted to the potentially deviating group of agents.

Key Theorem: There exists a stable outcome.

Proof:

Consider the following Linear Program: $\max \sum_{i,j} \alpha_{ij} x_{ij}$ subject to three conditions:

- i. $\sum_j x_{ij} \leq 1$ for all i
- ii. $\sum_i x_{ij} \leq 1$ for all j
- iii. $x_{ij} \geq 0$ for all i and j

Note that in this LP, we do not require x_{ij} to be integer, they can take any value in $[0, 1]$.

However, there always exists at least one solution in which every x_{ij} is either 0 or 1.

(The standard exposition obtains this as a corollary of “total unimodularity” of LP constraints, although is it not hard to show this directly without invoking other results.)

Last step in the proof of Key Theorem - using LP duality.

General “strong duality” theorem. Let $c \in \mathbb{R}^k$, $b \in \mathbb{R}^\ell$, and $A \in \mathbb{R}^{\ell \times k}$. Suppose the (primal) linear program ($\max c^T x$ s.t. $Ax \leq b$, $x \geq 0$) has a solution x^* . Then its dual, ($\min b^T y$ s.t. $A^T y \geq c$, $y \geq 0$) also has a solution, y^* , and moreover, $c^T x^* = b^T y^*$.

Recall our “primal” LP: $\max \sum_{i,j} \alpha_{ij} X_{ij}$ subject to three conditions:

- i. $\sum_j x_{ij} \leq 1$ for all i
- ii. $\sum_i x_{ij} \leq 1$ for all j
- iii. $x_{ij} \geq 0$ for all i and j

Its dual can be expressed as: $\min \left(\sum_{i=1}^m u_i + \sum_{j=1}^n v_j \right)$ subject to all u_i and v_j being non-negative and

- i. $u_i + v_j \geq \alpha_{ij}$ for all i and j .

Take any solution X^* of the primal LP and any solution (u^*, v^*) of the dual LP. By strong duality, $\sum_{i,j} \alpha_{ij} X^*_{ij} = \sum_{i=1}^m u_i^* + \sum_{j=1}^n v_j^*$, and so (u^*, v^*, X^*) is a feasible outcome. The constraints in the dual LP also imply that it is stable.

Lattice structure of the set of stable outcomes

Suppose (u, v, X_1) and (u', v', X_2) are stable outcomes and X is an optimal assignment.

For each i , let $\underline{u}_i = \min\{u_i, u'_i\}$, and define $\bar{u}$, $\underline{v}$, and $\bar{v}$ analogously.

Then $(\underline{u}, \bar{v}; X)$ and $(\bar{u}, \underline{v}; X)$ are feasible, stable outcomes.

Proof for $(\underline{u}, \bar{v}; X)$:

First, it is immediate that $\underline{u}_i \geq 0$, $\bar{v}_j \geq 0$ (individual rationality)

Second, it is also immediate that (u, v, X) and (u', v', X) are feasible, stable outcomes.

Third, observe that if for some i and j , $X_{ij} = 1$, then $u_i + v_j = u'_i + v'_j = \alpha_{ij}$, and so if $\underline{u}_i = u_i$, then $\bar{v}_j = v_j$, and if $\underline{u}_i = u'_i$, then $\bar{v}_j = v'_j$. Thus, $(\underline{u}, \bar{v}; X)$ is a feasible outcome.

Finally, to prove that it is stable, take any i and j . Either $\underline{u}_i = u_i$ or $\underline{u}_i = u'_i$. Without loss of generality, we'll work with the first case. Then

$$\underline{u}_i + \bar{v}_j = u_i + \bar{v}_j \geq u_i + v_j = \alpha_{ij}.$$

EXAMPLES

What are the sets of stable outcomes in the following three examples?

Example 1. There is only one agent in P ($m = 1$) and 4 agents in Q ($m = 4$). The values of potential matches are given by the following matrix of α_{ij} 's: $A = \begin{bmatrix} 6 & 8 & 3 & 5 \end{bmatrix}$.

Example 2(a). There are 10 agents in P and 10 agents in Q . The value of every possible match is 1 (i.e., for every i and every j , $\alpha_{ij} = 1$).

Example 2(b). There are 10 agents in P and 11 agents in Q . The value of every possible match is 1 (i.e., for every i and every j , $\alpha_{ij} = 1$).

Crawford-Knoer (1981)

Crawford and Knoer (1981) show that the similarities are not accidental. Their setup:

- $P \rightarrow \text{Firms } (i)$ and $Q \rightarrow \text{Workers } (j)$.
- $u_{ij} \rightarrow$ how much firm i would like worker j (with no transfers).
- $v_{ji} \rightarrow$ how much worker j would like firm i if matched to it.

Notice where this model comes from:

- If we think of only $u_{ij} + v_{ji}$ (and have continuous transfers), we get Shapley-Shubik.
- If we think of only individuals' preferences (and no transfers), we get Gale-Shapley.

Key idea: consider a market in which we *can* have transfers, but there is only a specific set of transfers allowed. A match must describe (i, j, s_{ij}) , where s_{ij} is a salary from a finite set of possible salaries S_{ij} .

In this setting, we can analyze matching with transfers using the Gale-Shapley approach. No LP needed.

The DAA becomes the *salary adjustment process*:

1. Firms propose to their favorite worker at the lowest salary.
2. Workers reject all offers except for their favorite one, which they hold on to.
3. Firms that were rejected propose again. If a firm is re-proposing to a worker who has previously rejected it, the firm must increase the offered salary.

By essentially the same argument as in Gale-Shapley, this process will eventually stop, at which point we finalize all of the matches, and end up with a stable outcome. (And one can show that many other properties also continue to hold.)

Of course, the set S_{ij} of possible salaries can be made as fine as we want it to be. If it only contains zeros, we are in the Gale-Shapley world. If we take the unit of discretization to zero, in the limit the set of stable outcomes converges to that in the Shapley-Shubik world.

Hatfield and Milgrom (2005): Matching with Contracts

(See also Kelso-Crawford, 1982, and Roth, 1984)

- firms f , workers w , contracts X (all finite);
- Each contract $x \in X$ is associated with one firm f and one worker w ;
- $1 - n$ matching (one firm – multiple workers);
- For each pair $f - w$, the matching specifies which contract x is chosen by that pair;
- A matching can be viewed simply as a set of contracts $Y \subset X$; feasible if no two contracts in Y involve the same worker w .

- Preferences of both sides are strict; $Ch_a(Y)$ denotes the optimal choice of agent a from set of contracts Y .
- $R_a(Y) = Y \setminus Ch_a(Y)$
- Firms' preferences are *substitutable* if for any firm f and any sets of contracts $X'' \subset X' \subset X$, $R(X'') \subset R(X')$.
- A matching is pairwise stable if it is individually rational and is not blocked by any firm-worker pair (and some associated contract).

Theorem There exists a pairwise stable matching.

Theorem There exists a pairwise stable matching.

Proof

Say that a matching is *simple* if any blocking pair involves an unemployed worker.

Take a firm-Pareto-optimal simple matching μ .

It must be stable. (Suppose not, and suppose there is a blocking contract involving an unemployed worker w . Of all the blocking contracts involving worker w , take his favorite one, x , and suppose it involves firm f . Consider matching μ' that adds contract x to μ , drops all contracts that f no longer wants to keep after adding x , and keeps all other firms' contracts unchanged. Then matching μ' is also simple, and also firm-Pareto-dominates μ , which is a contradiction.)

Summary of technical machinery in this setting:

- Deferred-Acceptance Algorithm and Salary Adjustment Process
- Linear Programming
- Sotomayor's non-constructive approach via simple matchings
- (Did not cover, beyond the scope of this mini-course) Lattice-theoretic approach, finds stable matchings as fixed points of an isotone operator on an appropriately defined space (Hatfield and Milgrom, 2005; Ostrovsky, 2008)

Applications:

- College admissions with endogenously determined financial aid (Biró, Hassidim, Romm, Shorrer, and Sóvágó, 2025)
- Matching U.S. Army cadets to military service branches with endogenously determined length of commitment (Greenberg, Pathak, and Sonmez, 2024)
- Pricing and matching (passengers with drivers) on rideshare platforms