CSC200: Lecture 18

- Today: Continue with Matching Markets:
  - Market clearing prices, properties, computation
  - Ch.10 (*including* advanced material in 10.6)
- Next week wrap up matching markets and then stable matchings
  - Stable matchings material not in text
- After the break:
  - The internet, search engines and internet advertising: Ch.13-15
  - various network phenomena at a more aggregate (population) level (Ch.16-18), e.g., information cascades, power laws, tipping points, etc.
- Announcements
  - Term test Friday; Scope of test: social-affiliation networks, Nash equilibria, auctions (including CAs). Usual page of handwritten notes. Test begins at 3:00 in SS 1088 and at 3:10 in SS 1069. 50 minutes
  - Final quiz on Friday, Dec 4
  - Final class this term on Monday, December 7
  - Slides 2 and 3 from Lecture 17 clarify social welfare and POA

Clarification of social welfare and POA

- After class, one student indicated that what I was saying was different from what he found on Wikipedia. This is true. There are two ways in which there is a difference.
- One is just a matter of convention (for maximization problems) in terms of whether the OPT value is in the numerator or the denominator and I prefer that it be in the numerator so that the ratio is greater than or equal to 1.
- But the more substantial difference is that say for an auction I am defining the social welfare as a sum of the values for the allocated items whereas Wikipedia is saying that it is the sum of the utilities.
- Given that I define an agents utililty (for a set of allocated items S) as value(S) –price(S), this does seem like a substantial difference. Turn to next slide for explanation.

Clarifying the definition of social welfare

- I define social welfare in terms of the buyers valuations.
- Wikipedia in terms of their utilities.
- I could have used a more neutral term: *payoff*
- The way to reconcile this is to decide whether or not to include the seller (or mechanism) as an agent. I define the payoff of each buyer as their valuation for any set allocated. I was not viewing the seller/mechanism as an agent and not considering the seller in the social welfare.
- But it is also reasonable to think of the seller/mechanism as an agent and if so then the sellers payoff is the sum of prices (for allocated sets) and a buyers payoff is their utilitity. Summing the payoff of all buyers and the seller is the same as summing all buyer values for sets allocated.

Recap: Bipartite Graph Model

- Consider simple one-sided model:
  - items *Y*, agents *X*
  - *simple preferences:* for any *x*, *y*, either *x* satisfied with *y*, or not satisfied with *y*
- Easy to visualize as a *bipartite graph*
  - nodes *Y*, nodes *X*, edge between *x* and *y* iff *x* is satisfied with *y*
- Example: five students assigned to five summer research projects/faculty advisers
- *One can model two-sided simple preferences with bipartite model; how?*

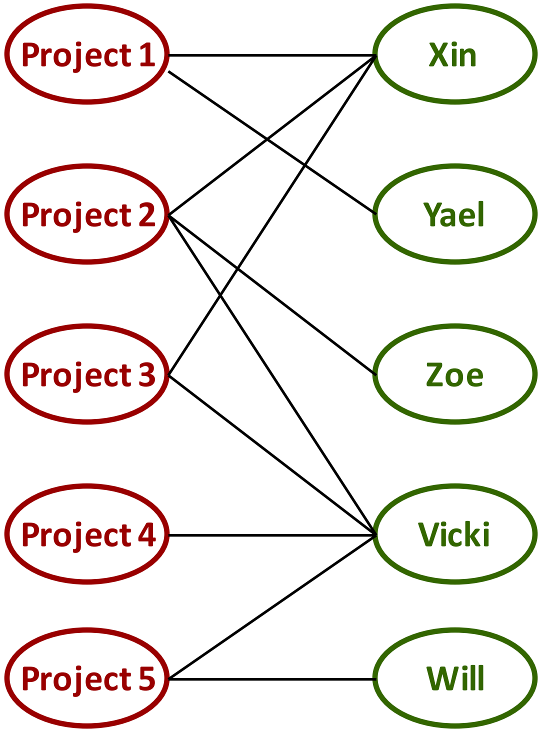

e.g., Vicki would be happy with four different projects (P2, P3, P4, P5); Will only wants one (P5)

Recap: Bipartite Matchings

- A *matching* in a graph *G=(V,E)* is subset $E' \subseteq E$ of the edges with no node participating in more than one edge
- *Matchings defined for general graphs*
- In a bipartite graph, a matching assigns:
  - no more than one agent to an item
  - no more than one item to an agent
- A matching is *maximal* if there is no way to add another edge to it (and still be a matching)
- Can also define max size, or max weight matchings (if edges have weights)
- A bipartite matching is *perfect* iff each node is matched
  - can only exist if graph has equal number of nodes on both sides

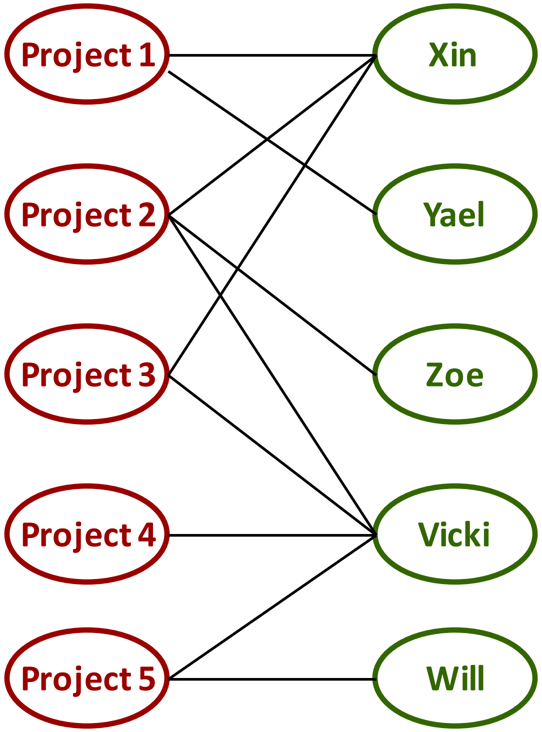

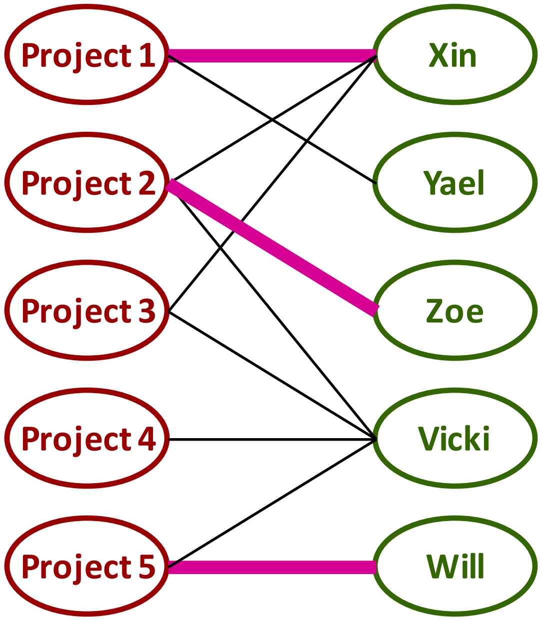

Matching: yes
Maximal: no
Maximum: no
Perfect: no

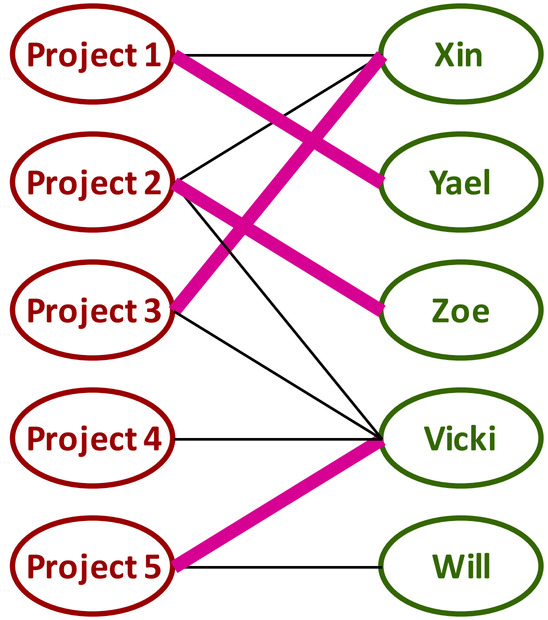

Matching: yes
Maximal: yes
Maximum: no
Perfect: no
Why?

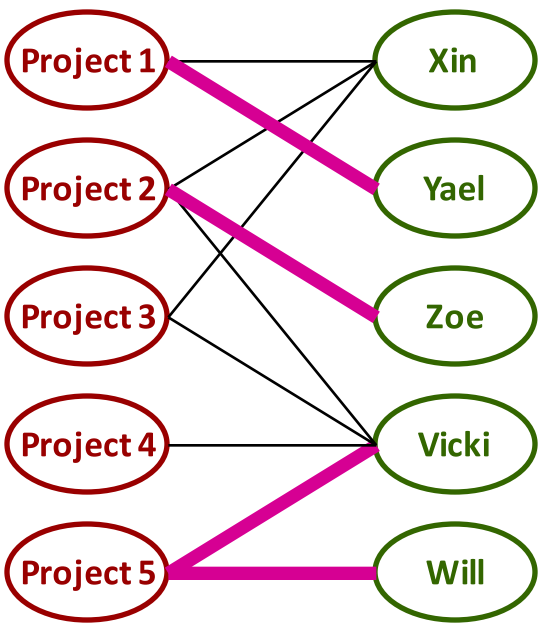

Matching: no

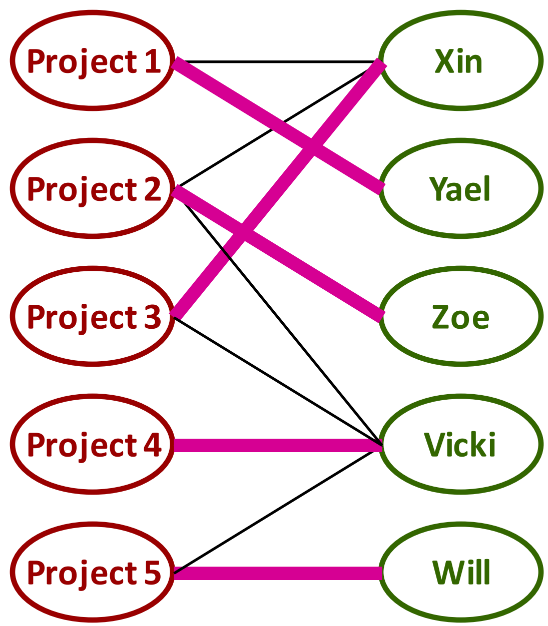

Matching: yes
Maximal: yes
Maximum: yes
Perfect: yes
Perfect Implies maximum
Converse?

Recap: Perfect Matchings and Social Welfare

- Assume equal number of agents, items
- A maximum matching gives us an assignment of items to agents that maximizes total preference (social welfare) in this simple model of preferences
- If perfect matching exists, then everyone is satisfied; if not someone will be left unsatisfied
- If one doesn't exist, how can you know?
  - *Don't want to enumerate all matchings!*
- A simple "test" can verify non-existence

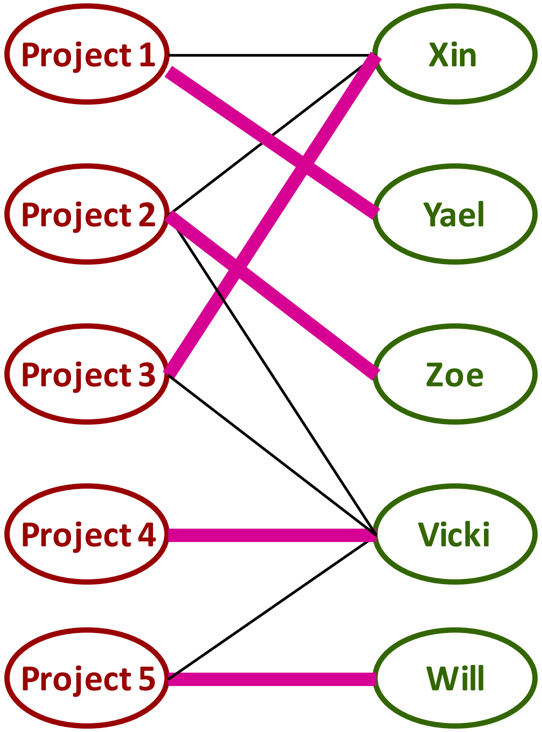

Recap: Constricted Sets

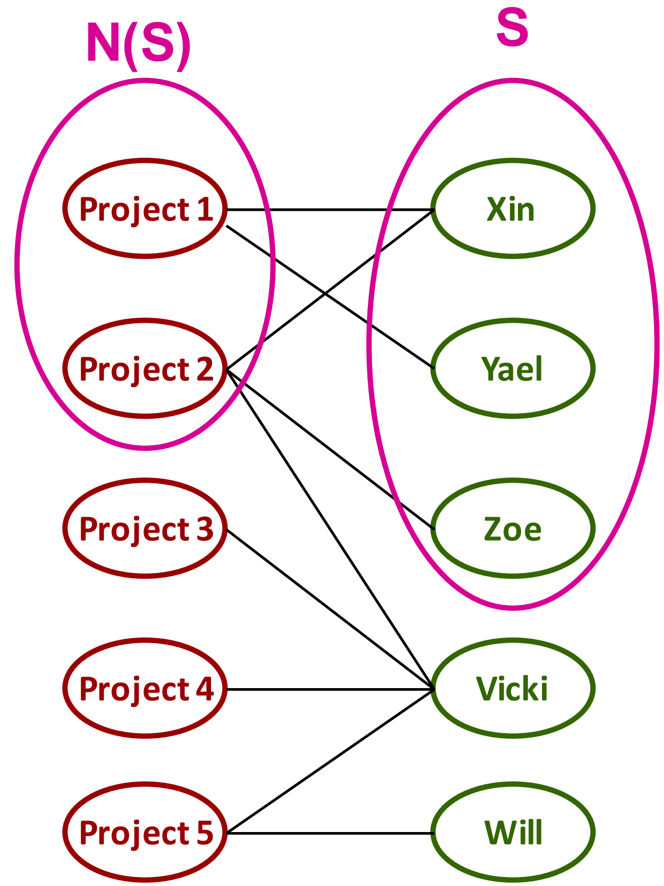

- New graph: it has no perfect matching
- Why?
  - *X, Y, Z* only collectively satisfied with *P1, P2*: not enough projects for the three of them
- If *S* is a subset of agents, let *neighbor set N(S)* be the set of all items they are connected to
- *Constricted set:* any set *S* whose neighbor set *N(S)* is smaller than *S* itself
- If *G* has a constricted set, then obviously there is no perfect matching
- **Matching Theorem:** If *G* has no perfect matching, then it must have a constricted set
- We'll look briefly at an algorithm for constructing matchings later
  - if it does not find a perfect matching, it will identify a constricted set

More General Preferences

- Listing satisfactory items not very "expressive"
  - so use more powerful preferences
- **Ranking (ordinal) models**
  - let agents *rank* all items
    - *e.g., from most to least preferred*
    - *we'll return to this type of model*
- **Valuation models** *(as in auctions)*
  - let agents express *valuations*
    - *i.e., willingness to pay (like auctions)*
    - *captures strength of preference*
    - *we assume all valuations ≥ 0*
  - goal: find matching that maximizes social welfare (total valuation)

Ranking Model of Preferences

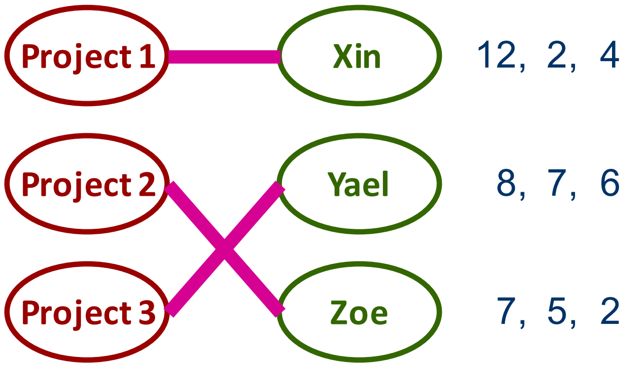

Valuation Model of Preferences

Maximum Weight Matchings

- With valuations, can view as a *complete* bipartite graph with edge weights: edge weight denotes valuation if that pair is matched
  - We put in edges with weight 0 if graph is to be complete
    - Could leave them out if we wanted (see next slide)
- Want to find matching with *maximum weight* (max sum of valuations, i.e., social welfare); how many possible perfect matchings? Cannot in general enumerate.

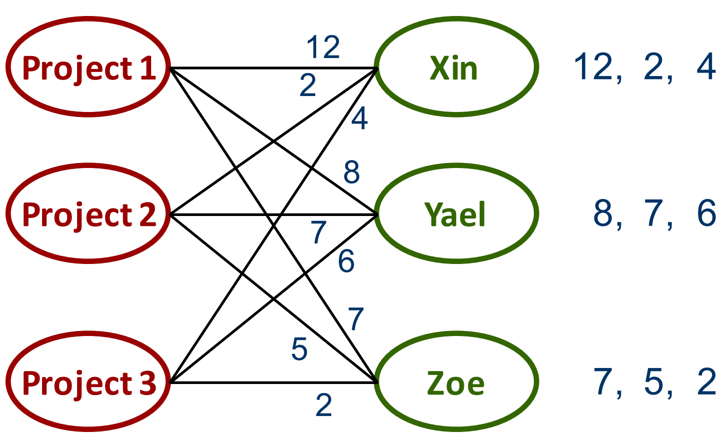

Social welfare of some matchings:

- X-P1, Y-P2, Z-P3: 12+7+2 = 21
- X-P2, Y-P3, Z-P1: 2+6+7 = 15
- X-P3, Y-P1, Z-P2: 4+8+5 = 17
- *X-P1, Y-P3, Z-P2: 12+6+5 = 23*

Max Weight Matchings Need not be "Perfect"

- Without loss of generality we can define max weight matchings on *complete* bipartite graph, we know that perfect matchings exist
  - any assignment of some item to each agent is perfect
- But maximum weight matching may select edges with weight zero
  - i.e., if we deleted weight-zero edges, the max weight matching may not be perfect (some agents will be "shut out"); but this doesn't change the weight of the matching (the social welfare)

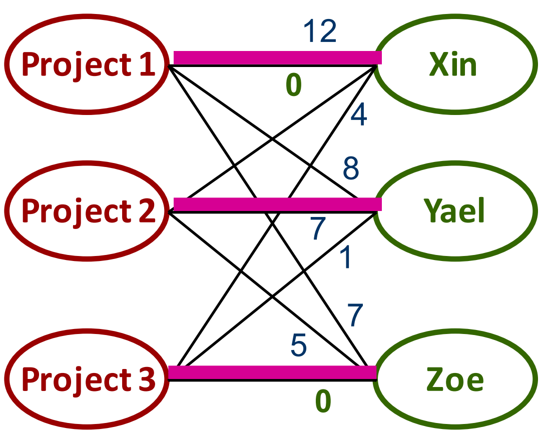

12, **0**, 4

8, 7, **1**

7, 5, **0**

With 0-edges shown

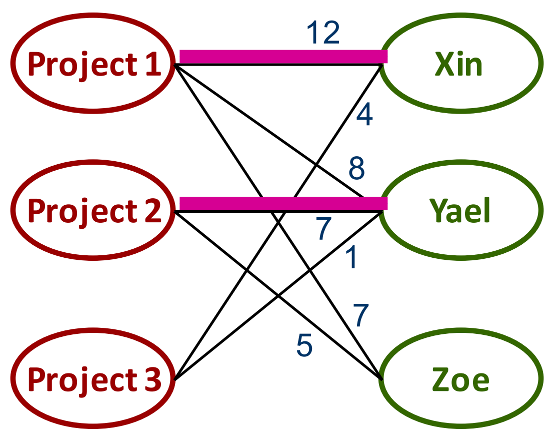

Without 0-edges

Social welfare of some matchings:

- Z-P2, X-P1, Y-P3: 18
- Z-P1, X-P3, Y-P2: 18
- *X-P1, Y-P2: 19*

Max Weight Matchings: Simple Prefs

- Notice that the simple preference model (items are satisfactory or not) is just a special case:
  - put valuation 1 on satisfactory items
  - put valuation 0 on unsatisfactory items
  - leave out 0-weight edges when drawing graph
  - max weight matching assigns items to as many agents as possible
    - thus, if *perfect* matching exists, it will have maximum weight

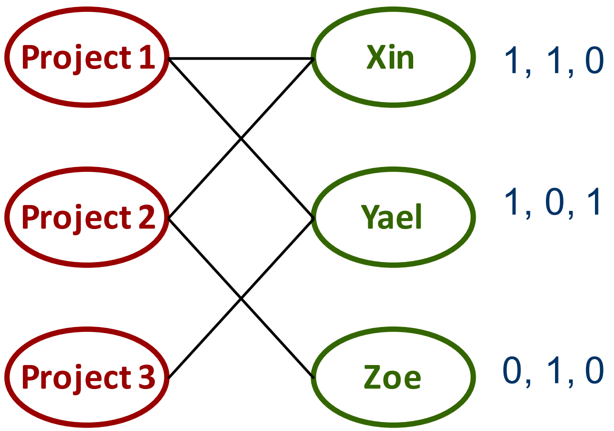

Prices

- Suppose we ask for valuations and then assign to maximize SW
- We can't expect people to reveal their true valuations
  - *e.g., Y likes P1 most, so why doesn't she overstate her value?*
  - *e.g., Z gets P2, but Y likes P2 more than Z; why wouldn't Y offer to pay more?*

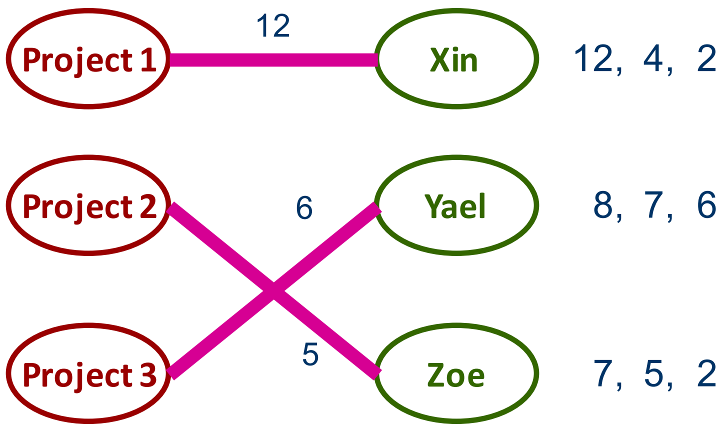

- Just as in single item auctions, prices (charging people) is a means of ensuring that agents "who value items most" get them
  - to maximize SW, we must view "collective value"
- Faculty don't take bribes to supervise projects, so let's sell condos
  - can (decentralized) sellers set prices in such a way that efficient (SW maximizing) outcomes emerge?
  - contrast with VCG: not a single centralized seller trying to elicit valuations; instead it's a free market with sellers setting prices

Prices and Utilities (Payoffs)

- Model is simple
  - agent x has valuation $v_{x,y}$ for item y
  - seller sets price p_y for item y
  - if x buys y, utility is $u(x,y,p) = v_{x,y} - p_y$; but x can buy nothing ($u = 0$)
- Agents clearly only interested in items that maximize their utility
  - once prices p are set, create new *preferred seller graph (PSG)*
  - edges join agents to items that maximize utility: only "satisfactory" items

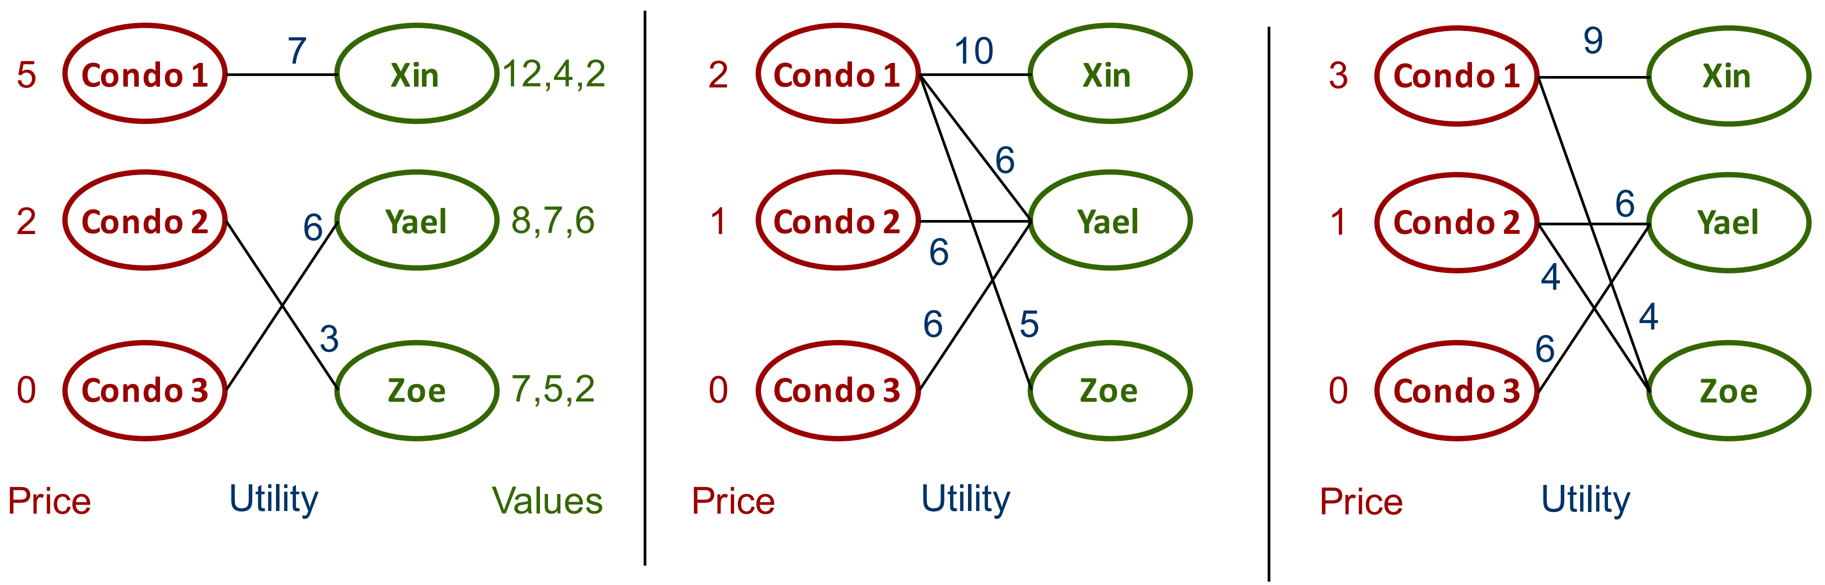

Market Clearing Prices

- Set of prices is *market clearing* iff they allow each item to be purchased (and in this auction, by a different agent)
    - agent must be *satisfied* with item at price paid (not "envious")
- Same as requiring *perfect matching in preferred seller graph (PSG)*

Prices are market clearing:
- *has a perfect matching*
- *no agent prefers different item*

Prices are ***not*** market clearing:
- *no perfect matching*
- *X, Z both want C1*

Prices ***are*** market clearing:
- *has a perfect matching*
- *X gets C1, Z C2, Y C3*

Notice raising price of C1 pushed Y away (reduced contention); ***and made C2 attractive to Z***

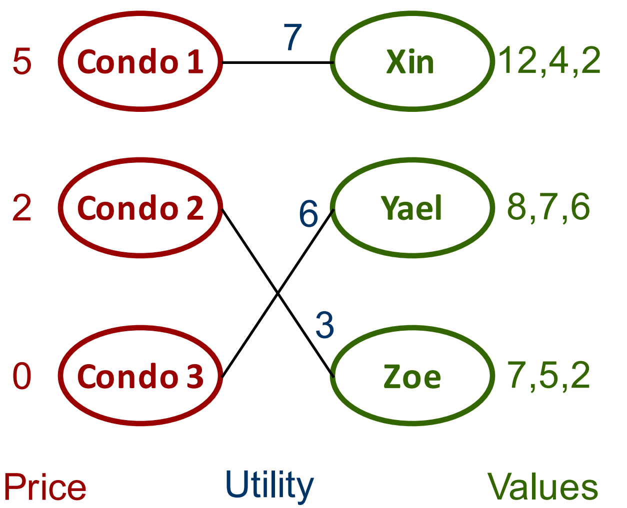

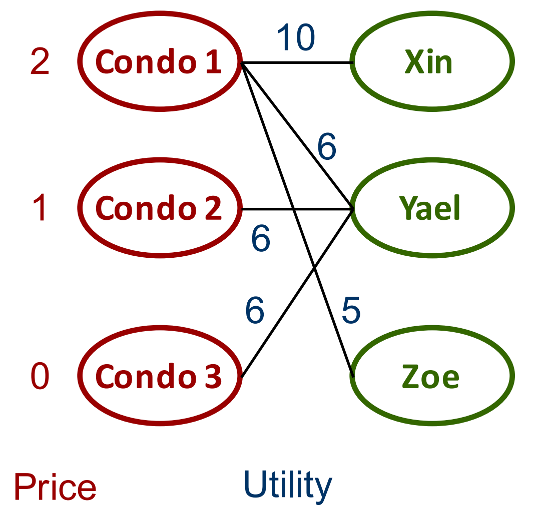

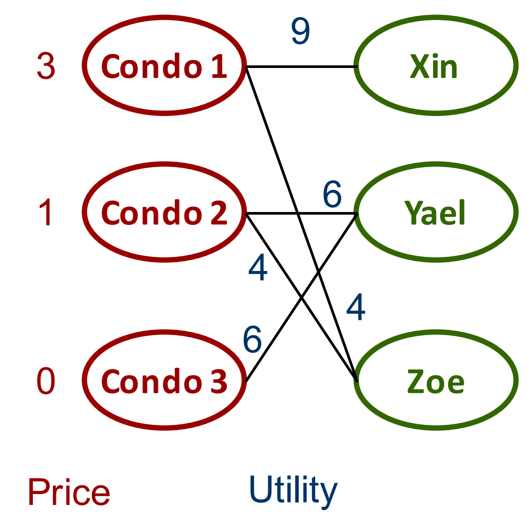

Some Remarkable Properties

- **Existence of Market Clearing Prices:** For any set of buyer valuations, there is a set of market clearing prices
  - notice: model assumes sellers have no values (reserve prices)
  - we'll discuss in shortly how to compute clearing prices, in a way that mimics how sellers might gradually adjust their prices in an open market

- This means a (somewhat) decentralized market will find a perfect matching; but will it be a good one? We've set prices so buyers are individually happy, but collectively? Will we maximize SW?

- **Optimality of Matching at Market Clearing Prices:** For any set of market clearing prices, any perfect matching maximizes the total valuation of items assigned to buyers

Why Matching is Socially Optimal (1)

- Suppose we have market clearing prices p
  - *This means there is a perfect matching in the PSG*
- Now consider the complete graph (each buyer, seller connected)
  - *Note: perfect matching in PSG is also a perfect matching in complete graph*
- In the perfect matching, each buyer x matched to some item y that maximizes her utility $v_{x,y} - p_y$ at the clearing prices p
  - this is quite obvious, since each edge for x in PSG has same utility, and any other edges we added back have lower utility
- So the perfect matching maximizes total utility (*TU*) to all buyers at the clearing prices, *even in the complete graph*

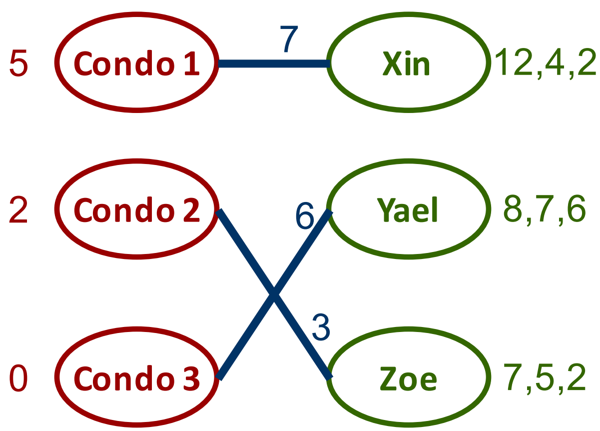

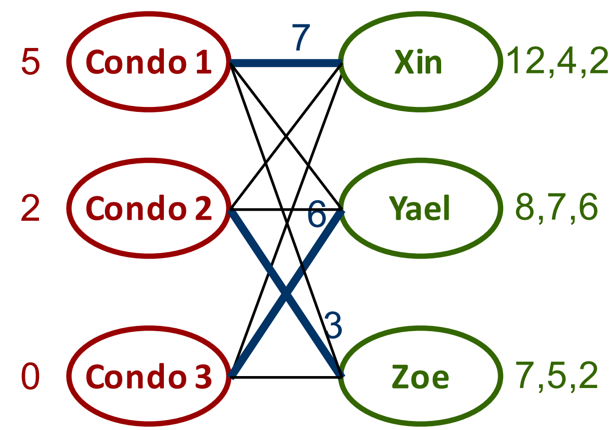

Why Matching is Socially Optimal (2)

- Now total utility is simply the total valuation of items to the matched buyers (*TV*) minus the total price paid by all buyers (*TP*)
  - in other words, *TU* = *TV* – *TP*
- But since prices *p* are fixed, any perfect matching—even in the complete graph, not just the PSG—has *same total price TP*
  - every item is sold, for same price, regardless of matching
- Since all perfect matchings have same *TP*, the matching that maximizes *TU must maximize total valuation TV*

Can Sellers Charge Higher Prices?

- Model ignores self-interest of sellers
  - adding seller reserve values is easy (but market may not "clear")
  - market-clearing prices not unique (as we already saw)
- Given market clearing, can clear at higher prices
  - e.g., C1 can raise price… if so, how high?
    - up to 10 (at 10, both C2 and C3 appeal to *X*)
  - could *all* sellers raise price by 1? 2? 3? 4?
    - PSG stays the same with any *uniform* price increase UNTIL
    - … *all items* too expensive for a buyer B (then B loses *all* edges: unmatched)

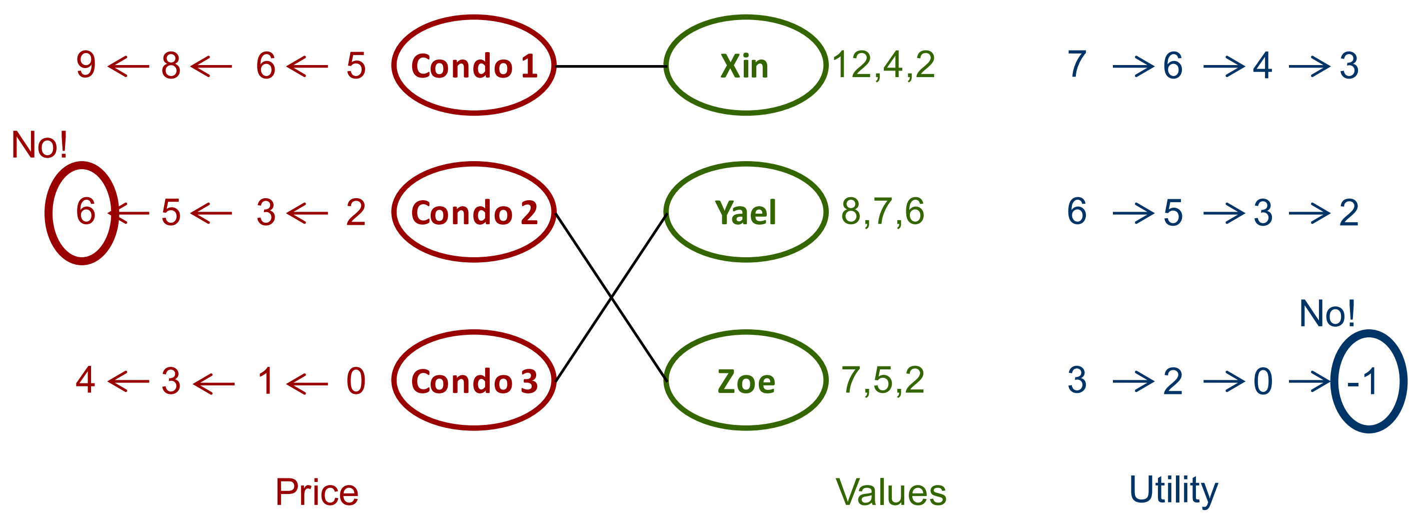

Illustration of Social Optimality

- Since Yael likes *C2* more than Zoe, why can't market clearing prices give *C2* to Yael (and leave Zoe with *C3*)?
- We'd need prices *p* for *C2*, *q* for *C3* such that:
  - $7 - p \geq 6 - q$ (so *Y* likes *C2* as much as *C3*)
  - $2 - q \geq 5 - p$ (so *Z* likes *C3* as much as *C2*)
  - not possible, it implies: $q \geq p - 1$; $p \geq q + 3$
- *C2* and *C3* are in competition for *Y* and *Z*
  - since *Y* likes *C2* and *C3* almost the same, their prices can't be very different
  - *Z*'s valuation for *C2* limits how much *C2* can charge
  - models of competition in networks are discussed in Ch.11 (we won't go into it)

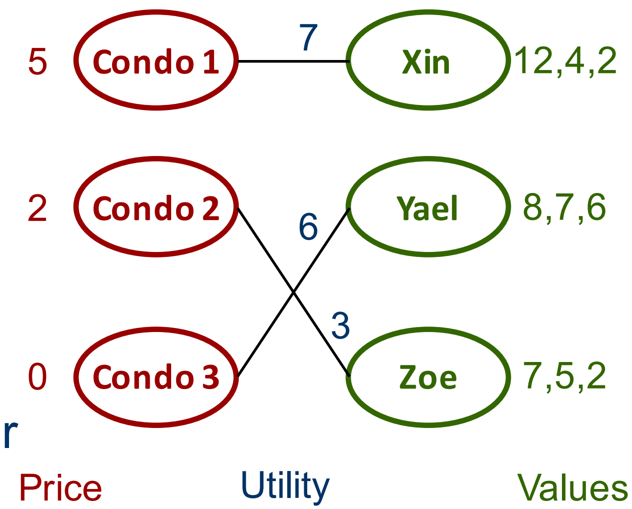

Computing Market Clearing Prices

- Why do market clearing prices always exist?
- We'll construct an *algorithm* that computes market clearing prices [Demange, Gale, Sotomayor, 1986]
- It will be a type of "ascending auction" for multiple items
  - sellers gradually raise their prices (starting from zero)
  - buyers indicate willingness to buy at current prices (pref. seller graph)
  - if any set of items is "overdemanded" (i.e., constricted set), the sellers of those items can raise their prices
    - forces some buyers to consider other items (reduce contention)
  - we'll continue until a perfect matching is found in the PSG
  - we'll add a simple price adjustment step (gives us *minimal* prices)
- It is really nothing more than a multi-item auction (but with different sellers *coordinating* their price adjustments explicitly)

A Simple, Coordinated Multi-item Auction

0. Each seller sets initial price of zero
1. Construct the preferred seller graph (PSG) using current prices
   - *buyers valuations may be known, or they bid for all preferred items*
2. Check if there is a perfect matching in this PSG
   a. If so, **stop**: current prices are market clearing (and matching is one reasonable assignment of items to buyers)
   b. If not, continue to next step
3. Find a constricted set of buyers *S*, and neighboring sellers *N(S)*
   - *in fact, we would like to find a minimal constricted set*
4. Raise price of each item in *N(S)*, *and only items in N(S),* by one
5. If price of each item is greater than zero, reduce each price by the amount of the lowest price (i.e., make lowest price zero)

Aside: I don't think this step is necessary since lowest price never more than 0 (!)

6. Return to Step 1 and repeat

Rough Intuitions

- Steps 1 and 2:
  - at current prices, if perfect matching is found, then we obviously have market clearing prices
- Step 3:
  - if no perfect matching, must have a constricted set: need to fix it
  - in economic terms, the set *N(S)* is *overdemanded*
- Step 4
  - if *S* is a constricted set, there is too much competition for *N(S)* at the current prices for the item in *N(S)*
  - by raising prices of each item in *N(S)*, we may make items outside of *N(S)* more attractive to some buyers in *S* (eventually this must happen!)
- Step 5
  - price adjustment ensures lowest price is always zero
    - *Aside: I don't think it will ever "be triggered"*
  - price adjustment does not affect the PSG at all (if *y* preferred by *x* before, still preferred when all prices decreased uniformly)

Illustration

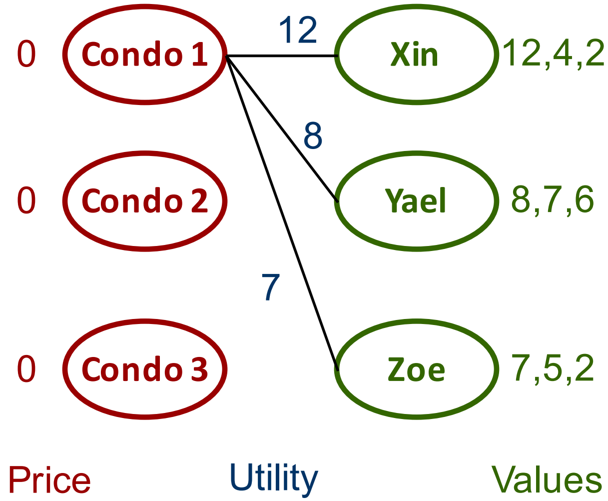

Increase Price of C1

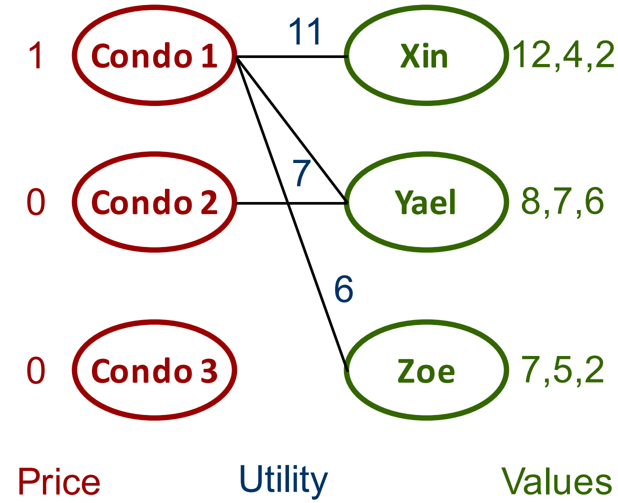

Increase Price of C1 *(or C1, C2)*

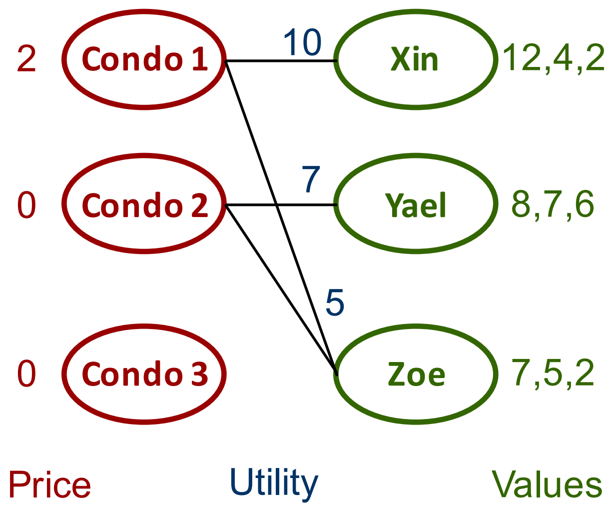

Increase Price of C1, C2

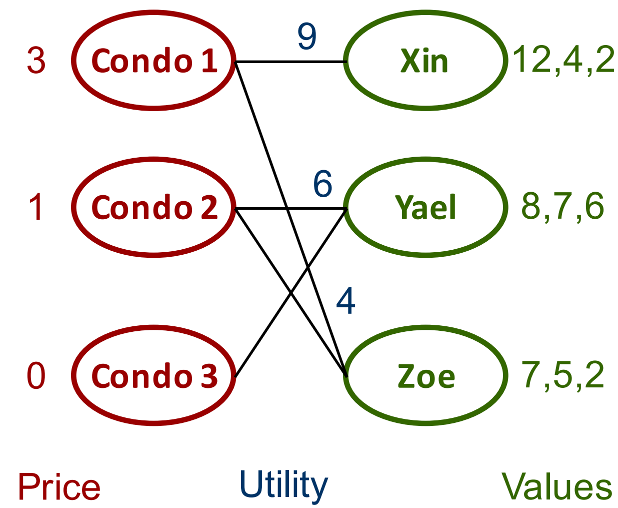

Market Clears