Truthful, Pareto Optimal Exchange Without Money

This lecture begins the second half of the course: up until now, we have studied the behavior of individuals in *already defined games* – i.e. we have been given a game, and have thought about way to predict what might happen when rational agents play the game. Starting today, we ask the reverse question: given how we believe rational agents will behave in an interaction, how can we *design the rules of the game* so as to get good outcomes of play? (And what is a "good" outcome?)

We begin our study with the classical "House Allocation Problem", studied by Shapley and Scarf (the "Top Trading Cycles" algorithm we will give is attributed to David Gale). In this problem, each individual comes to market with a single indivisible item, together with a strict preference ordering over all of the goods in the market. The question is how can we both:

1. Coordinate an exchange to arrive at a *good* allocation, and

2. Do so in a way such that it is a dominant strategy for everyone to report their true preferences.

Note that it is important that we accomplish both item 1 and 2. If we merely give an algorithm for finding a good allocation (more on what that means later) that does not incentivize individuals to report their preferences, then all we will know is that we have found a good allocation with respect to the reported preferences, which may have little bearing on the true preferences in the market.

It will turn out that we can solve this problem without resorting to having to use money (as we will do later in the class). This is useful, because for some problems, money is not practical (or legal) to use. A key example is the problem of *kidney exchange*. Patients who need kidney transplants come to market together with donors who have an incompatible blood type (and so cannot donate to the patient they know). Patients have preferences over kidneys as a function of blood and protein compatibility, and a good exchange will allow us to perform lots of kidney transplants. However, buying and selling organs is illegal in the US, and so a solution that does not require money is important.

Lets start with the model:

1. There are n agents $i = 1, \ldots, n$ who each come to market with a good h_i.

2. Each agent has a strict preference ordering over all of the goods $h_1, \ldots, h_n$, i.e. a list of goods from most preferred to least. The preference ordering for i is denoted $\succ_i$ and used as follows: given a pair j, k, either $h_j \succ_i h_k$ (if i prefers h_j over h_k) or $h_k \succ_i h_j$ (if i prefers to get h_k over h_j). The ordering is *strict* because only one of these two can be true (i cannot be indifferent). The ordering is transitive – so each agent just has a rank order list of goods. In particular, this ranking includes an agents own good h_i. We use $\succeq_i$ for a "weak" preference ordering. For example, we may write $h_i \succeq_i h_i$, which says "i weakly prefers h_i to h_i (which is true).

We wish to design an algorithm which will induce a game played by the players. The algorithm will take as input the reported preferences $\succ_i$ of each of the players, and output an assignment of the goods μ, where $\mu(i)$ is the good assigned to i. This induces a game in which the strategy space for each player is the set of preference orderings $\succ_i$ that they can report, and the utility function is defined by their true preference for the good they receive.

First, lets define what we mean by a good allocation.

Definition 1 *An allocation ν **Pareto-dominates** an allocation μ if, for every i:*

$$\nu(i) \succeq_i \mu(i)$$

and for some j;

$$\nu(j) \succ_j \mu(j)$$

*i.e. everybody is at least as happy with their allocation in ν as with μ, and at least one person is strictly happier. If there is no ν that Pareto-dominates μ, then we say μ is **Pareto optimal**.*

If μ is Pareto-optimal, then any attempt to improve on it (i.e. reassign items) that makes anyone better off will make at least one person worse off.

We want to find Pareto-optimal allocations. We would also like to find them using mechanisms that incentivize agents to report their true preferences. As a very minimal goal, they should not be able to come to harm by participating (truthfully) in the mechanism:

Definition 2 *A mechanism A is **individually rational** if for every player i, every preference order $\succ_i$, and every set of reports of the other players $\succ'_{-i}$, if $\mu = A(\succ_i, \succ'_{-i})$ then:*

$$\mu(i) \succeq_i h_i$$

i.e. the good that player i gets is at least as good (according to her own preferences) as the good that she started with.

But we want more than that:

Definition 3 *A mechanism A is dominant-strategy incentive compatible if it is a dominant strategy for everyone to report their true preferences. i.e. if for all $\succ_i$, all $\succ'_{-i}$, and $\succ'_i$, if*

$$\mu = A(\succ_i, \succ'_{-i}) \quad and \quad \mu' = A(\succ'_i, \succ'_{-i})$$

then $\mu(i) \succeq_i \mu'(i)$.

We will give a surprisingly simple and intuitive algorithm that will simultaneously achieve these two goals, the "Top Trading Cycles" algorithm.

Algorithm 1 The top trading cycles algorithm

$\mathbf{TTC}(\succ_1, \ldots, \succ_n)$

  Let $S_1 = P$ be the set of all agents. Set a counter $t = 1$.

  while $|S_i| > 0$ **do**

    Construct a graph $G_t = (V_t, E_t)$ where $V_t = S_t$ and for each $i, j \in V_t$, the directed edge $(i, j) \in E_t$ if and only if $h_j \succ_i h_k$ for all other $k \in V_t$. i.e. this is the graph that results when every agent "points to" their favorite remaining good.

    Find any cycle C_t in G_t and clear all trades along it: i.e. for every directed edge $(i, j) \in C_t$ set $\mu(i) = j$.

    Set $S_{t+1} = S_t$ and remove all cleared agents: for each $i : (i, j) \in C_t$, set $S_{t+1} \leftarrow S_{t+1} - \{i\}$. Increment t ($t \leftarrow t + 1$).

  end while

  Output μ.

We first observe that the algorithm indeed halts and outputs *some* allocation. Note that so long as we find a cycle at every round, at least one agent is removed from consideration at every round, and the algorithm halts after at most n steps. So it suffices to observe:

Lemma 4 *In each graph G_t constructed by the algorithm, there is at least one cycle C_t, and every agent is part of at most one cycle.*

Proof   This follows simply because by construction, G_t is a directed graph in which every vertex has out-degree exactly one. So by starting at any vertex and following edges forward, we must find a cycle – we cannot get "stuck", since there is always an outgoing edge, and hence must repeatedly visit some vertex. ∎

Lets do an example – consider 5 goods and 5 agents. Agent 1 initially has good A, while 2 has B, 3 has C, 4 has D, and 5 has E. The agents have the following preference ordering over items:

$$\succ_1: B \succ E \succ C \succ A \succ D$$
$$\succ_2: C \succ A \succ E \succ D \succ B$$
$$\succ_3: A \succ B \succ C \succ D \succ E$$
$$\succ_4: A \succ C \succ E \succ D \succ B$$
$$\succ_5: D \succ A \succ C \succ B \succ E$$

Its worthwhile to draw out the set of graphs and cycles the algorithm finds. The resulting matching is:

$$\mu(1) = B, \mu(2) = C, \mu(3) = A, \mu(4) = E, \mu(5) = D$$

Is this allocation any good? We will show that Top Trading Cycles always finds a Pareto optimal allocation.

Theorem 5 *The Top Trading Cycles algorithm produces a Pareto optimal allocation μ on every input.*

Proof   Suppose not. In that case, there is some other allocation ν that Pareto dominates μ. Lets think what ν must look like.

First, observe that every agent that was in cycle C_1 must receive an identical allocation in ν: since these agents are receiving their first choice good in μ, and must do at least as well in ν.

Next note that every agent in cycle C_2 must receive an identical allocation in ν: since these agents are receiving their first choice good from the set $P - C_1$ in μ, and $\nu(i) = \mu(i)$ for every $i \in C_1$, they can't do better in ν, and so must do identically.

We can continue by induction. We claim inductively that $\nu(i) = \mu(i)$ for every $i \in C_1 \cup \ldots \cup C_t$, and have shown the base case above. We must also have that $\nu(i) = \mu(i)$ for every $i \in C_{t+1}$ since every such agent is obtaining their favorite item in μ among the set $P \setminus C_1 \cup \ldots \cup C_t$, and by the inductive hypothesis, they can't do better in ν. Continuing through $t = n$, we see that it must be that $\mu = \nu$, a contradiction. ∎

Its not hard to see that top trading cycles is individually rational: truthful agents are only ever allocated goods that are their favorite among the remaining goods, and that includes an agent i's own good (which is always available to agent i). So it is not possible to allocate to any agent a good that she prefers less than her own.

Finally, we (sketch) the proof showing that truthful reporting of their preferences $\succ_i$ is a dominant strategy for all agents i.

Theorem 6 *The Top Trading Cycles Algorithm is Dominant Strategy Incentive Compatible.*

Proof [Sketch] We think about the algorithm as if player i can "decide" where to point in the construction of graph G_t at each round t, as a function of where everyone else is pointing. We conclude that it is always in player i's best interest to point to his favorite good among the ones remaining. Since this is exactly how the TTC algorithm constructs G_t assuming player i has reported his true preference ordering $\succ_i$, the theorem follows.

At round t, why would player i want to not point to his most preferred good? This could only be because if he pointed to a less preferred good, he would receive it, and the opportunity to receive it later would disappear (resulting in him eventually receiving an even less preferred good). But this can't happen:

Fixing the edges of the other players, consider the set of goods that agent i can get today if he points to them. These are all of the goods that form paths leading to agent i (so that they would form cycles if agent i pointed at them). Call this "agent i's choice set" at round t. But note that agent i's choice set can only *increase* and not *decrease* if agent i is not matched at round t. It can increase because agents who were not previously pointing at agent i might now begin to, once the goods they were previously pointing to were removed, leading to more goods forming chains into agent i. It cannot decrease, because since agent i was not matched, no good forming a chain into agent i was removed, and all of those edges remain at time $t + 1$. Hence there is no opportunity cost for agent i to point to his most preferred good at each time step. ■