Game Theory and Algorithms*
Lecture 3: Weak Dominance and Truthfulness

March 1, 2011

Summary: We introduce the notion of a *(weakly) dominant strategy*: one which is always a best response, no matter what the opponents play. If every player has a weakly dominant strategy, they form a Nash equilibrium. One example we give is Hotelling's classic example of voting. We then give two other examples (group decision-making, auctions) where the profile of weakly dominant strategies is also *truthful*, which has the advantage that despite lots of unknown information, each person can deduce their weakly dominant strategy, and it corresponds to playing their private "true" value.

1 Weak Dominance

Previously we defined when one strategy *strictly dominates* another. We now define a weaker[1] notion:

Definition 1. For player i, strategy a_i *weakly dominates* strategy a'_i if for all partial action profiles a_{-i} of the other players, $u_i(a_i, a_{-i}) \geq u_i(a'_i, a_{-i})$.

Equivalently, a_i weakly dominates a'_i when playing a_i is never worse than playing a'_i. We remark that weak dominance does not enjoy the same properties as strict dominance: for example, iterated elimination of weakly dominated strategies can delete a Nash equilibrium.

The most important notion related to weak dominance is the following:

Definition 2. For player i, strategy a_i is *(weakly) dominant* if it weakly dominates all other strategies of player i.

Equivalently, a weakly dominant strategy is a best response no matter what your opponents do. Consequently, we get the following fact.

Proposition 3. *If each player i has a weakly dominant strategy a_i, then $(a_1, a_2, \ldots, a_n)$ is a Nash equilibrium.*

* Lecture Notes for a course given by David Pritchard at EPFL, Lausanne.

[1] Aside: some sources also require "for at least one partial action profile a_{-i} of the other players, $u_i(a_i, a_{-i}) > u_i(a'_i, a_{-i})$." (It prevents two strategies from weakly dominating one another.)

2 Hotelling's Election Model

In this 1929 example of Harold Hotelling, we have a game where the players are candidates in an election. The general idea is that they are deciding their campaign platforms on a one-dimensional scale (e.g., they can be 100% liberal, 100% conservative, or any combination of the two). Intuitively, there are an infinite number of voters; each voter has their own preference, and will vote for the candidate whose campaign platform is "closest" to their preference.

We now make this idea precise. We will model the space of possible campaign platforms as the unit closed interval $[0, 1]$. The sets of possible strategies are thus $A_i = [0, 1], i = 1, \dots, n$.

We imagine the voters' preferences as uniformly distributed[2] over the interval $[0, 1]$. Each infinitesimal element $x \in [0, 1]$ of the voters votes for the candidate i for whom $|a_i - x|$ is minimal; in the case there are several, the votes are split equally. We give some example calculations of the votes in this model:

- The action profile $(0, 0.4, 0.6)$ splits the voters into three intervals: those in $[0, 0.2)$ vote for player 1, those in $(0.2, 0.5)$ vote for player 2, and those in $(0.5, 1]$ vote for player 3. So the vote distribution is $v = (0.2, 0.3, 0.5)$.

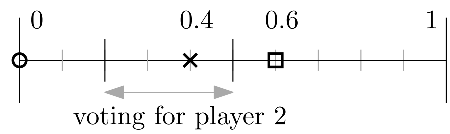

- The action profile $(0, 0.6, 0.6)$ splits the voters into two intervals: those in $[0, 0.3)$ vote for player 1, and those in $(0.3, 1]$ split their votes between players 2 and 3. So the vote distribution is $(0.3, 0.35, 0.35)$.

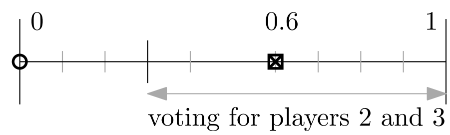

Finally, the utility for player i is determined in a way so that losing is worse than winning, but tying is worse than winning alone:

$$u_i = \begin{cases} 0, & \text{if player } i \text{ does not win (get the most votes);} \\ 1/k, & \text{if player } i \text{ wins in a } k\text{-way tie.} \end{cases}$$

2.1 Analysis

Claim 4. *If there are two players, for each candidate, choosing the action 0.5 is a (weakly) dominant strategy.*

Proof. Fix the action a_1 of player 1; we will prove that 0.5 is always a best response for player 2.

Case 1 : $a_1 = 0.5$. In this case, we claim the *only* best response of player 2 is 0.5. To see this, first look at what happens in the action profile (0.5, 0.5): it is a tie vote. Second, any other a_2, WOLOG $0 \leq a_2 < 0.5$, gives player 1 the interval $((0.5 + a_2)/2, 1]$ which is more than half of the votes, causing a_2 to lose. So the claim is true.

[2]This actually follows WOLOG by "rescaling" from a weak assumption on the voters, namely that their distribution has no "atoms" and no "holes."

Case 2 : $a_1 \neq 0.5$, WOLOG $0 \leq a_1 < 0.5$. This is the same as the above calculation: in the action profile $(a_1, 0.5)$, player 2 wins the vote alone, so 0.5 is a best response. □

Consequently (by Proposition 3) we see that $(0.5, 0.5)$ is a Nash equilibrium.

Exercise. Show that for two players, there are no other Nash equilibria in this game.

Thus, if you are willing to accept the axioms of the model, this says that in a two-party electoral system, both parties have an optimal (dominant) choice chosen by setting their policy to that of the "middle" voters, and there is no other stable situation. In some cases this agrees with experience, that there is no practical difference between the candidates.

Here is a question just for fun: what happens if the voters are not distributed on a line segment, but instead in a disc, or in a (Borg) cube?

Exercise. For three players (on $[0, 1]$), show that there is no Nash equilibrium of the form (x, x, x); find a Nash equilibrium of the form (x, y, y).

3 Truthfulness: Collective Decision-Making

We now give another setting with a dominant strategy equilibrium. It models the situation where multiple people need to combine their opinions to make a single joint decision. We have n players, and each one has a different opinion t_i on a single axis, modelled by real numbers as in Hotelling's game. However, players are allowed to lie: so while their action set is $\mathbb{R}$, it is not necessarily true that $a_i = t_i$.

We will actually describe several different games, which we will also call *mechanisms*, for picking a policy $p(a_1, a_2, \ldots, a_n)$ depending on the player's actions. The utility for a player decreases with the distance between the chosen policy and their true opinion:

$$u_i(a) = -|p(a) - t_i|.$$

A first likely candidate is to pick p to be the *mean* or *average*, i.e., we first consider defining $p(a) = (a_1 + a_2 + \cdots + a_n)/n$. However, this doesn't work very well:

Proposition 5. *When p is the mean, the collective decision-making game has no Nash equilibrium (unless all players have the same t_i value).*

Proof. Consider any action profile a and two players i and j with $t_i \neq t_j$. Since it is impossible for $p(a)$ to simultaneously equal both t_i and t_j, it must be that in a, at least one player has negative utility: say, $p(a) \neq t_i$ and so $u_i(a) < 0$. The crux is that player i can deviate in such a way that their utility increases: by picking

$$a'_i = nt_i - \sum_{k \neq i} a_k$$

we find that the mean of (a'_i, a_{-i}) equals t_i, the true opinion of player i, and so this deviation increased player i's utility from negative to 0. Consequently a was not a Nash equilibrium. □

More generally, this is a very impractical game to play: you can't really make any sensible judgement about what is a "good" action without knowing your opponents' future choices; and when the game is played, it may not reveal much about the players' true opinions.

3.1 The Median Mechanism

The above problem gives our first experience with *mechanism design*: we'd much prefer a rule p for choosing the group policy such that individuals have incentive to tell their true opinion. We will give an example of a p where the "true" action is weakly dominant for each player. Thus with this p each player can choose an optimal action without knowing anything about the other players (neither their t_i nor the a_i they will choose); and the person running the game has some confidence that they are eliciting the true opinions of the players.

Finding games where the truthful action is weakly dominant is a central theme in game theory and algorithmic game theory. Some terms for this type of concept are *incentive compatible* and *truthfully implementable in dominant strategies*. Here we'll call it a *truthful mechanism.*

We will assume the number of players is odd. Then to get our truthful mechanism define $p(a_1, a_2, \dots, a_n)$ to be the *median* of the a_i (sort these numbers, then the median is the $\frac{n+1}{2}$th one, i.e. the in the middle position). E.g., the median of $(7, 9, 3, 0, 9)$ is 7.

Theorem 6. *In the policy-setting game, if p is the median, then telling the truth is a dominant strategy for each player.*

Proof. Fix a player i and actions a_{-i} of the other players. We need to show that $a_i = t_i$ is a best response for player i. This entails determining $u_i(a_i, a_{-i})$ for each $x \in \mathbb{R}$, and in turn this requires computing the median (a_i, a_{-i}) as a_i varies.

How does the median of (a_i, a_{-i}) vary with a_i, when the a_{-i} are fixed? Let ℓ (resp. r) be the value in a_{-i} which is $\frac{n-1}{2}$th smallest (resp. $\frac{n+1}{2}$th smallest). Then the median is

$$p(a_i, a_{-i}) = \begin{cases} \ell, & \text{if } a_i < \ell; \\ a_i, & \text{if } \ell \le a_i \le r; \\ r, & \text{if } a_i > r. \end{cases}$$

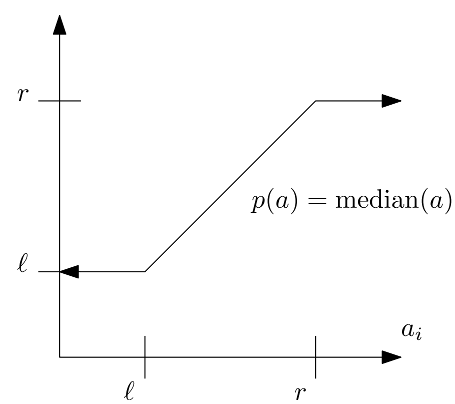

To check if the mechanism is truthful, we need to see if the choice $a_i = t_i$ maximizes u_i, i.e. if $a_i = t_i$ gets $p(a_i, a_{-i})$ as close to t_i as possible.

- If $t_i < \ell$, then the closest that p can get to t_i is if $p = \ell$; indeed, $a_i = t_i < \ell$ achieves this optimum.

- The case $t_i > r$ is analogous.
- If $\ell \leq t_i \leq \ell$, then the choice $a_i = t_i$ makes $p(a) = a_i = t_i$ and $u_i(a) = 0$, the maximum possible.

Thus, $a_i = t_i$ is indeed a dominant strategy. $\square$

In sum, the game is truthful. One consequence is that $(t_1, \ldots, t_n)$ is a Nash equilibrium. Note, however, that there are other Nash equilibria, such as $(13, 13, \ldots, 13)$. We mention one other truthful mechanism below.

Exercise. Show that the following *randomized* mechanism is *truthful in expectation* (i.e., if each player wants to maximize the expected value of their utility, then $a_i = t_i$ is a weakly dominant strategy): pick i uniformly at random from $\{1, \ldots, n\}$, and then set $p = a_i$. How does the expected *social welfare* $-\sum_i |t_i - p(t)|$ of this mechanism compare to the median mechanism?

4 Auctions

Auctions form a huge portion of the literature on algorithmic game theory. In this lecture we give only a small introduction; we will give a more thorough treatment later in the course.

In our simple example we have exactly one item to sell, and there are n players who will each make a bid on the item. Thus the action set for each player is $\mathbb{R}_{\geq 0}$. We will consider auctions which give the item to the player whose bid is maximal (using any form of tie-breaking, which we won't specify unless needed). As the auctioneer, we need to figure out what *price* or *payment* the bidder should pay for winning the item.

We model the players' utility functions in the following way: a player who does not win the item gets a utility of zero; and for each player, they have a private *valuation* v_i, and their utility upon winning and paying price p is $v_i - p$. Thus, getting an item and paying v_i is equivalent to not getting the item; paying less/more than v_i is better/worse than losing.

For auctions, we will often use the symbol b_i (bid) instead of a_i (action).

The most obvious type of auction would be to charge the winning player i a price $p = b_i$ equal to their bid. This is called a *first-price auction.*

Exercise. Consider a two-person auction with $v_1 > v_2$. Show that a first-price auction is not truthful, by showing that (v_1, v_2) is not a Nash equilibrium. Show that depending on how we break ties, there may be zero, one, or many NE.

Generally speaking, the problems are similar to what we encountered before: the best response for a player depends on the choices of their opponents.

4.1 Second-Price Auctions

In a second-price auction, we change the amount paid to be the *second-highest* bid. (In the case that the two largest bids are equal, say to x, then the second-highest bid is x.)

For example, suppose the vector of bids is $b = (3, 10, 7)$. Then player 1 has the highest bid and wins the item; their payment would be 7, the second-highest bid. This has the same nice properties as the median mechanism due to the following fact:

Proposition 7. *In a second-price auction, bidding your true valuation ($b_i = v_i$) is a dominant strategy.*

Proof. As with the median mechanism, we must show that for any collection b_{-i} of bids by the other players, v_i maximizes the utility $u_i(b)$ of player i.

Let b_j be any maximal bid amongst the players distinct from i, i.e. $b_j = \max\{b_{-i}\}$. As player i varies their bid, their utility changes as follows.

- When $b_i < b_j$, player i loses the auction and gets a utility of zero.
- When $b_i > b_j$, player i wins the auction and gets a utility of $v_i - b_j$.
- (For $b_i = b_j$, one of the above cases happens depending on the tie-breaking rule.)

(There is an important fact here which will help our intuition later: the price paid by player i does not depend on b_i! All that player i can change is whether she wins the item (and pays b_j) or loses it. This prevents the strategic manipulation which happened in first-price auctions.)

To finish the proof, consider two cases according to how $v_i - b_j$ compares to zero.

1. If $v_i > b_j$, then it is better (optimal) for player i to win the auction. Indeed, bidding the true valuation $b_i = v_i$ causes this.
2. If $v_i < b_j$, then it is better (optimal) for player i to lose the auction. Indeed, bidding the true valuation $b_i = v_i$ causes this.
3. (If $v_i = b_j$, all bids are optimal for player i.)

Thus, bidding the true valuation is weakly dominant. □

There are some other auctions which are essentially second-price auctions in disguise:

- In an *English auction/ascending auction*, bidders sit in a room and every bid is announced to everyone. Each individual can bid 0, 1, or multiple times; each bid must be larger than the previous one. The last bidder wins, and pays an amount equal to their last bid. Although the largest bid is paid (like a first-price auction), in order for you to pay more than the second-highest valuation, one of your opponents would have to risk getting a negative utility.

- eBay uses a proxy system which effectively runs an English auction, but you only need to enter one value and it will automatically place bids for you until nobody else competes, or the competing bids are larger than your value.

Another nice property of second-price auctions is that you cannot improve your situation by introducing a second "meat-puppet" voter to act as another player under your control. However, it is possible for two people to collaborate with a bribe and improve their situation: e.g. if there are two voters with valuations ($3, $2), whereas truthful bids give a utility of ($1, $0), if the voters collude and player 1 pays player 2 a one-dollar bribe to bid $0, then the utilities (including the bribe) become ($2, $1), an improvement for both players. These facets, and other distinct notions like *group-strategyproof* mechanisms, are studied in the literature.

4.2 Coming Up Later

When we return to auctions we will discuss several topics:

- What if we have several copies of an item to sell? What if we have several different items, and the bidders want different subsets of them? These will be handled by the Vickrey-Clark-Groves mechanism, a truthful mechanism which generalizes second-price auctions. There is a lot of literature which then stems from the fact that VCG is not always computationally tractable.

- We will consider the perspective of an auctioneer who wishes to maximize their profit, subject to the constraint of truthfulness. If the customers' preferences are drawn from random distributions, an approach of Myerson solves this problem (so-called *optimal mechanism design*) by adding a *reserve price* to the auction.

- We consider a generalization of the previous problem to when the bidders' distributions are unknown.