Topics on the Border of Economics and Computation	December 18, 2005
Lecture 8	
Lecturer: Noam Nisan	*Scribe: Ofer Dekel*

1 Correlated Equilibrium

In the previous lecture, we introduced the concept of correlated equilibria [2] with an example. We now give a more formal definition.

Definition 1 *A game consists of n players, each one has a set of strategies S_i and a utility function $u_i : S_1 \times \ldots \times S_n \to \mathbb{R}$. A probability distribution P over $S_1 \times \ldots \times S_n$ is a correlated equilibrium if for every strategy $s_i \in S_i$ such that $P(s_i) > 0$, and every alternative strategy $s'_i \in S_i$, it holds that,*

$$\sum_{s_{-i} \in S_{-i}} P(s_i, s_{-i}) u(s_i, s_{-i}) \;\geq\; \sum_{s_{-i} \in S_{-i}} P(s_i, s_{-i}) u(s'_i, s_{-i}) \quad . \tag{1}$$

An alternative formulation of the above is obtained by dividing both sides of this inequality by $P(s_i)$. This gives,

$$\sum_{s_{-i} \in S_{-i}} P(s_{-i}|s_i) u(s_i, s_{-i}) \;\geq\; \sum_{s_{-i} \in S_{-i}} P(s_{-i}|s_i) u(s'_i, s_{-i}) \quad ,$$

which is the expected payoff of player i if he chooses strategy s_i (left-hand side) or s'_i (right-hand side), and if the $n-1$ other players choose their strategy randomly from P.

We interpret P as the probability used by the trusted third-party to choose the strategy which he recommends to the n players. Our implicit assumption is that the $n-1$ other players follow this recommendation, and we ask ourselves whether it is in our best interest to follow the recommendation as well.

The constraints defined in Eq. 1 (for every s_i and s'_i) are linear constraints on P. So are the constraints $P(s_1, \ldots, s_n) \geq 0$ and $\sum P(s_1, \ldots, s_n) = 1$. We conclude that the set of correlated equilibria define a polyhedron, and moreover, that we can find a correlated equilibrium which maximizes any linear function of P using linear programming. Specifically, we can find the correlated equilibrium which maximizes the expected sum of utilities of the players.

We note that the set of Nash equilibria is obtained by intersecting the polyhedron described above with the additional constraint $P(s_1, \ldots, s_n) = P(s_1) \times \ldots \times P(s_n)$.

2 Secure Function Evaluation and Cheap Talk

As noted in the previous lecture, the fact that the trusted third party gives an individual recommendation to each player, and does not reveal his recommendation for s_i to anyone but player i, is a crucial component in the definition of the game. An interesting question is how to implement the ideas of correlated equilibria in a real system, and specifically how to insure that each player is isolated from the recommendations given to the other players.

The solution relies on cryptographic techniques for secure function evaluation. Generally speaking, player A has a value x and player B has a value y. Player A would like to calculate a function $f_A(x, y)$ and player B would like to calculate $f_B(x, y)$, without revealing their respective values to each other. In other words, they would like to evaluate f_A and f_B using a zero knowledge protocol. In our case, player i learns only the recommendation s_i, without learning anything about s_{-i}. If $n \geq 3$, every two players have a private channel between them, and less than $\frac{1}{3}$ of the players are dishonest, then there exist techniques for secure function evaluation which do not rely on cryptographic assumptions.

A related game-theoretic concept is *cheap talk*. Before playing the game, the players take part in a "talking" stage. As its name implies, this pre-game communication carries no costs and the players are not bound by anything said during this stage. For instance, each player can announce the strategy he is about to take, or the strategy he would like the other players to take. Under certain information theoretic assumptions, the cheap talk stage can be used to obtain common random bits.

3 Social Choice

Let A be a set of m alternatives ($m > 3$), and let L be the set of total orders over the elements of A (L is equivalent to S_m, the set of permutations over $1, \ldots, m$). There are n players, each one has a preference $<_i \in L$. Social choice addresses the problem of aggregating the n different preferences into a single output.

Definition 2 *A social choice function is a function $f : L^n \to A$. Namely, f takes the preferences of the n players and chooses a single alternative.*

Definition 3 *A social welfare function is a function $f : L^n \to L$. Namely, f takes the preferences of the n players and outputs a single aggregate preference.*

We use $<_{f(<_1,\ldots,<_n)}$ to denote the output of the social welfare function with respect to inputs $<_1, \ldots, <_n$. When the inputs are obvious from the context, abbreviate this by $<_f$. Every social welfare function naturally defines a social choice function, which simply selects its output to be the highest ranking alternative. For concreteness, we can think of a social choice function as an election, where A is the set of candidates, the players are the voters, and f is the protocol which tallies the votes.

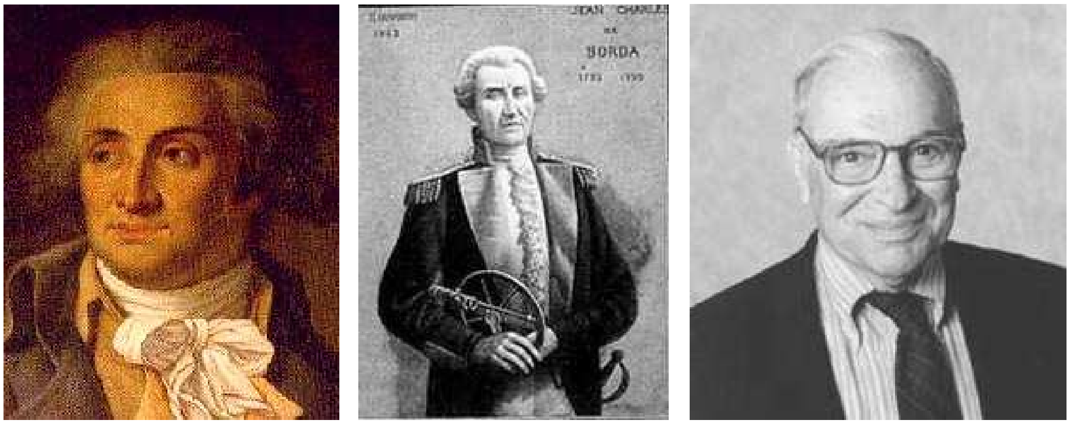

Figure 1: *Left:* Marie Jean Antoine Nicolas Caritat, The Marquis de Condorcet, 17/9/1743 - 28/3/1794. *Middle:* Jean-Charles Chevalier de Borda, 4/5/1733 - 19/2/1799. *Right:* Kenneth Joseph Arrow, 23/8/1921 -

Example - Plurality: The most common social choice function used in elections is the plurality function, which only considers the top-ranked alternative on each ballot, and chooses the candidate who received the most votes.

Example - The Condorcet Paradox: Intuitively, it seems reasonable to require a social welfare function to be consistent with the majority of the voters in the following sense: the social welfare function ranks a above b iff at least half of the voters ranked a above b. In the late 18th century, the Marquis de Condorcet noted a situation where every social welfare function violates this requirement. Here is an example of such a situation. We have three alternatives: a, b, and c.

- 25% of the voters prefer $a > b > c$.
- 35% of the voters prefer $b > c > a$.
- 40% of the voters prefer $c > a > b$.

Alternative a cannot be ranked highest by the social welfare function since 75% of the voters prefer c over a. Alternative b cannot be ranked highest since 65% of the voters prefer a over b. Alternative c cannot be ranked highest since 60% of the voters prefer b over c. In other words, the set of constraints imposed by this requirement leads to cycles.

Example - Borda Count and Independence of Irrelevant Alternatives: The Borda count is a social welfare which functions as follows. For every ballot, the candidate in position i receives $n - i$ points. Then, the points of each candidate are added up across all ballots, and the candidates are sorted in descending order according to their respective points. Ties are broken in some arbitrary predefined way. Compared to the ill-defined method presented in the previous example, the Borda count is well-defined for all inputs. Nevertheless, it suffers from another deficiency, as it violates the *independence of irrelevant alternatives* principle.

Definition 4 *A social welfare function f has the independence of irrelevant alternatives (IIA) property if it satisfies the following constraint. Let a and b be two alternatives in A and let $<_1, \ldots, <_n$ and $<'_1, \ldots, <'_n$ be two sets of preferences such that,*

$$\forall\ i \qquad a <_i b \quad \Longleftrightarrow \quad a <'_i b \ \ .$$

Then $a <_{f(<_1,\ldots,<_n)} b \quad \Longleftrightarrow \quad a <_{f(<'_1,\ldots,<'_n)} b$. In other words, the social preference of a over b depends only on how each voter ranked a with respect to b.

To see why the Borda count violates IIA, take for example the following ballots:

- Preference of voter 1: $b \ <_1 \ c \ <_1 \ a$.
- Preference of voter 2: $c \ <_2 \ a \ <_2 \ b$.

Candidate a receives 3 points, b receives 2 points, c receives 1 point. therefore, the Borda count f ranks the alternatives $c <_f b <_f a$. Now, the two players change their preference to:

- New preference of voter 1: $c \ <'_1 \ b \ <'_1 \ a$.
- New preference of voter 2: $a \ <'_2 \ c \ <'_2 \ b$.

Note that neither player changed his respective ranking of a vs. b. Nevertheless, the output of the social welfare function becomes $c <_f a <_f b$.

4 Arrow's Theorem

The following theorem due to Arrow[1] states that any method that tries to aggregate the preferences of the n players will suffer from one of the problems described above. Moreover, the only method which does not suffer from these problems is one where a single predetermined voter (a dictator) determines the social preference.

Definition 5 *A social welfare function f is unanimity respecting (also called "Pareto efficient") if for every pair of alternatives, a and b, which satisfy $b <_i a$ for all $1 \leq i \leq n$, it holds that $b <_{f(<_1,\ldots,<_n)} a$. In words, if some alternative is unanimously preferred over another by all voters, then it is also preferred by the social choice function.*

Definition 6 *A social welfare function f is called a dictatorship if there exists some i such that $a <_i b \quad \Longleftrightarrow \quad a <_f b$.*

Theorem 7 (Arrow) *Let f be a social welfare function which is: (1) unanimity respecting, and (2) independent of irrelevant alternatives, then f is a dictatorship.*

The proof is given in the next lesson.

Definition 8 *A social choice function f is monotonic if for every $<_1, \ldots, <_n$ and $<_i'$ such that $f(<_i, <_{-i}) = a$, $f(<_i', <_{-i}) = b$, and $a \neq b$ it holds that $b <_i a$ and $a <_i' b$.*

Corollary 9 (Arrow's strong theorem) *If f is a monotonic social choice function onto A then f is a dictatorship.*

The proof of this corollary is also diverted to the next lesson.

5 Manipulation

Arrow's theorem tells us that in order to get a meaningful social welfare function, we must relax at least one of the two requirements in Thm. 1. Since the unanimity respecting property is so elementary, it is common to relax the the IIA property. The disadvantage of this that it can make the voting mechanism "manipulable". That is, there are situations where some of the voters get a better result by mis-reporting their preferences. The voters are now strategic players, in the game-theoretic sense.

Definition 10 *A social choice function f is manipulable if there exist $<_1, \ldots, <_n$ and $<_i'$ such that $f(<_i, <_{-i}) \ <_i \ f(<_i', <_{-i})$.*

We now prove an equivalence between manipulability and non-monotonicity.

Lemma 11 *Let f be a social choice function. f is manipulable iff f is non-monotonic.*

Proof: Assume that f is non-monotonic. Namely, there exist $a, b \in A$, $<_1, \ldots, <_n$ and $<_i'$ such that $f(<_i, <_{-i}) = a$, $f(<_i', <_{-i}) = b$ and $a <_i b$. Then player i can manipulate f by choosing $<_i'$ instead of $<_i$. On the other hand, assume that f is manipulable, then there exist $<_1, \ldots, <_n$ and $<_i'$ such that $f(<_i, <_{-i}) <_i f(<_i', <_{-i})$. Let $a = f(<_i, <_{-i})$ and $b = f(<_i', <_{-i})$, and we have a counterexample which disproves monotonicity. ■

References

[1] K.J. Arrow. *Social Choice and Individual Values.* Yale University Press, second edition, 1962.

[2] R. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1:67–96, 1974.

[3] P.J. Reny. Arrow's theorem and the Gibbard-Satterthwaite theorem: A unified approach. *Economics Letters*, 70:99–105, 2000.