Dirk Bergemann
Joan Feigenbaum
Yale University

Economics and Computation

Computer Science 455/555 and Economics 425/563
Second In-Class Exam
12/04/08

This exam is a closed-book exam. The exam time is 75 minutes. The exam has 100 points and each subquestion has equal points.

1. Consider a sponsored search auction. The click-through rate μ_i is assumed to only depend on the position i of the link and we have

$$\mu_1 \geq \mu_2 \geq \cdots \geq \mu_k > 0.$$

There are $n > k$ competitors and the value of competitor j per click is $v_j > 0$ and for simplicity assume

$$v_1 \geq v_2 \geq \cdots \geq v_n > 0$$

The sponsor (Google, Yahoo, etc.) considers running a first price rather than a second price auction for the sponsored search. The auction format considered therefore ranks the bids b_j such that:

$$b_1 > \cdots > b_n \geq 0,$$

and the equilibrium revenue for each bidder is

$$\mu_i \left(v_i - b_i\right).$$

Consider the complete information environment in which the valuations of all the bidders are common knowledge. Suppose further that the search engine is running a first price auction with the tie-breaking rule that if two bidders make the same bid, then the bidder with the higher valuation is awarded the position with probability one (this is referred to as an efficient tie-breaking rule).

(a) Describe the necessary equilibrium conditions which have to hold for a pure strategy Nash equilibrium.

We first show that the equilibrium, if it exists must have an efficient allocation. Suppose not, i.e. $b_i < b_{i+1}$**, then in equilibrium we would have for some** k, k'**and**

$$\mu_k > \mu_{k'} \text{ and } \quad v_i > v_{i+1}, \tag{1}$$

the following equilibrium utilities:

$$\mu_{k'} (v_i - b_i) \text{ and } \mu_k (v_{i+1} - b_{i+1})$$

and by the equilibrium hypothesis, we also have

$$\mu_{k'} (v_i - b_i) \geq \mu_k (v_i - b_{i+1})$$

and

$$\mu_k (v_{i+1} - b_{i+1}) \geq \mu_{k'} (v_{i+1} - b_i) .$$

But adding the above two inequalities we get

$$(\mu_{k'} - \mu_k) (v_i - v_{i+1}) \geq 0$$

which is a contradiction to the hypothesis in (1). Now given that we must have an efficient allocation it follows that the winner in position i **would never want to pay more than the winner in position** $i+1$ **in the efficient tie-breaking rule, and thus we get**

$$b_i = b_{i+1}. \tag{2}$$

(b) Derive the equilibrium prices that the bidders will pay in a pure strategy Nash equilibrium.

With the insight from (2), we find that

$$b_1 = ... = b_k = b_{k+1}$$

and since the marginal looser would also bid up to this valuation of the click we have

$$b_1 = ... = b_k = b_{k+1} = v_{k+1}$$

(c) Does there always exist a pure strategy Nash equilibrium? Give a precise argument for your assertion.

Now we conclude by observing that since in equilibrium all slots will have to pay the same price, it must be that they are all equally valuable, and hence there is only a pure strategy Nash equilibrium if

$$\mu_1 = \mu_2 = \cdots = \mu_k > 0,$$

and hence almost always there does not exists a pure strategy Nash equilibrium.

2. We considered the notion of reputation in class. Consider the following normal form stage game:

	C	D
C	$2,2$	$0,3$
D	$3,0$	$1,1$

(a) Describe the Nash equilibrium in the static game given by the normal form above.

This is a version of the Prisoner's Dilemma game and it can be solved by dominant strategies and the unique Nash equilibrium of this game, pure or mixed is (D, D).

(b) Suppose now that the game is repeated finitely many times and ends after T periods, with $1 < T < \infty$. The discount factor δ satisfies $0 < \delta < 1$. Does there exist a Nash equilibrium (in every subgame) which leads to play that is different from the prediction of the Nash equilibrium of the static game? Argue carefully.

We would like to support an outcome different from (D, D). But now notice that with a finite horizon, we can solve this game by (backward) induction. The only equilibrium game in the final period T is (D, D). Thus in period $T - 1$, the players can anticipate the play in period T, and as it is independent of the play in the current period, the current period is like the last period, and now by induction it follows that in all periods, the play will have to be (D, D) to form a Nash equilibrium in every subgame.

(c) Suppose now that the stage game is repeated infinitely often, so $T = \infty$ and the discount factor δ satisfies $0 < \delta < 1$. Specify a complete strategy (i.e. for all possible histories of the game) for the repeated game in which a player chooses C in every period as long as all the players cooperated perfectly in the past and punishes failure to play C (by the other player or by herself) with playing D for a some finite number of periods N, and reverts back to play C if the players have followed the punishment phase for N periods.

The following is a complete specification of a repeated game strategy with the desired property:

$$s_i = \begin{cases} C, & \textbf{if} \quad h_0 = \varnothing, \\ C, & \textbf{if} \quad h_t = \{CC, ..., CC\}, \\ C, & \textbf{if} \quad h_t = \left\{ ...., \overbrace{DD, ..., DD,}^{N \textbf{ times}} \right\}, \\ C, & \textbf{if} \quad h_t = \left\{ ...., \overbrace{DD, ..., DD}^{N \textbf{ times}}, CC, ..., CC \right\}, \\ D, & \textbf{if} \quad \textbf{else.} \end{cases}$$

(d) For a given number N (after which players forgive each other), can you compute a discount factor δ such that the strategy described in (c.) actually forms a Nash equilibrium (in every subgame)?

We have to compute the no profitable deviation property along the equilibrium path

$$\frac{2}{1-\delta} \geq 3 + \delta \left(1 + ... + \delta^{N-1}\right) 1 + \delta^{N+1} \frac{2}{1-\delta}$$

and it is immediate that we can cancel and get

$$2 \left(1 + ... + \delta^{N}\right) \geq 3 + \delta \left(1 + ... + \delta^{N-1}\right) 1$$

or

$$\delta \left(1 + ... + \delta^{N-1}\right) 1 \geq 1$$

or

$$\delta \frac{1 - \delta^{N}}{1 - \delta} \geq 1$$

or

$$2\delta - 1 \geq \delta^{N+1}. \tag{3}$$

3. This question concerns the simple model of price formation in an information market given in Chapter 26 of AGT. The description of the model follows. There are n traders, each with a single bit x_i of private information; we use $\mathbf{x}$ to denote the vector $(x_1, x_2, \ldots, x_n)$. We are interested in learning the value of a Boolean function $f : \{0,1\}^n \to \{0,1\}$ of the combined information $\mathbf{x}$. To do this, we set up a market in a security F that will pay \$1 if $f(\mathbf{x})$ is ultimately revealed to be 1 and \$0 otherwise. The form of f (description of the security) is common knowledge among agents. The x_i are referred to as *input bits.* We assume that there is a prior distribution on the input bits and that it, too, is common knowledge among agents.

   The market proceeds in synchronous rounds. In each round, each agent i submits a bid b_i and a quantity q_i. The semantics are that agent i is supplying a quantity q_i of the security and an amount b_i of money to be traded in the market. For simplicity, we assume that there are no restrictions on credit or short sales, and so an agent's trade is not constrained by her possessions. The market clears in each round by settling at a single price that balances the trade in that round: The clearing price is $p = \sum_i b_i / \sum_i q_i$. At the end of the round, agent i holds a quantity q'_i proportional to the money she bid: $q'_i = b_i/p$. In addition, she is left with an amount of money b'_i that reflects her net trade at price p: $b'_i = b_i - p(q'_i - q_i) = pq_i$. Note that agent i's net trade in the security is a purchase if $p < b_i/q_i$ and a sale if $p > b_i/q_i$.

   After each round, the clearing price p is publicly revealed. Agents then revise their beliefs according to any information garnered from the new price. The next round proceeds as the previous. The process continues until an equilibrium is reached. In order to force trade, we assume that $q_i = 1$ for each agent i.

   Chapter 26 of AGT gives a necessary and sufficient condition on f such that, for any prior distribution on $\mathbf{x}$, the equilibrium price of F will reveal the true value of $f(\mathbf{x})$.

   (a) Agents in this model are assumed to be *risk-neutral*, *myopic*, and *truthful*. Give brief definitions of each of these three terms.

   Risk-neutral agents are those who seek to maximize their expected payoff. Myopic agents treat each round

as if it were the final round; they do not reason about how their bids may affect the bids of other agents in future rounds. Truthful agents are those who always bid their current valuations of the security (rather than lying about their current valuations in a strategic effort to influence the equilibrium price); in this model, an agent's current valuation is simply his current estimation of the expected value of the security.

(b) Show how to use this simple model to predict the results of an election in which there are two candidates. That is, give a Boolean function f that captures the semantics of a two-way election and prove that it satisfies the sufficient condition under which the equilibrium price of F will reveal the true value of $f(\mathbf{x})$. For simplicity, assume that the number n of bidders is odd.

Let f be the majority function; that is, $f(x_1, \ldots, x_n)$ is 1 if and only if at least $n/2$ of the input bits x_i are 1. (Because we've made the simplifying assumption that n is odd, that means that $f(x_1, \ldots, x_n)$ is 1 if and only if more of the input bits are 1 than are 0.) This captures the semantics of a two-way election, because the 0 inputs correspond to candidate A and the 1 inputs correspond candidate B; the majority function f will evaluate to 0 on vectors x that correspond to wins by A, and f will evaluate to 1 on vectors x that correspond to wins by B. Theorem 26.12 tells us that the price-formation process described above will converge to the correct answer in this case, because this f is a weighted threshold function: Let $\omega_0 = 0$ and $\omega_i = 2/n$, for $1 \leq i \leq n$; then, $f(x_1, \ldots, x_n) \equiv \omega_0 + \sum_{1 \leq i \leq n} \omega_i x_i \geq 1$ if and only if the majority of the inputs bits are 1 (i.e., if and only if the input vector x corresponds to a win by candidate B). Because of our assumption that n is odd, we could also use $\omega_i = 2/(n+1)$.

(c) For the Boolean function f that you provided in 3b, show how the price-formation process reveals the value of $f(\mathbf{x})$ on the input vector $(0, 1, 0, 0, 1)$ and the uniform common prior. That is, give the bids of each agent in each round and the clearing price in each round, give the equilibrium price, and explain why the process terminates when it does.

For clarity, we will express bids and clearing prices in

fractions of \$1 rather than in cents. In round 1, agents 1, 3, and 4 will each bid $5/16$**, because five out of the 16 equally probable observations by the other four agents contain at least three 1's. On the other hand, each of agents 2 and 5 will bid** $11/16$ **in round 1, because 11 out of the 16 equally probable observations by the other four agents contain at least two 1's. Thus the clearing price in round 1 will be** $(5/16+11/16+5/16+5/16+11/16)/5 = 37/80$**. From this clearing price, all agents will be able to deduce that two input bits are 1 and that three are 0; thus they will all bit \$0 in round 2, and the process will converge with an equilibrium price of \$0. The reason that they can all deduce that exactly two input bits are 1 is that there are exactly six possible clearing prices:** $1/16 = 5/80$ **(corresponding to zero input bits that are 1),** $31/80$ **(corresponding to one 1),** $37/80$ **(corresponding to two 1s),** $43/80$ **(corresponding to three 1s),** $49/80$ **(corresponding to four 1s), and** $11/16 = 55/80$ **(corresponding to five 1s).**