UC Berkeley
Haas School of Business
Economic Analysis for Business Decisions
(EWMBA 201A)

Game Theory III
Applications (part 2)

Lectures 8-9
Sep. 19, 2009

Outline

[1] Midterm

[2] Strictly competitive games and maximinzation

[3] Oligopolistic competition – the main ideas

[4] Auctions

[5] Bargaining / negotiations

[6] Observational learning – open discussion (if time permits)

Strictly competitive games

In strictly competitive games, the players' interests are diametrically opposed, like in Bertrand's duopoly game.

More precisely, a strategic two-player game is strictly competitive if for any two outcomes a and b we have

$$a \succsim_1 b \text{ if and only if } b \succsim_2 a.$$

A strictly competitive game can be represented as a zero-sum game

	L	R
T	$A, -A$	$B, -B$
B	$C, -C$	$D, -D$

This class of games is important for a number of reasons:

– A simple decision making procedure leads each player to choose a Nash equilibrium action.

– Unfortunately, there are innumerable social and economic situations which are strictly competitive.

– In the game of business, a successful strategy is <u>avoiding</u> the zero-sum games by <u>reshaping</u> the game.

$\Longrightarrow$ See, Brandenburger and Nalebuff, and Hermalin (Chapter 6).

Maxminimization

A maxminimizing strategy is a (mixed) strategy that maximizes the player's minimal payoff.

> A strategy that maximizes the player's expected payoff under the (very pessimistic) assumption that whatever she does the other player will act in a way that minimizes her expected payoff.

A pair strategies in a strictly competitive game is a Nash equilibrium if and only if each player's strategy is a maxminimizer (or a minimaximzer).

An example

	L	R
T	$2, -2$	$-1, 1$
B	$-1, 1$	$1, -1$

The maxminimizing strategy of player 1 is $(2/5, 3/5)$, which yields her a payoff of $1/5$.

Some history: the theory was developed by von Neumann in the late 1920s but the idea appeared two centuries earlier (Montmort, 1713-4).

Review of oligopolistic competition

- In an oligopolistic market, a small number of firms compete for the business of consumers.

- Barriers to entry restrict the number of firms, possibly allowing firms to earn substantial profits.

- In models of oligopoly, decisions involve strategic considerations – each firm must best response to the actions of its rivals.

Cournot's model

– Firms produce a homogeneous good, and all firms decide simultaneously how much to produce.

– Each firm chooses a profit-maximizing output, treating the outputs of its rivals as fixed.

– In the Nash equilibrium, each firm's profit is higher (resp. lower) than it is under perfect competition (resp. collusion).

Stackelberg's model

– Firms move sequentially so we use backward induction to find a subgame perfect equilibrium.

– The firm that sets its output first has a strategic advantage and earns a higher profit in equilibrium.

– Which model, Cournot or Stackelberg, is more appropriate depends on the industry (operating advantage / leadership position).

Bertrand's model

– Cournot

  Each firm chooses an output, and the price is determined by the market demand in relation to the total output produced.

– Bertrand

  Each firm chooses a price, and produces enough output to meet the demand it faces, given the prices chosen by *all* the firms.

Suppose again that there are two firms (a "duopoly") and that the cost for firm $i = 1, 2$ for producing q_i units of the good is given by cq_i (equal constant "unit cost").

Assume that the demand function is

$$D(p) = A - p$$

for $A \geq p$ and zero otherwise, and that $A > c$ (the demand function in PR 12.3 is different).

Because the cost of producing each until is the same, equal to c, firm i makes the profit of $p_i - c$ on every unit it sells.

Thus firm i's profit is

$$\pi_i = \begin{cases} (p_i - c)(A - p_i) & \text{if} \quad p_i < p_j \\ \frac{1}{2}(p_i - c)(A - p_i) & \text{if} \quad p_i = p_j \\ 0 & \text{if} \quad p_i > p_j \end{cases}$$

where j is the other firm.

The unique subgame perfect equilibrium is $(p_1, p_2) = (c, c)$:

- If one firm charges the price c, then the other firm can do no better than charge the price c.

- If $p_1 > c$ and $p_2 > c$, then each firm i can increase its profit by lowering its price p_i slightly below p_j.

$\Longrightarrow$ Cournot: the market price decreases toward c (zero profits) as the number of firms increases.

$\Longrightarrow$ Bertrand: the price is c even if there are only two firms (and the price remains c when the number of firm increases).

Auctions (PR 13.8)

From Babylonia to eBay, auctioning has a very long history.

- Babylon:
  - women at marriageable age.
- Athens, Rome, and medieval Europe:
  - rights to collect taxes,
  - dispose of confiscated property,
  - lease of land and mines,

  and more...

The word "auction" comes from the Latin *augere*, meaning "to increase."

The earliest use of the English word "auction" given by the *Oxford English Dictionary* dates from 1595 and concerns an auction "when will be sold Slaves, household goods, etc."

In this era, the auctioneer lit a short candle and bids were valid only if made before the flame went out – Samuel Pepys (1633-1703) –

- Auctions, broadly defined, are used to allocate significant economics resources.

    Examples: works of art, government bonds, offshore tracts for oil exploration, radio spectrum, and more.

- Auctions take many forms. A game-theoretic framework enables to understand the consequences of various auction designs.

- Game theory can suggest the design likely to be most effective, and the one likely to raise the most revenues.

Types of auctions

Sequential / simultaneous

Bids may be called out sequentially or may be submitted simultaneously in sealed envelopes:

– English (or oral) – the seller actively solicits progressively higher bids and the item is soled to the highest bidder.

– Dutch – the seller begins by offering units at a "high" price and reduces it until all units are soled.

– Sealed-bid – all bids are made simultaneously, and the item is sold to the highest bidder.

First-price / second-price

The price paid may be the highest bid or some other price:

– First-price – the bidder who submits the highest bid wins and pay a price equal to her bid.

– Second-prices – the bidder who submits the highest bid wins and pay a price equal to the second highest bid.

Variants: all-pay (lobbying), discriminatory, uniform, Vickrey (William Vickrey, Nobel Laureate 1996), and more.

Private-value / common-value

Bidders can be certain or uncertain about each other's valuation:

– In private-value auctions, valuations differ among bidders, and each bidder is certain of her own valuation and can be certain or uncertain of every other bidder's valuation.

– In common-value auctions, all bidders have the same valuation, but bidders do not know this value precisely and their estimates of it vary.

First-price auction (with perfect information)

To define the game precisely, denote by v_i the value that bidder i attaches to the object. If she obtains the object at price p then her payoff is $v_i - p$.

Assume that bidders' valuations are all different and all positive. Number the bidders 1 through n in such a way that

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

Each bidder i submits a (sealed) bid b_i. If bidder i obtains the object, she receives a payoff $v_i - b_i$. Otherwise, her payoff is zero.

Tie-breaking – if two or more bidders are in a tie for the highest bid, the winner is the bidder with the highest valuation.

In summary, a first-price sealed-bid auction with perfect information is the following strategic game:

- Players: the n bidders.

- Actions: the set of possible bids b_i of each player i (nonnegative numbers).

- Payoffs: the preferences of player i are given by

$$u_i = \begin{cases} v_i - \bar{b} & \text{if} \quad b_i = \bar{b} \text{ and } v_i > v_j \text{ if } b_j = \bar{b} \\ 0 & \text{if} \quad b_i < \bar{b} \end{cases}$$

  where $\bar{b}$ is the highest bid.

The set of Nash equilibria is the set of profiles $(b_1, ..., b_n)$ of bids with the following properties:

$$\begin{array}{ll} [1] & v_2 \leq b_1 \leq v_1 \\ [2] & b_j \leq b_1 \ \underline{\text{for all}} \ j \neq 1 \\ [3] & b_j \leq b_1 \ \underline{\text{for some}} \ j \neq 1 \end{array}$$

It is easy to verify that all these profiles are Nash equilibria. It is harder to show that there are no other equilibria. We can easily argue, however, that there is no equilibrium in which player 1 does not obtain the object.

$\Longrightarrow$ The first-price sealed-bid auction is socially efficient, but does not necessarily raise the most revenues.

Second-price auction (with perfect information)

A second-price sealed-bid auction with perfect information is the following strategic game:

- Players: the n bidders.
- Actions: the set of possible bids b_i of each player i (nonnegative numbers).
- Payoffs: the preferences of player i are given by

$$u_i = \begin{cases} v_i - \overline{b} & \text{if} \quad b_i > \overline{b} \text{ or } b_i = \overline{b} \text{ and } v_i > v_j \text{ if } b_j = \overline{b} \\ 0 & \text{if} \quad b_i < \overline{b} \end{cases}$$

  where $\overline{b}$ is the highest bid submitted by a player other than i.

First note that for any player i the bid $b_i = v_i$ is a (weakly) dominant action (a "truthful" bid), in contrast to the first-price auction.

The second-price auction has many equilibria, but the equilibrium $b_i = v_i$ for all i is distinguished by the fact that every player's action dominates all other actions.

Another equilibrium in which player $j \neq 1$ obtains the good is that in which

$$
\begin{array}{ll}
[1] & b_1 < v_j \text{ and } b_j > v_1 \\
[2] & b_i = 0 \ \underline{\text{for all}} \ i \neq \{1, j\}
\end{array}
$$

Common-value auctions and the winner's curse

Suppose we all participate in a sealed-bid auction for a jar of coins. Once you have estimated the amount of money in the jar, what are your bidding strategies in first- and second-price auctions?

The winning bidder is likely to be the bidder with the largest positive error (the largest overestimate).

In this case, the winner has fallen prey to the so-called the winner's curse. Auctions where the winner's curse is significant are oil fields, spectrum auctions, pay per click, and more.

Spectrum auctions

- Spectrum licenses in the US were originally allocated on the basis of hearings by the Federal Communications Commission (FCC).

- This process was inefficient and time-consuming. A large backlog prompt a switch to lotteries.

- A winner of a license to run cellular telephones in Cape Cod sold the license shortly after the lottery for $41.5 million.

- The first four spectrum auctions generated revenues of over $18 billion, and the winners were able to obtain the set of licenses they wanted.

Budget and exposure is spectrum auctions (Milgrom's case)

- Bidding teams often face budget constraints and yet have considerable freedom in deciding which licenses to buy within their budgets.

- A bidder's exposure is the sum of all of its standing bids, whether provisionally winning or not.

- It is the largest amount that a bidder might have to pay if all of its bids were to become winning.

- If a bidder faces a binding budget constraint and has broad interests, then as prices increase from round to round, its total exposure will eventually level at an amount approximating its budget.

- If all bidders were to fall in this category, then the total exposure of all bidders in the auction would rise to the level of the aggregate bidder budgets and level off, forecasting the final auction prices!

- A winning play – saving nearly $1.2 billion on spectrum license purchases compared to the prices paid by other large bidders...

Total exposure and revenues I

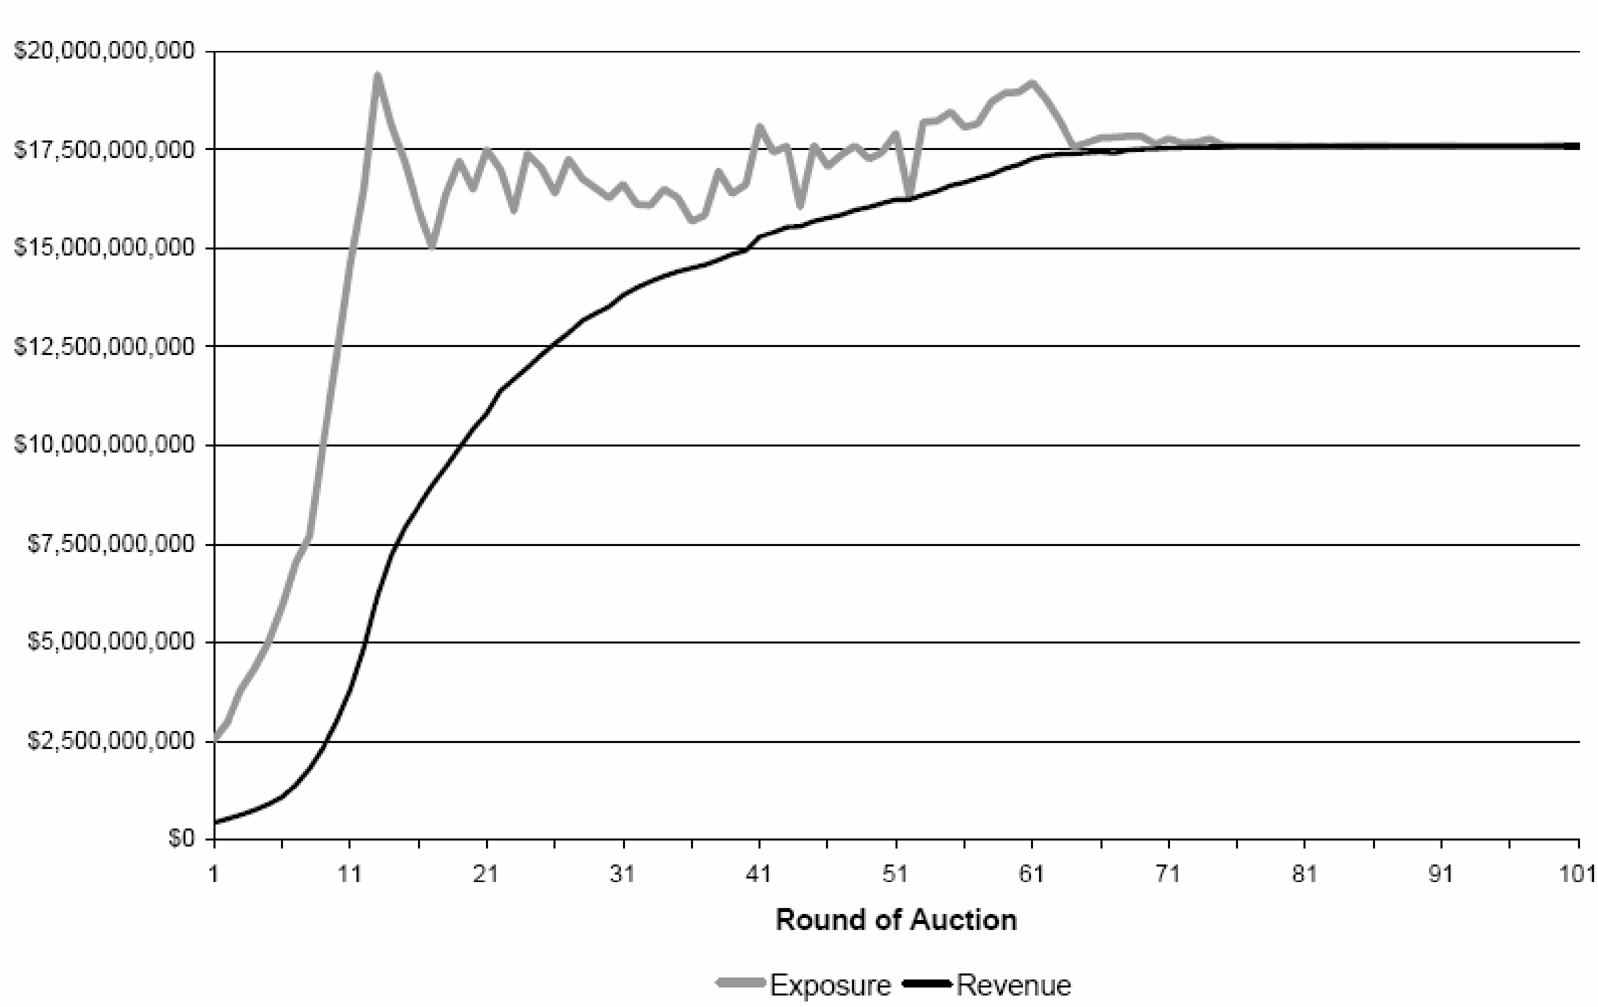

Total exposure and revenues II

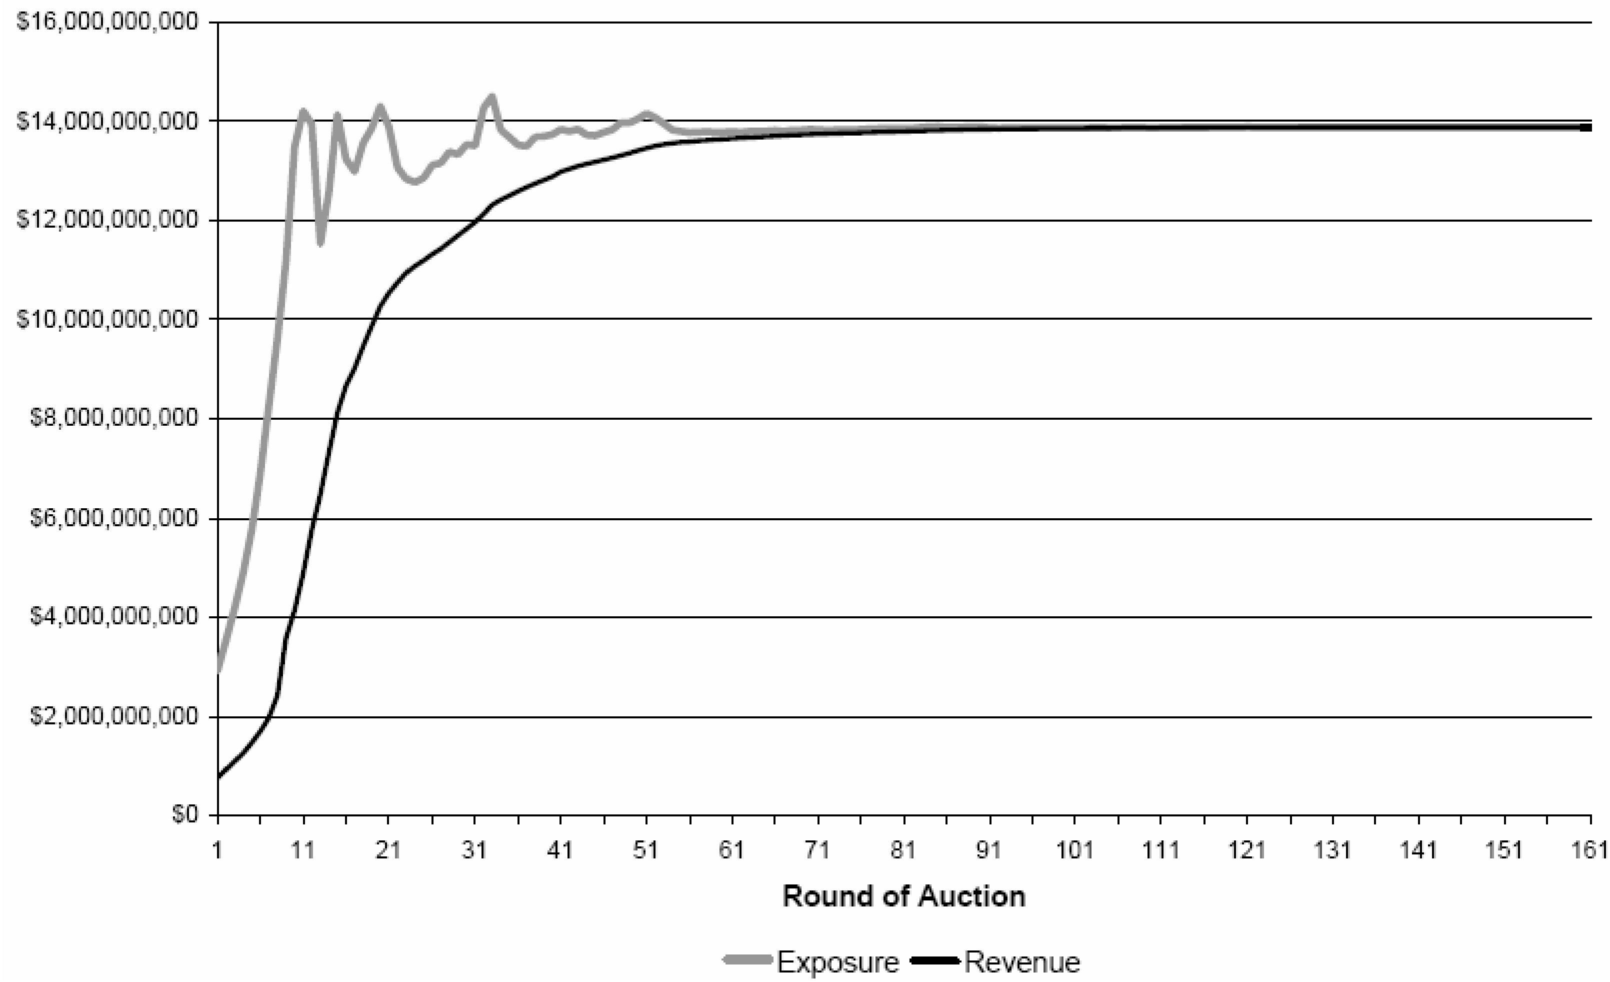

Bidder exposure I

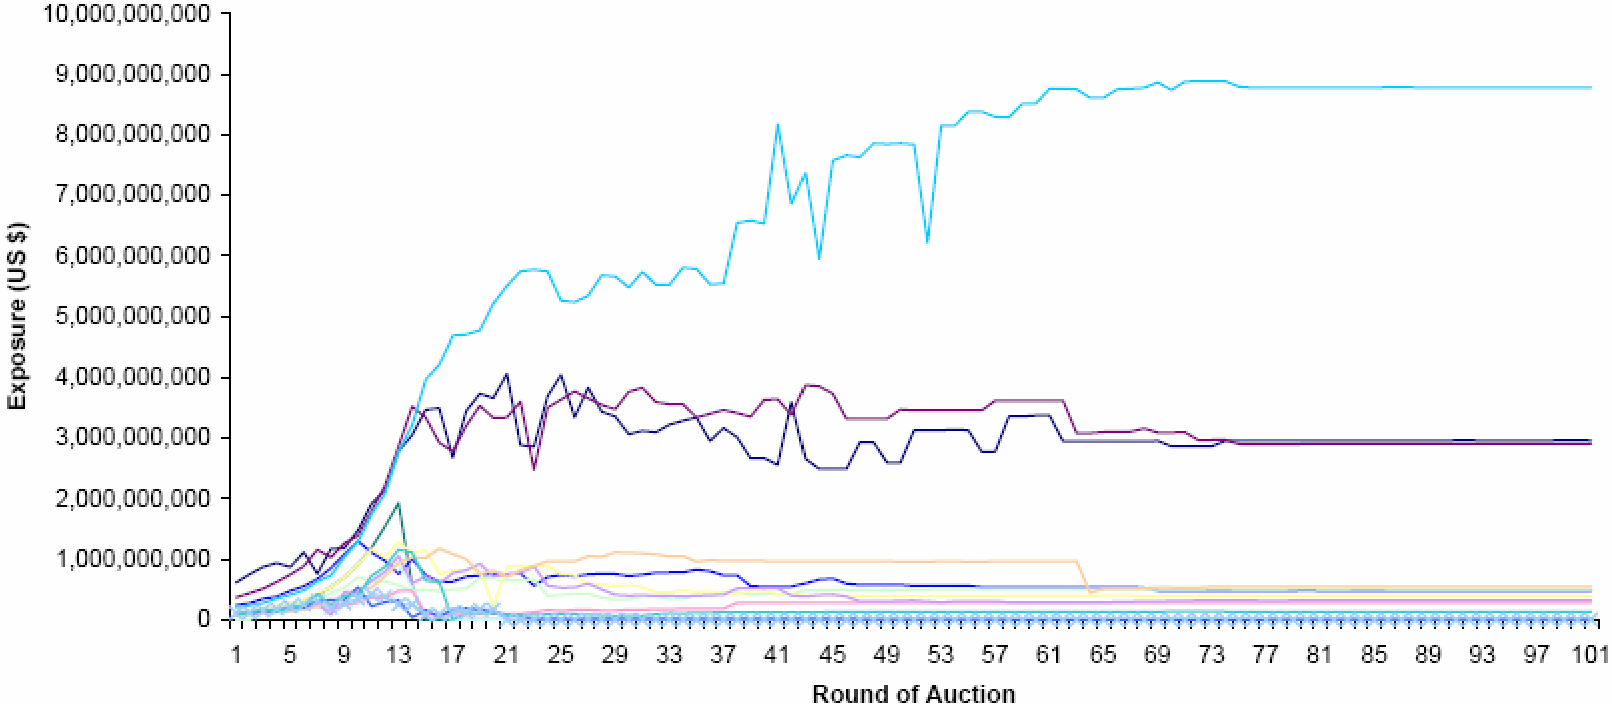

Bidder exposure II

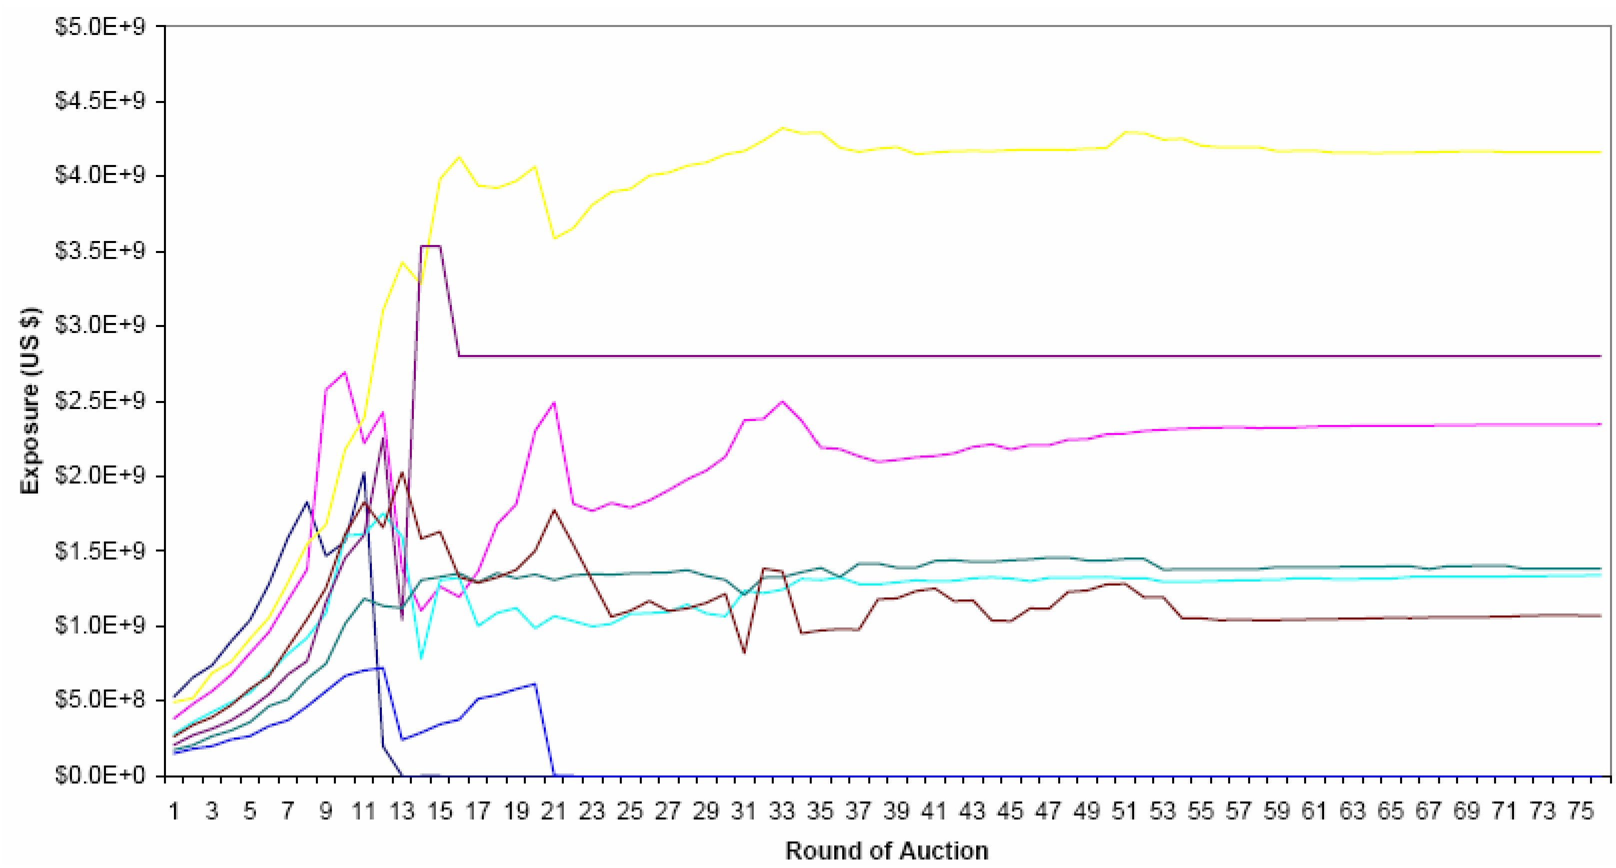

Purchases in the auction
FCC Auction 66, Aug-Sep 2006

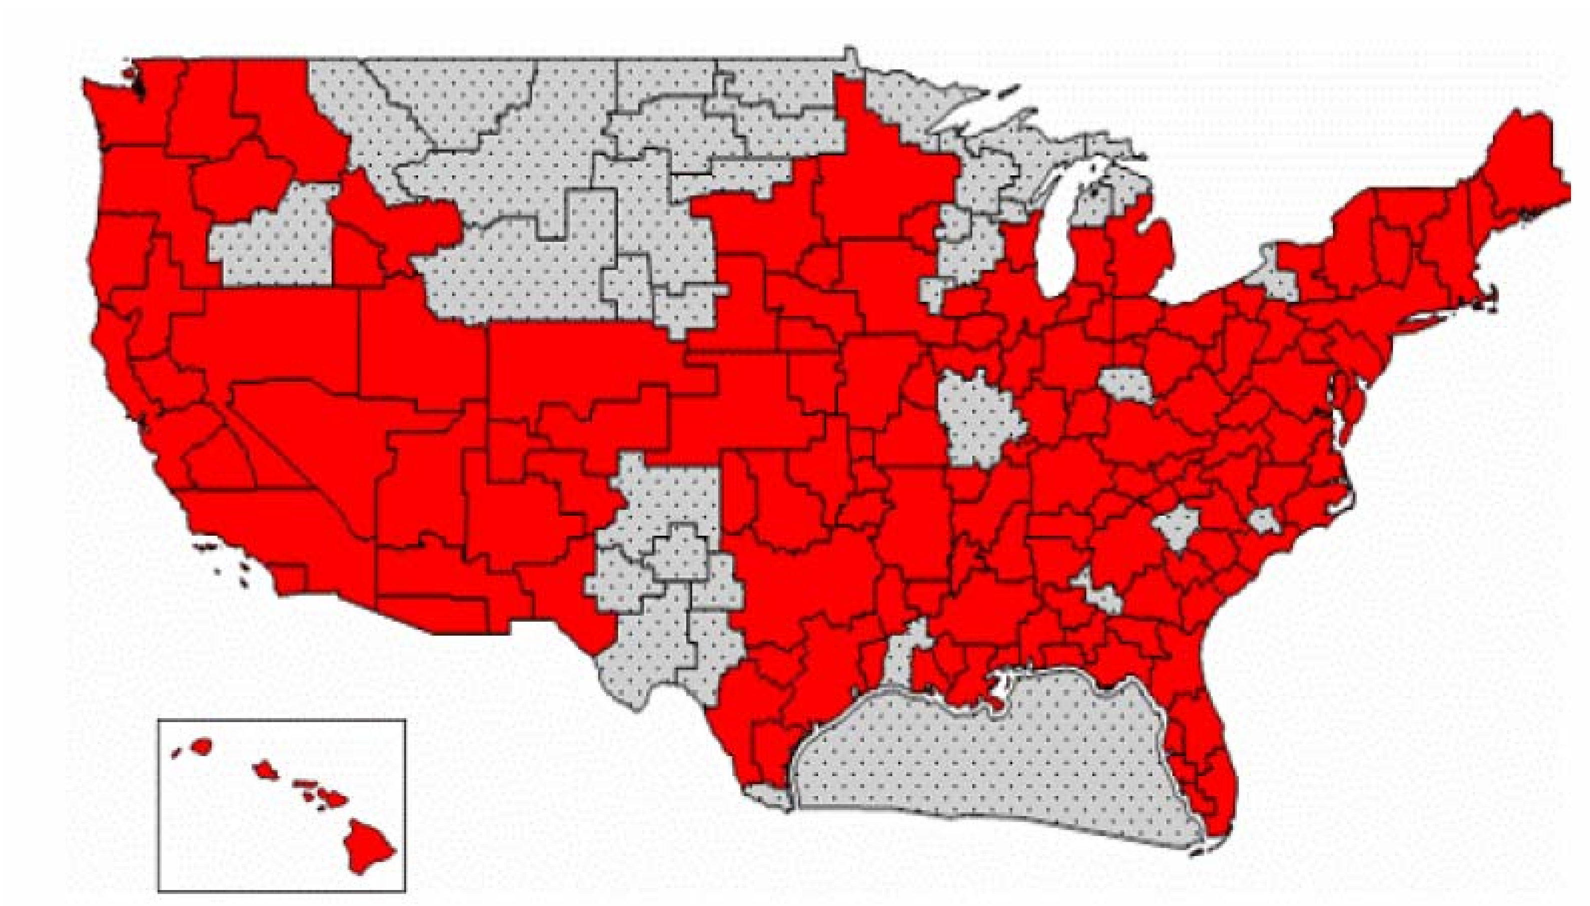

Performance of the five bidders spending over $1 billion

Bidder	Total winning bids (billions)	Price per MHz-Pop
SpectrumCo	$2.4	$0.451
Cingular	$1.3	$0.548
T-Mobile	$4.2	$0.630
Verizon wireless	$2.8	$0.731
MetroPCS	$1.4	$0.963

Bargaining

The strategic approach (Rubinstein, 1982)

- Two players $i = 1, 2$ bargain over a "pie" of size 1.

- An agreement is a pair (x_1, x_2) where x_i is player i's share of the pie.

- The set of possible agreements is $x_1 + x_2 = 1$ where for any two possible agreements x and y

$$x \succsim_i y \text{ if and only if } x_i \geq y_i$$

The bargaining procedure

- The players can take actions only at times in an (infinite) set of dates.

- In each period t player i, proposes an agreement (x_1, x_2) and player $j \neq i$ either accepts (Y) or rejects (N).

- If (x_1, x_2) is accepted (Y) then the bargaining ends and (x_1, x_2) is implemented. If it is rejected (N) then the play passes to period $t + 1$ in which j proposes an agreement (alternating offers).

Preferences

The preferences over outcomes alone may not be sufficient to determine a solution. Time preferences (toward agreements at different points in time) are the driving force of the model:

– - Disagreement is the worst outcome.

  - The pie is desirable and time is valuable.

  - Increasing loss to delay.

Under this assumptions, the preferences of player i are represented by

$$\delta_i^t u_i(x_i)$$

for any $0 < \delta_i < 1$ and u_i is increasing and concave function

Any two-player bargaining game of alternating offers in which players' preferences satisfy the assumptions above has a unique (!) subgame perfect equilibrium.

$\Longrightarrow$ Player 1 (moves first) always proposes

$$(x_1^*, x_2^*) = (\frac{1-\delta_2}{1-\delta_1\delta_2}, \frac{\delta_2(1-\delta_1)}{1-\delta_1\delta_2}),$$

and accepts an offer (y_1, y_2) of player 2 if and only if $y_1 \geq y_1^*$.

$\Longrightarrow$ Player 2 always proposes

$$(y_1^*, y_2^*) = (\frac{\delta_1(1-\delta_2)}{1-\delta_1\delta_2}, \frac{1-\delta_1}{1-\delta_1\delta_2}).$$

and accepts an offer $(x_1; x_2)$ of player 1 if and only if $x_2 \geq x_2^*$.

The unique outcome is that player 1 proposes (x_1^*, x_2^*) at the first period and player 2 accepts (no delay!).

When players have the same discount rate $\delta_1 = \delta_2 = \delta$ then

$$(x_1^*, x_2^*) = (\frac{1}{1+\delta}, \frac{\delta}{1+\delta}),$$

and

$$(y_1^*, y_2^*) = (\frac{\delta}{1+\delta}, \frac{1}{1+\delta}).$$

$\Longrightarrow$ Properties of subgame equilibrium: efficiency (no delay), first-mover advantage (perfect information), effects of changes in patience.

The axiomatic approach (Nash, 1950)

Nash's (1950) work is the starting point for formal bargaining theory.

– Bargaining problem: a set of utility pairs $(s_1; s_2)$ that can be derived from possible agreements, and a pair of utilities (d_1, d_2) which is designated to be a disagreement point.

– Bargaining solution: a function that assigns a *unique* outcome to *every* bargaining problem.

Let S be the set of all utility pairs $(s_1; s_2)$. $\langle S, d \rangle$ is the primitive of Nash's bargaining problem.

Nash's axioms

One states as axioms several properties that it would seem natural for the solution to have and then one discovers that the axioms actually determine the solution uniquely – Nash 1953 –

[1] Invariance to equivalent utility representations (INV)

If $\langle S', d' \rangle$ is obtained from $\langle S, d \rangle$ by "monotonic" transformations then $\langle S', d' \rangle$ and $\langle S, d \rangle$ represent the same situation.

INV requires that the utility outcome of the bargaining problem co-vary with representation of preferences. The physical outcome predicted by the bargaining solution is the same for $\langle S', d' \rangle$ and $\langle S, d \rangle$.

[2] Symmetry (SYM)

A bargaining problem $\langle S, d \rangle$ is *symmetric* if

$$d_1 = d_2$$

and

$$(s_1, s_2) \text{ is in } S \text{ if and only if } (s_2, s_1) \text{ is in } S.$$

If the bargaining problem $\langle S, d \rangle$ is symmetric then the bargaining solution must assign the same utility.

Nash does not describe differences between the players. All asymmetries (in the bargaining abilities) must be captured by $\langle S, d \rangle$.

[3] Independence of irrelevant alternatives (IIA)

If $\langle S, d \rangle$ and $\langle T, d \rangle$ are bargaining problems, S is a strict subset of T, and the solution to $\langle T, d \rangle$ is in $\langle S, d \rangle$ then it is also the solution to $\langle S, d \rangle$.

Put diffidently, if T is available and players agree on (s_1, s_2) in S then they also agree on the same (s_1, s_2) if only S is available.

IIA excludes situations in which the fact that a certain agreement is available influences the outcome.

Pareto efficiency (PAR)

If $\langle S, d \rangle$ is a bargaining problem where $(s_1; s_2)$ and (t_1, t_2) are in S and $t_i > s_i$ for $i = 1, 2$ then the solution is *not* (s_1, s_2).

Players never agree on an outcome (s_1, s_2) when there is an outcome (t_1, t_2) in which both are better off.

After agreeing on the outcome (s_1, s_2), players can always "renegotiate" and agree on (t_1, t_2).

Nash's solution

There is precisely *one* bargaining solution, satisfying SYM, PAR, INV and IIA.

The unique bargaining solution is the utility pair that maximizes the product of the players' utilities

$$\underset{(s_1,s_2)}{\arg\max}\, s_1 s_2$$

$\Longrightarrow$ Application: wage bargaining.