

Game Theory and Competition

R. Preston McAfee and John McMillan

March 21, 1998

1. Introduction

"Game theory is hot," the *Wall Street Journal* has proclaimed.¹ As game theory has swept through economics and related disciplines, among its many payoffs have been new insights into that most fundamental of business processes, competition.

The classical theory of perfect competition takes a broad-brush approach: it depicts the outcome of competition, but not the activity of competing. Much of what is interesting and important about competition is hidden in the background. The *Oxford English Dictionary* defines competition as "Rivalry in the market, striving for custom between those who have the same commodities to dispose of." A perfect competitor, as depicted in the textbooks, does not actually do any competing in this sense. "Striving for custom" implies a dynamic process, the action of competing. A perfectly competitive firm does not pay attention to what any of the other firms in the industry are doing. Instead, it passively accepts the going market price. Any "rivalry in the market" is assumed away. The new game-theoretic models, by contrast, view competition as a process of strategic decision-making under uncertainty; they depict people and

¹ *Wall Street Journal*, February 13 1995, p.A19.

firms engaged in competition. Revealing information is seen to be one of competition's effects.

Competition, as we will discuss, helps sellers to price items when buyers' willingness to pay is unknown. Being faced with competition on the other side of the market is a source of bargaining power. Competition can be used to generate incentives for productive effort. Competition in product markets disciplines managers and in financial markets promotes efficient commitments.² The game-theory models of competition do not supplant the textbook wisdom; they supplement it. They do not undermine the crucial finding that, under certain precisely specified conditions (Debreu, 1959), competition sets prices right, thereby generating an efficient allocation of resources. What the new models do is add micro-detail to the classical theory, modeling, among other things, how prices get to be set through competition.

2. Competition Reveals Hidden Information

The simplest setting in which to study competition is in the sale of a single item by open outcry auction. What is learned from this can be generalized to richer and more complex competitive mechanisms.

Suppose that each of the competing bidders attaches a different, subjective value to the item for sale. They all know exactly how much it would be worth to themselves to own it, but not how highly any of the others values it. (This is called the independent-private-values case.) The seller is ignorant of the bidders'

² For a more on the new concepts of competition, see Vickers (1995). On theories of competitive bidding, see McAfee and McMillan (1987), McMillan (1992, Ch.11), Milgrom (1989), and Wilson (1992).

valuations, so would find it difficult to set a price for it. Each bidder's rational strategy in the open auction is to be willing to bid until the price passes his or her valuation. Rather than let the item sell to someone else for a price below one's valuation, it is better to bid and possibly obtain a profit. Conversely, should the price go above one's valuation, winning would mean paying more than value; bidding above one's valuation is unprofitable. The selling price therefore equals (approximately) the second-highest valuation. The bidding process reveals information. The seller may have initially known nothing about the value of the item. After the auction, the price gives the seller an estimate of value. It is an underestimate, since the price is the second-highest valuation. But if the number of bidders is reasonably large, the price reached in the auction comes close to the winning bidder's valuation.

Bidding competition can also serve to reveal another kind of information. Let us change the nature of the uncertainty about valuations, so as to make valuations objective rather than subjective. The item, in this model (the common-value case), is worth the same no matter who wins it. At the time of bidding, however, no one knows what this value is. Each bidder has an estimate of value, which is subject to error. (This model applies to the bidding for an oil well, for example.) Bidders are now faced with the risk of the winner's curse. The winner is likely to be the bidder who has most overestimated the item's value. There is a tendency, therefore, for the winning bidder to overbid. Bidders can avoid the winner's curse by bidding cautiously: recognizing they will win only if they have relatively high estimates, and discounting their estimates. The winning bid then is on average below the item's true but unknown value; but with more and more bidders the price approaches the true value, and with a very large number of bidders, comes very close to it (Wilson, 1977; Milgrom, 1979). This is a

remarkable result. No one person knows the true value; each individual's estimate may be highly imperfect. But the price is an accurate value estimate. The competitive process serves to aggregate the scattered information.

Auction theory generates practical recommendations for an organizer of a bidding competition. Open bidding rather than sealed bidding should be used, for this encourages higher bids. A minimum price should be set, high enough to force a high bid but not so high that no one is likely to bid. Royalties or other forms of revenue sharing might be used to make the bidding more competitive. The seller should release any information she has that is relevant to the bidders' assessment of the item's value, for doing so will increase the competitiveness of the bidding.

3. Competition Works Better than Bargaining

Competition is a good substitute for bargaining skill.³ Consider a seller of an indivisible item who has the choice of selling by means of either bargaining with a single buyer, or offering the item for competitive bidding between two potential buyers. In the bargaining case, we will allow the seller to be a determined, ruthless negotiator, who is well-informed in that, although she does not know the buyer's valuation exactly, she does know the range of possible valuations. In the bidding case, the seller passively accepts bids without setting a minimum price, knowing nothing about even the range of possible bidder valuations.

³ For an introduction to game-theoretic models of bargaining, see McMillan (1992, Chs. 5, 6).

To be specific, let us suppose that a buyer's valuation could be either \$40 or \$100, each with equal probability; only the buyer himself knows which it actually is. The seller attaches no value to the item if it is unsold. Under bargaining, the best the seller can do is set a take-it-or-leave-it price of \$100, selling the item only if the buyer happens to have the higher valuation, for an expected profit of \$50. (Given that the seller is not very risk averse, she will prefer this to the alternative of pricing the item at \$40 and selling it for sure, for a profit of \$40.) Under bidding, on the other hand, there are four equally likely possibilities for the two bidders' valuations: (\$40, \$40), (\$40, \$100), (\$100, \$40), and (\$100, \$100). Bidding competition drives the price up to the second-highest valuation, \$40 in the first three cases and \$100 in the fourth, so the seller's expected profit is $\$40 \times 0.75 + \100×0.25 , which equals \$55. The price reached by competitive bidding, on average, is more than the negotiated price.

The intuition behind this result (which is due to Bulow and Klemperer (1996)) is that a good bargainer functions like an artificial competitor. The bargainer's main source of bargaining power, the threat to refuse to sell the item if the price is not high enough, is analogous to going to another bidder in the competitive case. A real competitor is more effective than a fake one. The result generalizes to describe varying degrees of competition. An extra bidder helps the seller, in that she is better off with $n+1$ bidders than with n bidders and actively stimulating the bidding competition by setting a minimum price.⁴

⁴ Bulow and Klemperer (1996) show this result to hold in far more general settings than here, though a caveat is needed. The conclusion that the expected price is higher with two bidders and a passive seller than with a single buyer and a committed seller assumes the probability distribution function of valuations satisfies what is called the monotone-hazard-rate condition; this will hold in most but not all cases.

The seller is better off, therefore, being a passive acceptor of bids from competing potential buyers than a tough, clever bargainer without buyer competition. Moreover, being a bid-taker puts much less of a burden on the seller's knowledge and abilities than being a negotiator. The bargaining game requires the seller to know the range of possible buyer valuations and to compute the best price to charge, while in the bidding game the seller needs no information about valuations and needs put no thought into her strategy. To implement the optimal bargaining strategy the seller has to be able credibly to commit herself to her \$100 price demand, even in the case that the seller finds out the buyer's valuation must be \$40 through the buyer's rejecting the offer. The seller, upon failing to negotiate a price of \$100, will be tempted to cut the price to \$40. But any such willingness to undercut would reduce the seller's initial bargaining power, for the buyer, in the event that he values the item at \$100, would now be tempted to wait for a price reduction (McAfee and Vincent, 1995). In the bidding game, by contrast, the seller's commitment is merely to sell to the highest bidder, a much easier commitment to stick to. Competition, therefore, economizes on knowledge, computation, and commitment abilities, and yields higher prices. Competition is more effective than bargaining.

4. Competition Creates Effort Incentives

Competition can be used to create incentives for productive effort. Salespeople compete to win a prize such as a trip to Hawaii. Assistant professors compete for limited tenure slots. Corporate vice-presidents compete to be promoted to chief executive officer. These contests or tournaments, which make remuneration depend not only on people's own performance but also that of their peers, can be structured as incentive devices (Lazear and Rosen, 1982).

Suppose two agents, A and B, each produce an output Q according to $Q = E + u$, where E is the effort the agent exerts and u is a random variable representing the agent's luck. The personal cost to the agent of exerting the effort is $C(E)$, and there are diminishing returns to effort. A prize of W is offered to the agent who produces the most; the other agent receives nothing. Agent A wins the prize if she produces more than B, i.e., $Q_a > Q_b$; this occurs if $E_a - E_b > u_b - u_a$. Represent the probability distribution of the random terms by $G(u_b - u_a)$: this reflects the degree of correlation of the two agents' luck. The probability of A's winning is then $G(E_a - E_b)$, and agent A's expected net return from exerting effort is therefore $W \cdot G(E_a - E_b) - C(E)$. Differentiation of this with respect to E_a yields A's optimal effort level E_a^* (for any given effort by B, E_b), satisfying the condition $W \cdot g(E_a^* - E_b) - C'(E_a^*) = 0$ (where the prime denotes a derivative and g denotes the probability density function corresponding to G). This shows agent A's best response E_a^* to agent B's effort E_b . Agent B has a similar best-response function. The symmetric Nash equilibrium has each winning with 50% probability, and each exerting the effort level E^* that satisfies $C'(E^*) = W \cdot g(0)$.

The larger the prize, therefore, the more effort it engenders. The designer of the incentive scheme can elicit from the two agents whatever effort she wishes by setting the prize W appropriately. The size of the prize that is needed reflects the extent to which luck affects output. If the randomness in measured performance is large (i.e., $g(0)$ is small), then the prize must be large if it is to generate much effort from the agents. If luck has a big effect on output, the agents will tend to sit back and rely on luck to get them the prize, so a large prize is needed to counter this incentive. If, on the other hand, an agent's output is tightly related to effort, with little variability uncontrolled by the agent (i.e., $g(0)$ is large), a

smaller prize will suffice. Considerable disparities across ranks in employees' pay might be explained by the effort incentives they induce at the lower ranks.

5. Competition Provides Market Discipline

"The best of all monopoly profits is a quiet life," as John Hicks said. Do monopolists produce above minimum cost, causing a welfare loss? Conversely, does competition force lean production, mitigating the weak incentives that come from the separation of ownership and control? Generations of economists have believed that product-market competition disciplines managers to make relatively efficient production decisions. Despite the familiarity of the idea that product-market competition promotes efficiency and the fact that it has empirical support, however, it still lacks a rigorous theoretical basis. The proposition is so fundamental to economics that at some level it just has to be true. But it is surprisingly difficult to obtain it as an unambiguous prediction from a principal-agent model. (For some recent such modeling, see for example Aghion et al. (1996), McAfee and McMillan (1995), Stole and Zwiebel (1996).)

Competition in financial markets also can have efficiency-enhancing effects. In a monopolized financial system, borrowers' budget constraints may be soft: they can go on getting further loans even after it becomes clear to the bank that the project should be jettisoned. Competition among banks (according to Dewatripont and Maskin (1995)) creates the ability to make commitments: a decentralized banking system can commit to abandon any project if it is revealed to be unprofitable, while a monopoly bank cannot. A monopoly bank might have an incentive, once started, to invest resources in propping up an unprofitable project. If banks are small enough that it takes more than one to

fund a project, on the other hand, then the bank that makes the initial loan has no incentive to persist with the project after it sees it is a money-loser, as the benefits from its efforts will accrue to the other lending banks. A competitive financial system will therefore have a more efficient loan portfolio than a monopolized one.

6. Competitive Mechanisms Are Robust

Competitive mechanisms tend to be robust: they work well even in the face of mistakes or irrational behavior by the market participants. This theme runs through the new models of competition, though it has yet to be fully explored. A well designed (or appropriately evolved) competitive mechanism works efficiently even when the the people whose behavior it determines lack the knowledge and computational abilities that are usually assumed in economic modeling.

The outcome of market competition is more likely to be in conformance with game-theoretic rationality than the outcome of a bilateral negotiation, according to the experiments of Roth et al. (1991). The competitiveness of the market process constrains participants in their choice of strategies, so the outcome is usually the game-theoretic equilibrium, whereas the indeterminacy of bargaining⁵ leaves scope for bargainers to engage in complicated game-playing, with the result that outcomes are more idiosyncratic.

⁵ That is, the range of (Pareto efficient) outcomes that are Nash equilibria but not perfect equilibria.

Different competitive mechanisms differ in their robustness. An open auction, in which bidders call their bids and go on raising them until only one remains, is more robust than a sealed-bid auction, in which bidders make a single bid in ignorance of their rivals' bids. This is because, in the open auction, bidders can learn from the others' bids, while in the sealed-bid auction they cannot. To make a good bid decision in the sealed-bid auction, a bidder must conjecture the others' bid decisions. Bad outcomes--either losing by bidding too low, or overpaying through bidding higher than turns out to be needed to win--can result from incorrect conjectures. In the open auction the bidder's decisions are easier and the outcome is more controllable. For these reasons bidders typically say they prefer open auctions to sealed-bid auctions, even though, according to auction theory, open auctions produce higher prices.

The rules for competition, if well designed, can ensure that a market produces an allocation that is close to efficient even with traders who are incapable of calculating what is in their interest, according to experiments by Gode and Sunder (1993). The wisdom of the market compensates for the lack of rationality of the market participants.

Competitive tournaments, as noted earlier, provide effort incentives. They are not the only way of generating incentives; instead, people could simply be paid according to their own performance. But tournaments are more robust than direct pay-for-performance, in that they economize on the principal's information gathering. Observing whether one person has done a better job than another may be easier than measuring each person's output precisely. Also, if the randomness that affects output is correlated across the agents, the

tournament filters out the common randomness, permitting more accurate inferences about each agents' effort.

7. Designing New Competitive Mechanisms

The ultimate test of a theory of markets, more stringent than confronting the theory with data from existing markets, is to use the theory in designing innovative market mechanisms. The new theoretical understanding of how competitive markets function, together with the use of high-speed computers, has made possible the invention of novel markets (McAfee and McMillan, 1996b). Electronic markets can allocate goods efficiently in circumstances where simple supply and demand works poorly: because goods are idiosyncratic and differentiated, or there are multiple goods with synergies among them, or buyers' preferences are ill-behaved by the criteria of standard theory, or there is a need to match particular buyers and sellers. Economic theory has been used in the design of competitive bidding mechanisms to sell spectrum licenses (McAfee and McMillan, 1996a), to devise railroad schedules (Brewer and Plott, 1995), and to trade electric power and long-term contracts for the supply of industrial chemicals (Chao and Wilson, 1996). These new competitive mechanisms permit decentralized decision-making in situations where, before, decisions were necessarily made centrally, and inefficiently.

References

- Aghion, Philippe, Dewatripont, Mathias, and Rey, Patrick, "Competition, Financial Discipline and Growth," unpublished, ENSAE, Paris, April 1995.
- Brewer, Paul, and Plott, Charles R., "A Binary Conflict Ascending Price Mechanism for the Decentralized Allocation of the Right to Use Railroad Tracks," unpublished, Caltech, Feb. 1995.

- Bulow, Jeremy, and Klemperer, Paul, "Auctions vs. Negotiations," *American Economic Review*, to appear, 1996.
- Chao, Hung-po and Wilson, Robert, "Multiple-Unit Auctions," unpublished, Stanford University, Nov. 1995.
- Debreu, Gerard, *Theory of Value*, New Haven, Yale University Press, 1959.
- Dewatripont, Mathias, and Maskin, Eric, "Credit and Efficiency in Centralized and Decentralized Economies," *Review of Economic Studies* 62, October 1995, 541-555.
- Gode, Dhananjay K., and Sunder, Shyam, "Allocative Efficiency of Markets with Zero-Intelligence Traders," *Journal of Political Economy* 101, Feb. 1993, 119-138.
- McAfee, R. Preston, and McMillan, John, "Auctions and Bidding," *Journal of Economic Literature*, June 1987, 25, 699-738.
- McAfee, R. Preston, and McMillan, John, "Analyzing the Airwaves Auction," *Journal of Economic Perspectives* 26, Winter 1996a, to appear.
- McAfee, R. Preston, and McMillan, John, "Electronic Markets," in *Readings in Electronic Commerce*, ed. Andrew Whinston. New York: Addison-Wesley, 1996b, to appear.
- McAfee, R. Preston, and Vincent, Daniel, "Sequentially Optimal Auctions," Northwestern University Discussion Paper No. 1104, 1995.
- McMillan, John, *Games, Strategies, and Managers*, New York, Oxford University Press, 1992.
- Milgrom, Paul R., "A Convergence Theorem for Competitive Bidding with Differential Information," *Econometrica* 47, 1979, 670-688.
- Milgrom, Paul, "Auctions and Bidding: A Primer," *Journal of Economic Perspectives*, Summer 1989, 3, 3-22.
- Roth, Alvin E., Prasnikar, Vesna, Okuno-Fujiwara, Masahiro, and Zamir, Shmuel, "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo," *American Economic Review* 81, Dec. 1991, 1068-96.
- Stole, Lars, and Zwiebel, Jeffrey, "Organizational Design and Technology Choice under Intrafirm Bargaining," *American Economic Review*, to appear, 1996.
- Vickers, John, "Concepts of Competition," *Oxford Economic Papers* 47, 1995, 1-23.
- Wilson, Robert B., "Strategic Analysis of Auctions." In Aumann, R. J., and S. Hart, eds., *Handbook of Game Theory with Economic Applications*, Vol. 1. Amsterdam: North-Holland, 1992, pp.227-280.