AN INTRODUCTION TO

THE THEORY OF CONTESTS

by

John G. Riley

University of California, Los Angeles

AN INTRODUCTION TO THE THEORY OF CONTESTS

by

John G. Riley

University of California, Los Angeles

Abstract

Over the last few years, economists have successfully applied the theory of contests to a broad range of econonic phenomena. In contrast to standard neoclassical models, the reward to an agent in a contest is dependent upon his relative, rather than absolute performance. For example, in patent races, bidding for an object d'art, or competing in the political arena for a franchise, it is the winner who takes all.

This paper examples several examples of contests and, in so doing, develops the central theoretical insights of the recent literature.

TABLE OF CONTENTS

Acknowledgment

The last decade has seen very significant progress in the analysis of strategic choice. One important strand of this recent literature examines the equilibrium behavior of agents competing for an indivisible payoff. Such "contests" take place in a wide range of environments. Interest groups compete for political favors. Firms enter Research and Development races knowing that there will be a monopoly patent for the winner. Animals compete for a mate or territory. And, as a final example, buyers compete in auctions for antiques and paintings.

In the following pages I have tried to illustrate the scope of the contest model and to present, in some detail, the ideas behind the central results. In so doing, my objective has been to provide not only the intuition behind the results but also insights into the way they are derived.

Section I explores a model of competition for political favors. Differences in preferences are known to all competitors however opponents' strategies are unobservable. Then, in Section II incomplete information is introduced not just about opponents' play but also about their preferences. The example studied is the "war of attrition." This has been used both by biologists and economists to explain strategic competition which takes place over time and which ends only when all but one contestant withdraw.

Section III turns to the theory of auctions. An extended example is examined in order to illustrate the contrast between equilibrium bidding in an open ascending bid auction and in a sealed high bid auction. Section IV then presents the general contest model and the assumptions commonly employed in the literature. This general structure is then used in Section V to place in perspective the central results in auction theory.

I. The Pursuit of Public Payoffs

Partly to introduce basic concepts and partly because it is of considerable independent interest, we begin with a simple model describing a contest for some political payoff. In its simplest form, n agents confront the opportunity of influencing the outcome of a political allocation decision. This decision will be to assign some valuable right (such as a franchise) to one of the n agents. Each of the agents must decide how large an outlay to make in an attempt to influence the outcome of the contest in his favor. In contrast with bidding in an auction, all the outlays are irretrievably lost. That is, they are used up in various promotions, fund raising events, etc.

Let agent i outlay x_i to influence the outcome of the political contest in his favor. The probability that agent i will be the successful contender is

$$P_i(x) = p_i(x_1, \ldots, x_n)$$

where

$$\sum_{i=1}^{n} p_i(x) = 1$$

and p_i is nondecreasing in x_i and nonincreasing in x_j, $j \neq i$.

Of central interest to analysts of such political competition or "rent seeking behavior" is the relationship between the value of the payoff and the total value of the outlays $\Sigma_{j=1}^{n} x_j$, given the different valuations v_i (i=1,...,n) which agents assign to winning the contest. To make the analysis as simple as possible we assume that the individual valuations v_i are known to all agents. We also label agents according to their valuations, that is $v_1 < v_2 < v_3 < \ldots < v_n$.

To evaluate the resource cost of political contestability, we require a specification for the probability function $p_i(x)$. Here we assume that contests are perfectly discriminating. That is, the political process awards the prize to the individual making the greatest outlay in seeking to influence the outcome in his favor. If more than one contender makes the highest outlay the winner is chosen at random. Formally, the probability of winning is then

$$p_i(x) = \begin{cases} 0 & \text{if } x_i \text{ is not a maximal element of} \\ & \{x_1, x_2, \ldots, x_n\} \\ \dfrac{1}{m} & \text{if } x_i \text{ is one of } m \text{ maximal elements of} \\ & \{x_1, \ldots, x_n\} \end{cases}$$

While we shall consider complete symmetry as a limiting case, we shall focus on the effect of asymmetric valuations on the strategies of the different agents.

Obviously this is rather an extreme model and Tullock (1980) has correctly argued that a better assumption is that the probability of winning rises continuously with an agent's outlay. Elsewhere (Hillman and Riley (1987)) it is argued that the basic insights are not critically dependent upon the assumptions of the model to be explored here.

In trying to predict the conclusions, the following two ideas make useful starting points. First of all, in the early literature on rent seeking (Posner (1975)), an analogy is made with bidding. Consider then open bidding for the prize, as in the usual ascending bid auction. As long as the asking price is below the second highest valuation, at least two agents have an incentive to remain in the bidding. Therefore the bidding continues until the price reaches the second highest valuation, v_{n-1}.

There is a second argument suggesting that total outlays might be v_{n-1}. No agent ever has an incentive to spend more than his valuation. Therefore agent n, by spending just a bit more than v_{n-1} can guarantee himself a payoff of $v_n - v_{n-1}$. The problem for agent n, however, is that each agent must choose his own outlay prior to observing others' outlays. That is, even though there is complete information about preferences, each agent must try to make an inference as to what his opponents' strategies will be. If agent n says he will spend v_{n-1}, another agent might reason as follows. "If I were to believe agent n there would be no point in competing. My best response would be to stay out of the contest. But, knowing this, agent n would then have an incentive to make only a very small outlay and achieve a much larger payoff. Since agent n has this incentive, he is just bluffing when he says he will spend v_{n-1}."

In game theoretic terms, the strategies chosen by each of the n agents are equilibrium strategies if and only if, for all i, agent i's strategy is his best response given the strategies of the other n-1 agents. For our example, each of the other agents' best response to agent n's strategy of spending v_{n-1} is to stay out of the contest. But then agent n's best response to agents 1,...,n-1 is to spend some very small amount and not v_{n-1}. From all this it seems reasonable to conclude that agents other than agent n will make positive outlays and that agent n will not spend as much as v_{n-1}.

The next point to be made is a technical one. In order to write down a mathematical expression for each agent's expected payoff, it is helpful to be able to rule out certain strategies. Specifically, we now argue that no agent will, in equilibrium, ever spend a positive amount β with a strictly positive probability. That is, equilibrium strategies are continuous mixed

strategies.

To see this, suppose agent i does spend β with strictly positive probability. Then the probability that a rival agent j beats agent i rises discontinuously as a function of x_j at $x_j = \beta$. Therefore there is some $\epsilon > 0$ such that agent j will bid on the interval $[\beta-\epsilon, \beta]$ with zero probability, for all $j \neq i$. But then agent i is better off spending $(\beta-\epsilon)$ rather than β since his probability of winning is the same, contradicting the hypothesis that $x_i = \beta$ is an equilibrium strategy.

Given this result, it follows immediately that if there are just two agents, they must have the same maximum spending level. For if $\bar{x}$ is agent 1's maximum spending level, agent 2 wins with probability 1 by spending $\bar{x}$ and vice versa.

A similar argument establishes that the minimum spending level is zero for each agent. To see this, suppose to the contrary that agent i spends less than β with zero probability, where $\beta > 0$. Then any spending level between zero and β yields a negative payoff since the probability of winning is zero. Since other agents can always spend zero it follows that no other agent will spend in the interval $(0,\beta)$. But then agent i could reduce his spending level below β without altering the probability of winning, contradicting the hypothesis that agent i could, in equilibrium, do no better than take β as his minimum spending level.

Given these results, if we define $1 - G_i(x_i)$ to be the probability that agent i spends more than x_i, then $G_i(x_i)$ is continuous over $(0,\infty)$. If $0 < G_i(0) < 1$ then agent i spends a strictly positive amount with probability less than 1. His remaining alternatives are to spend zero or stay out of the contest. Where it is important to make the distinction between these two alternatives we shall do so. One point to be made is that

if one agent spends zero with positive probability, no other agent will do so. The argument is similar to those made earlier. If two agents spend zero with positive probability they force a tie with positive probability. With an arbitrarily small positive bid one agent can therefore increase his win probability by a finite amount.

To summarize the above arguments, we have established that

(1.1) $G_i(x_i)$ is continuous over $(0, \infty)$

(1.2) $G_i(x_i) > 0$ for all $x_i > 0$

(1.3)   With only 2 active agents the maximum spending levels are the same

(1.4)   At most one agent spends zero with strictly positive probability.

We are now ready to characterize the equilibrium bidding strategies. To begin, we consider the two agent case with $v_1 < v_2$. Given condition (1.1), for all positive spending levels the probability of a tie is zero therefore agent 1's expected payoff is

(1.5) $U_1 = \left\{ \begin{matrix} \text{probability} \\ \text{of winning} \end{matrix} \right\} \left\{ \begin{matrix} \text{value of} \\ \text{prize} \end{matrix} \right\} - \left\{ \begin{matrix} \text{spending} \\ \text{level} \end{matrix} \right\}$

$$= G_2(x_1)v_1 - x_1 \ .$$

Similarly, agent 2's expected payoff is

(1.6) $U_2 = G_1(x_2)v_2 - x_2.$

Since agent 1 has the option of staying out of the contest, he will, in equilibrium, never spend more than his valuation and so incur a sure loss. Agent 2 can therefore guarantee himself a profit of $v_2 - v_1$ by spending v_1. It follows that agent 2's equilibrium expected payoff must be at least $v_2 - v_1$. Then, from (1.6), $G_1(0) > 0$. Moreover, since agent 2's expected payoff is strictly positive, he enters the contest with probability 1. We

now show that the equilibrium expected payoff for agent 1 must be zero. For, if not, from (1.5)

$$G_2(0) > 0.$$

Moreover, with a strictly positive expected payoff, agent 1 also enters the contest with probability 1. But with both $G_1(0)$ and $G_2(0)$ strictly positive and both agents always competing, both agents must spend zero with strictly positive probability. But this contradicts condition (1.4). Then agent 1's equilibrium payoff is zero and so, from (1.5)

$$(1.7) \qquad G_2(x_2) = \frac{x_2}{v_1}, \qquad x_2 \in [0,v_1]$$

That is, agent 1's equilibrium mixed strategy is to spend according to the uniform distribution over $[0,v_1]$.

Since both agents have the same maximum spending level v_1, we know that v_1 is in support of agent 1's bid distribution. Moreover $G_1(v_1) = 1$. Setting $x_2 = v_1$ in (1.6) we obtain

$$(1.8) \qquad U_2 = G_1(x_2)v_2 - x_2 = v_2 - v_1$$

Rearranging, it follows that agent 1's equilibrium mixed strategy is

$$(1.9) \qquad G_1(x_1) = \left(1 - \frac{v_1}{v_2}\right) + \frac{x_1}{v_2} = \left(1 - \frac{v_1}{v_2}\right) + \left(\frac{v_1}{v_2}\right)\frac{x_1}{v_1}$$

Note that agent 1 makes a strictly positive bid with probability

$$1 - G_1(0) = v_1/v_2 < 1$$

The most natural interpretation of this result is that agent 1 stays out of the contest with probability $(1-v_1/v_2)$ and enters with probability v_1/v_2. From (1.9), conditional upon entering the contest, agent 1 also adopts a

uniform mixed strategy over the interval $[0, v_1]$.

It is now easy to compute expected total spending. Agent 2's spending is uniformly distributed on $[0, v_1]$ and so his expected spending is $(1/2)v_1$. Conditional upon entering, agent 1's spending is also $(1/2)v_1$. Multiplying the latter by the probability of entry, v_1/v_2, expected total spending is therefore

$$(1.10) \qquad E\{x_1 + x_2\} = \frac{1}{2} v_1 + \frac{1}{2} v_1 \left(\frac{v_1}{v_2}\right) = \frac{v_1}{2} \left[1 + \frac{v_1}{v_2}\right]$$

We conclude, therefore, that only for the limiting case, when differences in valuations approach zero, is it true that expected total spending equals the second highest valuation.

Having examined the 2 agent case it is easy to characterize equilibrium with more than two agents. Suppose agent 0 has a valuation v_0 where $v_0 < v_1 < v_2$. Suppose agents 1 and 2 assume that agent 0 will always remain inactive. They will therefore spend according to (1.7), and (1.10). We now show that, given such spending behavior, agent 0 will indeed wish to remain out of the contest.

If he spends x_0 his expected payoff is

$$U_0(x_0) = \text{Prob}\{x_1 \text{ and } x_2 \text{ are both less than } x_0\}v_0 - x_0$$

$$= G_2(x_0)G_1(x_0)v_0 - x_0$$

$$= \left(\frac{x_0}{v_1}\right)\left[1 - \frac{v_1}{v_2} + \frac{x_0}{v_2}\right]v_0 - x_0$$

$$= \frac{v_0}{v_1} x_0\left[1 - \frac{v_1}{v_2} + \frac{x_0}{v_2} - \frac{v_1}{v_0}\right]$$

$$< 0 \qquad \text{for all} \qquad x_0 \in (0, v_0].$$

Since agent 0 will never spend more than his valuation, it follows that he is

strictly better off remaining inactive than entering the competition. Since this is true for all agents with lower valuations, we can conclude that, with n agents, if $v_n > v_{n-1} > v_{n-2} \geq \ldots \geq v_1$, equilibrium spending strategies are for the two higher valued agents to behave as if there were no other potential competitors and for all other agents to remain passive.

While we shall leave this example at this point, the obvious next question is whether there might be other equilibria in which more than two agents are active. In Hillman and Riley (1987) it is established that this is not the case. That is, we have characterized here the unique equilibrium.

II. The War of Attrition

Another contest in which winner and loser incur costs is one in which these costs are associated with the duration of the contest. As first applied by biologists, the competition is between animals competing for food, territory or a mate. Typically such contests involve ritualistic combat or "displays" rather than an all-out battle. Eventually one contestant concedes and leaves the other to enjoy the prize.

If the contest is of duration x, the rate of resource utilization is c and the discount rate is ρ, the present cost of entering the contest is

$$(2.1) \qquad C(x) = \int_0^x ce^{-\rho t}dt$$

If the prize is of value v, the present value of the payoff to the winner is

$$(2.2) \qquad e^{-\rho x}v - C(x)$$

while the present value to the loser is $-C(x)$.

For an economic application of this type of contest, consider two firms producing a product at a constant marginal cost of b and a fixed cost c per unit of time. With both firms in the market, and assuming price competition, the product is sold at marginal cost, that is

$$p(t) = b$$

Each firm thus loses the fixed cost c per unit of time.

With only one firm left in the market, this firm charges the monopoly price p^m. If we let v be the present value of the stream of profits earned by a monopolist at this price, we have a model which is identical to the biologist's war of attrition.[1]

In the simplest version of this model, and the one emphasized by the biologists (Maynard Smith (1976), Bishop and Cannings (1978)), competitors value the prize equally. The only uncertainty in the model is therefore strategic.

Here, however, we introduce further uncertainty by making the value of the prize private information. Each competitor's value is assumed to be an independent draw from a distribution with cumulative distribution function $F(v)$. We take the minimum valuation to be zero and the maximum to be $\bar{v}$ and assume that $F(v)$ is continuously differentiable over $[0,\bar{v}]$.

Given the symmetry of the model we seek a symmetric equilibrium. Arguing, essentially as in the previous section, no competitor will in equi- librium choose a particular duration $\hat{x}$ with strictly positive probability. It is tempting, then, to think that the equilibrium might require mixed strategies. However, given the informational asymmetry, any strictly

---

[1]The "war of attrition" has also been used in economics to explain the exit decision of firms in a declining industry (Fudenberg and Tirole 1983)).

monotonic mapping from a competitor's valuation (more generally, his "type") onto his action yields a distribution of duration levels which contains no mass points.

We therefore seek, as an equilibrium, such a mapping,

$$(2.3) \qquad x = B(v) \ .$$

Before doing so, a short digression is in order regarding the choice of an equilibrium concept. Certainly when modelling animal behavior it is hardly appropriate to rely on the calculating optimization implicit in the standard (Nash) noncooperative equilibrium. In economics, also, it is often argued that profit maximizing decisions are made because it is profit-maximizers which survive in the marketplace.

Consider then two populations. When a member from each meet in competition, the population 1 agent adopts the strategy s, from among some feasible set of possible strategies S. Similarly the population 2 agent (hereafter, agent 2) adopts some strategy $t \in T$. Let $u(s|t)$ be the expected payoff to agent 1 playing s against t and let $v(t|s)$ be the expected payoff to agent 2 playing t against s.

The pair of strategies $(\underline{s},\underline{t})$ is evolutionary stable with respect to a particular deviant strategy s if, when the proportion, p, adopting the deviant strategy is small, the deviants have a lower expected payoff than the nondeviants. As a result the survival rate of the nondeviants is higher and the proportion of deviants declines. The pair $(\underline{s},\underline{t})$ is an _evolution-ary equilibrium_ if it is evolutionarily stable with respect to all feasible deviant strategies.

With two different populations the difference in expected payoffs for population 1 when a proportion p adopt s is

$$D_1(s;p) = u(s|\underline{t}) - u(\underline{s}|\underline{t}) \ .$$

Similarly, if a proportion p of population 2 adopts strategy t, the difference in expected payoffs in population 2 is

$$D_2(t;p) = v(t|\underline{s}) - v(\underline{t}|\underline{s})$$

It follows immediately that if $(\underline{s},\underline{t})$ is a strong Nash equilibrium in the sense that $\underline{s}$ is the unique best reply to $\underline{t}$ and vice-versa, both D_1 and D_2 are negative for all $s \neq \underline{s}$ and $t \neq \underline{t}$.

With a single population the argument is only slightly more complicated. If the proportion adopting strategy s is p, a competitor adopting the deviant strategy finds himself competing against another deviant with probability p. The expected payoff of a deviant is therefore

$$pu(s|s) + (1-p)u(s|\underline{s}) \ ,$$

while the expected payoff of a nondeviant is

$$pu(\underline{s}|s) + (1-p)u(\underline{s}|\underline{s}) \ .$$

The net "advantage" of the deviant strategy is therefore

(2.4) $\qquad D(s;p) = (1-p)[u(s|\underline{s})-u(\underline{s}|\underline{s})] + p[u(s|s) - u(\underline{s}|s)] \ .$

If $\underline{s}$ is a strong Nash equilibrium the first bracketed expression is strictly negative for each deviant strategy s. Therefore $D(s;p)$ is negative for sufficiently small p. It follows that for $\underline{s}$ to be an evolutionary equilibrium strategy it is again sufficient for $\underline{s}$ to be a strong Nash equilibrium.

Essentially the same argument establishes that an evolutionary equilibrium is necessarily a Nash equilibrium. For, if not, the first bracketed expression is positive for some $\underline{s}$ and so $D(s;p) > 0$ for small p.

Finally, if $\underline{s}$ is a Nash equilibrium and, if for all s,

$$u(s|\underline{s}) - u(\underline{s}|\underline{s}) \Rightarrow u(s|s) < u(\underline{s}|s)$$

it follows from equation (2.4) that $\underline{s}$ is again evolutionarily stable. The intuition is clear. Even if a deviant has the same expected payoff as a nondeviant when competing with a nondeviant, the deviant will operate under a disadvantage if it is hurt more when it competes with another deviant.

We now return to the war of attrition. Since the equilibrium that we examine is a strong Nash equilibrium it is also an evolutionary equilibrium. Rather than attempt to characterize the equilibrium strategy $x = B(v)$ directly, it is helpful to define the inverse mapping

$$(2.5) \qquad v = \phi(x) = B^{-1}(x)$$

Since an agent's valuation is a draw from the distribution $F(\cdot)$, his implied distribution of waiting times is

$$(2.6) \qquad G(t) = F(\phi(t)) .$$

Then if agent 2 adopts the equilibrium strategy, while agent 1 plans to drop out at time x, agent 1's expected payoff is

$$U_1(x) = - (1-G(x))C(x) + \int_0^x e^{-\rho t}(v-C(t))G'(t)dt$$

Differentiating by x, agent 1's best reply is to choose $x*$ so that

$$(2.7) \qquad U_1'(x) = -(1-G(x))ce^{-\rho x} + ve^{-\rho x}G'(x)$$
$$= 0$$

That is,

$$(2.8) \qquad \frac{G'(x*)}{1 - G(x*)} = \frac{c}{v}$$

But, for $B(\cdot)$ to be the equilibrium, it must be the case that $x^* = B(v)$ and hence $v = \phi(x^*)$. Substituting for v in (2.8) we obtain the following necessary condition for an equilibrium

$$(2.9) \qquad \frac{\phi(x)G'(x)}{1 - G(x)} = c .$$

Substituting from (2.6) we obtain

$$\frac{\phi F'(\phi)}{(1 - F(\phi))} \frac{d\phi}{dx} = c$$

Integrating the expression we have finally

$$(2.10) \qquad x = B(v) = \int_0^v \frac{\phi F'(\phi)}{c(1-F(\phi))} \, d\phi$$

To complete the argument we must show that the necessary condition, (2.9), is also sufficient. Substituting (2.9) into (2.7)

$$U_1'(x) = e^{-\rho x}G'(x)[v-\phi(x)]$$

Since $\phi(x)$ is increasing we have

$$U_1'(x) > 0, \qquad x < v$$
$$< 0, \qquad x > v$$

Therefore choosing $x = B(v)$, as given by equation (2.10) is agent 1's unique best reply. It follows that we have characterized a strong Nash equilibrium and hence an evolutionary equilibrium.

It is perhaps puzzling that the equilibrium strategy is independent of the discount rate. To understand this better, consider the following informal characterization of the necessary condition. If an agent with valuation v has a best reply of dropping out at x, he must be indif-

ferent between dropping out at x and staying in just a brief time Δt

longer. The conditional probability that this competitor will drop out over

this time interval is $G'(x)\Delta t/(1-G(x))$. The net expected payoff to staying

beyond time x is therefore

$$\frac{vG'(x)}{1-G(x)} \Delta t - c\Delta t$$

Since this must be zero we obtain the necessary condition (2.8).[2]

For the biological war of attrition, it is natural to seek a symmetric

equilibrium, at least if competing animals are observationally equivalent.

However, in many circumstances, there will be characteristics which distin-

guish competitors. Within a species it might be age or even something as

simple as whether an animal is up- or down-wind of its opponent. As soon as

an observable difference is introduced, there is no compelling reason to

focus attention on symmetric equilibria. As we shall now see, the war of

attrition has a continuum of asymmetric equilibria in which one competitor

is more aggressive (with longer waiting times) and the other is more passive

than in the symmetric equilibrium.

Arguing exactly as before, if agent 2 adopts the strategy $x_2 = B_2(v_2)$,

agent 1 faces a distribution of waiting times

(2.11) $G_2(t) = F(\phi_2(t))$

where

$$\phi_2(t) = B_2^{-1}(t)$$

The necessary conditions for $\phi_1(t)$ to be the inverse of agent 1's

_____

[2]The expression $G'(x)/(1-G(x))$ is just the hazard rate at time x. In
equilibrium this must equal $c/\phi(x)$ for all x. Since $\phi(x)$ is strictly
increasing, the equilibrium hazard rate declines with the length of the
contest.

equilibrium best reply is then, from (2.9)

$$\frac{\phi_1(x)G_2'(x)}{1 - G_2(x)} = c \ .$$

Substituting from (2.11)

$$(2.12) \qquad \frac{\phi_1(x)F'(\phi_2(x))\phi_2'(x)}{1 - F(\phi_2(x))} = c \ .$$

An identical argument for agent 1 yields the second necessary condition

$$(2.13) \qquad \frac{\phi_2(x)F'(\phi_1(x))\phi_1'(x)}{1 - F(\phi_1(x))} = c \ .$$

Define

$$H(v,a) = \int_a^v \frac{F'(x)}{x(1-F(x))} \ , \qquad a \in (0,\bar{v})$$

Dividing (2.12) by (2.13) and rearranging we obtain,

$$\frac{\partial H}{\partial v}(\phi_1, a_1) = \frac{\partial H}{\partial v}(\phi_2, a_2), \quad \text{for all} \quad a_1, a_2 \in (0,\bar{v}).$$

Integrating we obtain finally

$$(2.14) \qquad H(\phi_1, a_1) = H(\phi_2, a_2), \quad a_1, a_2 \in (0,\bar{v}).$$

Clearly one solution is $a_1 = a_2$ and hence $\phi_1 = \phi_2$. This is the symmetric equilibrium defined above. However, as is demonstrated in Nalebuff and Riley (1985), equation (2.14) has a solution for all other parameter values as well. Here we illustrate with a simple example. Suppose valuations are distributed exponentially, that is,

$$F(v) = 1 - e^{-\alpha v}$$

Then $H(v,a) = \ln(v/a)$ and so equation (2.14) becomes

$$\frac{\phi_1(x)}{a_1} - \frac{\phi_2(x)}{a_2}$$

The equilibrium inverse bid function of one agent is therefore any linear function of the other agent's inverse bid function.

As economists, we usually feel frustrated by failure to either establish uniqueness or, in the case of multiple equilibria, pin down a particular equilibrium as being especially plausible. One reason for this is that it is hard to justify the assumption that opponents will choose their equilibrium strategies unless the equilibrium is unique. However, for the war of attrition, we have seen that each of the equilibria is a stationary point of a dynamic adjustment process. Therefore perhaps it should simply be accepted that, in such environments, quite different equilibria may emerge. Indeed, as is argued in Nalebuff and Riley, the absence of a unique equilibrium may well help to explain the remarkable variety of the natural world.

III. Open and Sealed Bid Auctions

We now begin a rather more detailed review of the theory of auctions. In this section we present some simple examples to illustrate equilibrium bidding behavior.

There are two very common methods used to sell an item which is infrequently traded. The first, used for the sale of painting and other objects d'art is the "open ascending bid" auction. Here the auctioneer calls out an ever higher asking price until there are no takers. The last bidder is then the winner.

The second type of auction, used in competitive bidding for private and government contracts is the "sealed high bid" auction. Here each buyer must

submit a sealed bid. The bids are then opened and the high bidder pays his bid in return for the item. In each auction, if the bidding ends in a tie, the winner is selected randomly.

To see that these two auctions are strategically very different, consider a situation in which valuations are entirely private. That is, buyer i has some valuation v_i which he would not change, even if he were to learn of others' valuations. In the open ascending bid auction there is a very simple strategy that he can follow. As long as the asking price is less than v_i, buyer i is better off remaining in the auction regardless of other buyers' strategies. Of course once the asking price reaches v_i, buyer i will drop out of the bidding.

Since this is true for each buyer, the equilibrium is for each buyer to take his valuation as his maximum price. As a result the seller receives the second highest valuation and the winning buyer has a payoff equal to the difference between his and this second highest valuation.

The situation is much more complicated in the sealed high bid auction. Buyer i is in the bidding hoping to make a profit. Therefore he will make a bid somewhat less than his valuation. Just how much depends on his beliefs about the valuations of his opponents. The more likely it is that they have low valuations, the greater is buyer i's incentive to lower his bid. But this is only part of the story. For buyer i will also recognize that other buyers will be shading their bids below their valuations. This further increases his incentive to lower his bid.

Beyond the issue of what constitutes an equilibrium bidding strategy in the sealed high bid auction, there is a further important issue facing the seller. Given his opportunity to choose the way the item is sold, which form of auction will generate a higher expected payment? As we shall see,

at least under some central simplifying assumptions, expected payments are the same under the two common auctions.

To understand the sealed high bid auction we begin with some very simple examples all involving just two buyers. Suppose first that each buyer's valuation v_i, $i = 1,2$ is either zero or unity. Let p_{ij} be the joint probability that buyer i has valuation v_i and buyer j, valuation v_j. Without loss of generality, we impose the further restriction,

(3.1) $$p_{10} \geq p_{01}$$

on the joint probability matrix

$$[p_{ij}] = \begin{bmatrix} p_{00} & p_{01} \\ p_{10} & p_{11} \end{bmatrix} .$$

Therefore, either beliefs are symmetric or buyer 1 is more likely to have a high value than buyer 2.

Since a buyer with a zero valuation has no incentive to bid we shall assume that he remains out of the auction. A buyer with a unit valuation will wish to bid less than 1 since there is a positive probability that his opponent has a zero valuation. However, arguing almost exactly as in Section I, no buyer will, in equilibrium, make some strictly positive amount bid $\hat{b}$ with positive probability. For if buyer 1 were to do so, buyer 2 would increase his probability of winning discretely if he were to bid just greater than $\hat{b}$. Then there is some interval $[\hat{b}-\epsilon,\hat{b}]$ over which buyer 2 would not bid. But then buyer 1 could lower his bid below $\hat{b}$ without altering his probability of winning and so $\hat{b}$ is not an equilibrium bid.

Moreover, while one buyer with a unit valuation might bid zero with positive probability, both would not do so. The argument is almost the

same. As long as buyer 1 is bidding zero with positive probability, buyer 2's expected payoff rises discontinuously at $b = 0$.

We therefore seek, as an equilibrium, a continuous mixed strategy $G^i(b)$, $i = 1,2$, for each player, where $G^i(0) = 0$ for at least one bidder.

Suppose buyer 2 is bidding according to his equilibrium mixed strategy. Then, if buyer 1, with unit valuation, bids b he wins for sure if buyer 2 has a zero valuation and with probability $G^2(b)$ if buyer 2 has a unit valuation. His expected payoff is therefore

$$(3.2) \qquad U^1(1,b) = (1-b) \left[\frac{P_{10}}{P_{10}+P_{11}} + \frac{P_{11}}{P_{10}+P_{11}} G^2(b) \right].$$

Similarly, if buyer 1 adopts his equilibrium mixed strategy, buyer 2's expected payoff, when he has a unit valuation, is

$$(3.3) \qquad U^2(1,b) = (1-b) \left[\frac{P_{01}+P_{11}G^1(b)}{P_{01}+P_{11}} \right].$$

If buyer 1's maximum bid is $\bar{b}$, buyer 2 has no incentive to bid higher. Moreover, if buyer 1's minimum bid is $\underline{b}$, buyer 2 has no incentive to bid lower. Since the same argument holds in reverse, it follows that the two buyers must have the same maximum and minimum bids.

If the minimum bid were strictly positive, buyer 1's expected payoff would be

$$U^1(1,\bar{b}) = (1-\bar{b}) \frac{P_{10}}{P_{10}+P_{11}}$$

since buyer 2 would make the bid $\underline{b}$ with probability zero. But buyer 1 would thus be strictly better off bidding less than $\underline{b}$ since he only wins against buyer 2 if the latter has a zero valuation. We conclude, therefore,

that the minimum bid must be zero.

In equilibrium each buyer must be indifferent between all his bids. Setting $b = 0$ and $\bar{b}$ in (3.2) and (3.3) we can therefore conclude that

$$(3.4) \qquad \frac{P_{10} + P_{11} G^2(0)}{P_{10} + P_{11}} = (1-b)\left[\frac{P_{10} + P_{11} G^2(b)}{P_{10} + P_{11}}\right] = 1 - \bar{b}$$

and

$$(3.5) \qquad \frac{P_{01} + P_{11} G^1(0)}{P_{01} + P_{11}} = (1-b)\left[\frac{P_{01} + P_{11} G^1(b)}{P_{01} + P_{11}}\right] = 1 - \bar{b}.$$

We have already argued that only one buyer can bid zero with positive probability. Since $P_{10} \geq P_{01}$, it follows from (3.4) and (3.5) that

$$(3.6) \qquad 0 = G^2(0) \leq G^1(0)$$

and, hence, that

$$(3.7) \qquad \bar{b} = 1 - \frac{P_{10}}{P_{10} + P_{11}} = \frac{P_{11}}{P_{10} + P_{11}}.$$

Combining (3.4), (3.5) and (3.7), one can then solve for the equilibrium mixed strategies. We obtain

$$(3.8) \qquad \begin{cases} G^1(b) = \left[\dfrac{P_{10}}{P_{10} + P_{11}}\right]\left[\dfrac{P_{01} + P_{11}}{P_{11}}\right]\left(\dfrac{1}{1-b}\right) - \dfrac{P_{01}}{P_{11}} \\[3mm] G^2(b) = \dfrac{P_{10}}{P_{11}}\left(\dfrac{b}{1-b}\right) \end{cases}$$

The two distribution functions are illustrated in Figure 1. Note, in particular, that the bid distribution for buyer 2 exhibits first order stochastic dominance over that for buyer 1. That is, the buyer who is more likely to have a higher valuation and hence is up against an opponent who is more likely to have a lower valuation, tends to bid less aggressively.
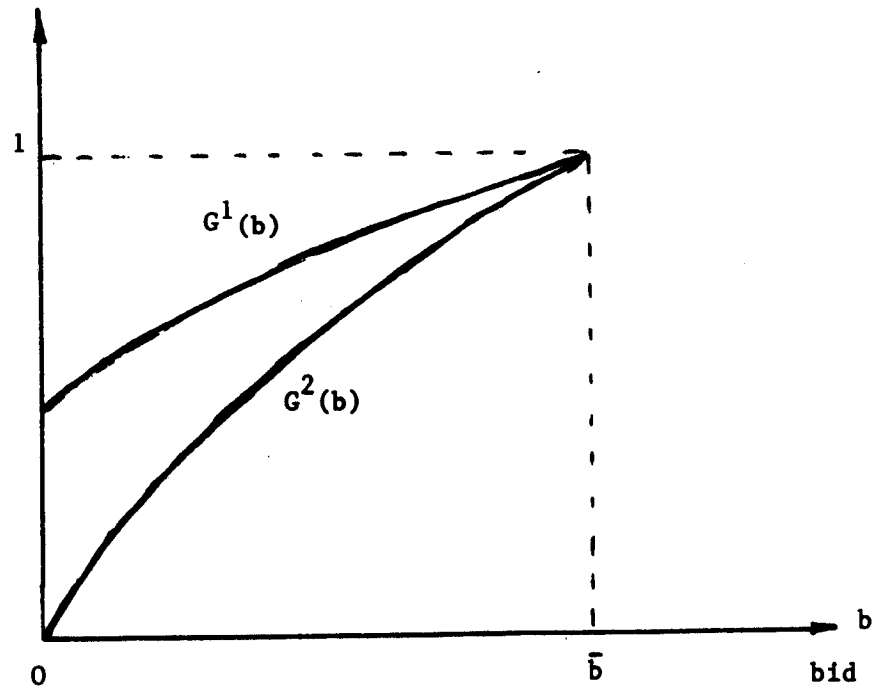
Figure 1: Equilibrium mixed strategies in the sealed high bid auction

The reason for this is that the more likely it is that the opponent has a low valuation, the greater is the incentive to make a very small bid and so win against a low valued buyer.

While we shall return to the asymmetric model later, we now focus on the symmetric case, that is, the joint probability matrix is symmetric. Then, setting $P_{10} = P_{01}$ in (3.8) the equilibrium bidding strategies of the two agents are the same, that is,

$$G^1(b) = \frac{P_{10}}{P_{11}} \left(\frac{b}{1-b}\right), \quad b \in [0,\bar{b}].$$

To obtain an expression for the expected payment by a high value buyer we rewrite (3.2) as

$$\text{payoff when bid is } b = U^1(1,b) = 1 \text{ Prob} \left\{\begin{matrix}\text{win and}\\\text{pay } b\end{matrix}\right\} - b \text{ Prob} \left\{\begin{matrix}\text{win and}\\\text{pay } b\end{matrix}\right\}.$$

Taking the expectation over bids

$$\text{Expected payoff} = \underset{b}{E} U^1(1,b) = 1 \text{ Prob} \left\{\begin{matrix}\text{buyer 1}\\\text{wins}\end{matrix}\right\} - \text{Expected Payment}.$$

But, from (3.7) the equilibrium payoff is $1-\bar{b} = P_{01}/(P_{01}+P_{11})$. Moreover, a high value buyer 1 always wins against a low value opponent and, by symmetry, half the time against a high value opponent. Therefore

$$\frac{P_{01}}{P_{01}+P_{11}} = \frac{P_{01} + \frac{1}{2} P_{11}}{P_{01}+P_{11}} - \text{Expected Payment}$$

and so the expected payment by a high value buyer is

$$\frac{1}{2} \left(\frac{P_{11}}{P_{01}+P_{11}}\right).$$

We now compare this with the expected payment by a high value buyer in the open ascending bid auction. In the latter, the bidding stops at zero unless a high value buyer is up against an opponent who also has a high valuation. This occurs with probability $P_{11}/(P_{01}+P_{11})$. In this case the bidding proceeds until the price reaches unity and then the winner is selected randomly. The probability that a particular buyer wins and pays 1 is therefore 1/2 and so his expected payment is

$$\frac{1}{2} \left(\frac{P_{11}}{P_{01}+P_{11}} \right)$$

that is, exactly as in the sealed high bid auction.

Of course the two point example is rather special and, as we shall see, revenue equivalence does not generally obtain except in the two point case.

A further question when buyers have more than two possible valuations is whether or not agents will necessarily bid monotonically. To understand the issues involved we consider a case in which there are only three possible valuations, v_0, v_1, v_2, where

$$0 = v_0 < v_1 < v_2.$$

The underlying joint probability distribution,

$$\begin{bmatrix} P_{00} & P_{01} & P_{02} \\ P_{10} & P_{11} & P_{12} \\ P_{20} & P_{21} & P_{22} \end{bmatrix}$$

is taken to be symmetric, thereby preserving the symmetry of the auction. Moreover, reflecting the idea that a buyer with a higher valuation himself will assign higher probabilities to higher valuations for his opponent, we assume that the following condition holds.

Definition: Conditional Stochastic Dominance

$$\text{Prob} \left\{ \begin{array}{l} \text{buyer 1's valuation} \\ \text{is less than } \alpha \end{array} \middle| \begin{array}{l} \text{buyer 1's valuation is \underline{less than}} \ \beta \ \text{and} \\ \text{buyer 2's valuation is \underline{equal to}} \ \gamma \end{array} \right\}$$

is, for all α, β and γ nonincreasing in γ.

In particular, for the three point example, conditional stochastic dominance implies that

$$\text{Prob} \left\{ \begin{array}{l} \text{buyer 2's valuation} \\ \text{is less than } \alpha \end{array} \middle| \begin{array}{l} \text{buyer 2's valuation is less than } 1+\alpha \\ \text{buyer 1's valuation} = 1 \end{array} \right\}$$

$$\geq \text{Prob} \left\{ \begin{array}{l} \text{buyer 2's valuation} \\ \text{is less than } \alpha \end{array} \middle| \begin{array}{l} \text{buyer 2's valuation is less than } 1+\alpha \\ \text{buyer 1's valuation} = 2 \end{array} \right\}$$

Setting $\alpha = 1,2$ we obtain

$$(3.9) \qquad \frac{P_{10}}{P_{10}+P_{11}} \geq \frac{P_{20}}{P_{20}+P_{21}} \quad \text{and} \quad \frac{P_{10}+P_{11}}{P_{10}+P_{11}+P_{12}} \geq \frac{P_{20}+P_{21}}{P_{20}+P_{21}+P_{22}} \ .$$

Given the symmetry of the model we seek a symmetric equilibrium. Let $G_t(b)$ be the equilibrium c.d.f. for a buyer with valuation v_t. We begin by supposing that equilibrium bidding in the sealed high bid auction is monotonic, as depicted in Figure 2. If a type 1 buyer bids $b \in [0,\bar{b}_1]$ against an opponent adopting the equilibrium bidding strategy, his expected payoff is

$$(3.10) \qquad \bar{U}_1(v,b) = (v_1-b) \ \frac{P_{10}+P_{11}G_1(b)}{P_{10}+P_{11}+P_{12}}.$$

The equilibrium mixed strategy of a type 1 buyer is then determined completely by the requirements that $\bar{U}_1(v_1,b)$ is constant over $[0,\bar{b}_1]$ and that $G_1(0) = 0$. Similarly, if a type 2 buyer bids $b \in [\bar{b}_1,\bar{b}_2]$ against an
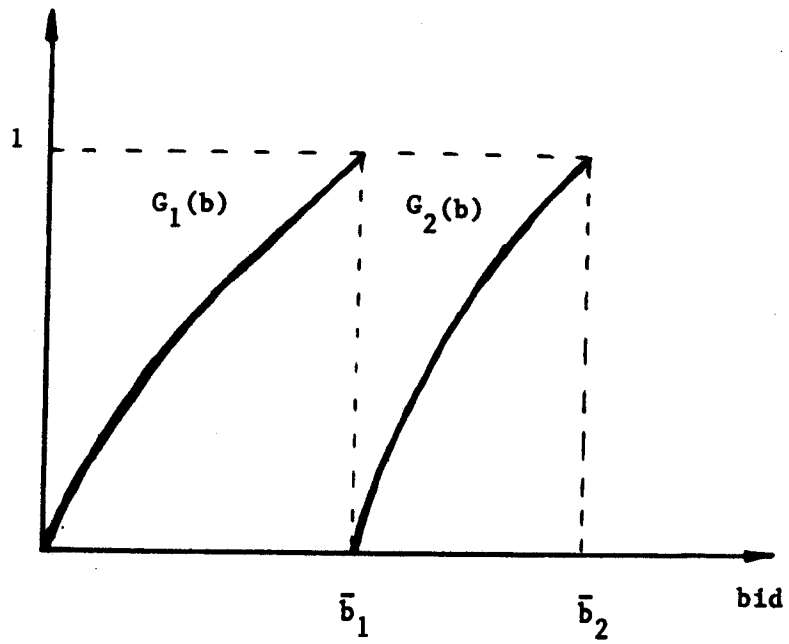
Figure 2:   Monotonic bidding strategies in the
            Sealed high bid Auction

opponent adopting the equilibrium bidding strategy, his expected payoff is

$$(3.11) \qquad U_2(v_2,b) = (v_2-b) \frac{P_{20}+P_{21}+P_{22}G_2(b)}{P_{20}+P_{21}+P_{22}} \; .$$

The equilibrium strategy for a type 2 buyer is then completely determined essentially as for a type 1 buyer.

To complete the analysis we must confirm that a type 1 buyer would be strictly worse off bidding in the interval $(\bar{b}_1, \bar{b}_2)$ and that a type 2 buyer would be strictly worse off bidding less than $\bar{b}_1$. If a type 1 buyer bids more than $\bar{b}_1$ against a buyer adopting the proposed equilibrium strategy his expected payoff is

$$(3.12) \qquad U_1(v_1,b) = (v_1-b) \frac{(P_{10}+P_{11}+P_{12}G_2(b))}{P_{10}+P_{11}+P_{12}}, \quad b > \bar{b}_1 \; .$$

Dividing (3.12) by (3.11) we obtain

$$(3.13) \qquad \frac{U_1(v_1,b)}{U_2(v_2,b)} = \left[\frac{v_1-b}{v_2-b}\right] \left[\frac{\alpha_1+(1-\alpha_1)G_2(b)}{\alpha_2+(1-\alpha_2)G_2(b)}\right], \quad b > \bar{b}_1$$

where

$$\alpha_i = \frac{P_{10}+P_{11}}{P_{10}+P_{11}+P_{12}} \; .$$

It is readily verified that, since $v_1 < v_2$, the first bracketed expression is strictly decreasing in b. Moreover, from (3.9), $\alpha_1 \geq \alpha_2$. Since G_2 is increasing in b it follows that the second bracketed expression is nonincreasing in b. Therefore $U_1(v_1,b)/U_2(v_2,b)$ strictly decreases with b for $b > \bar{b}_1$. But $U_2(v_2,b)$ is constant over this interval and so $U_1(v_1,b)$ strictly decreases with b. It follows that all bids greater than $\bar{b}_1$ yield a strictly lower expected payoff to a type 1 buyer than $\bar{b}_1$.

Therefore, a type 1 buyer is strictly worse off attempting to bid against a type 2 buyer.

An almost identical argument can be used to show that, given conditional stochastic dominance, a type 2 buyer is also strictly worse off bidding below $\bar{b}_1$. The monotonic bidding strategy is therefore an equilibrium bidding strategy as claimed.

We conclude this section by noting that, in the absence of the conditional stochastic dominance assumption, there is no reason to expect bidding to be monotonic in type. Consider the following joint probability matrix.

$$\begin{bmatrix} 0.2 & 0.2 & 0.1 \\ 0.2 & 0.2 & 0 \\ 0.1 & 0 & 0 \end{bmatrix}.$$

In this example, if one buyer has a high valuation he knows that his opponent has a zero valuation. The zero valuation buyer has no incentive to bid so the high value buyer's best reply is to bid zero. However a buyer with an intermediate value (v_1) knows he is equally likely to be opposed by a type zero or type 1 buyer. He therefore adopts a mixed strategy exactly as described at the beginning of the section.

IV. The "General" Contest Model

In this section we outline a rather more general model of a contest. Each of the examples of the previous section and essentially all of the models in the rapidly growing literature on auctions are special cases.

There are n competing agents, each of whom is willing to make a nonnegative payment in order to receive a single prize. Agent i, $i = 1,\ldots,n$ receives a private signal $s_i \in S_i$ prior to the start of the contest. There may also be a signal $s_0 \in S_0$ which (eventually) becomes

public information. Agent i's prior beliefs about the distribution of signals is given by the cumulative distribution function $H^i(s_0, \ldots, s_n)$, $i = 1, \ldots, n$. If he has an initial wealth w and incurs cost x and wins the contest, agent i's expected utility is

$$U^i = U^i(w-x, s_0, \ldots, s_n), \quad i = 1, \ldots, n.$$

That is, in general we allow agent i's valuation of the prize to depend not only on his own signal but also others' private signals as well.

If agent i does not win the contest and incurs a cost x, his expected utility is

$$\bar{U}^i = \bar{U}^i(w-x).$$

Simultaneously, each agent, using only his own private information (and the public signal s_0, if available), must decide how large a cost to incur in his pursuit of the prize. The rules of the contest are summarized by n probability functions

$$(4.1) \qquad p^i = p^i(x_1, \ldots, x_n)$$

which are nonnegative and sum to unity. Moreover, $p^i(x)$ is nondecreasing in x_i and nonincreasing in $x_{-i} = (x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n)$.

To illustrate the applicability of this more general structure, consider n agents bidding for an oilfield tract. Agent i incurs a research cost k to obtain an estimate s_i of the value of the oil in the ground. This estimate is distributed around the true value θ according to the density function $f(s_i|\theta)$. Prior beliefs about θ are given by the density $f_0(\theta)$. Agent i's joint density function for the n signals is then

$$h(s) = \int [\prod_{j=1}^{n} f(s_j|\theta)] f_0(\theta) d\theta$$

and the density function for θ conditional upon signal levels $s_1, \ldots, s_n$ is

$$(4.2) \qquad g(\theta|s) = \frac{[\prod_{j=1}^{n} f(s_j|\theta)]f_0(\theta)}{h(s)}$$

Finally, suppose each agent has an identical utility function $u(\cdot)$. Using (4.2) if agent i bids b in a sealed high bid auction and he is successful his expected utility is

$$(4.3) \qquad U^i = U(w-b,s) = \int u(w-k-b+\theta)g(\theta|s)d\theta.$$

If, on the other hand, b is not the winning bid, agent i's utility is

$$(4.4) \qquad \bar{U}^i = u(w-k).$$

Actually, the "oil field" problem discussed most often in the literature contains one further simplification. Bidders are risk neutral so that (4.3) can be rewritten as

$$(4.5) \qquad U^i = v^i(s) - b - k + w$$

where

$$(4.6) \qquad v^i(s) = E\{\theta|s\} = \int \theta g(\theta|s)d\theta.$$

In order to extract results from the general model, a variety of simplifying assumptions are typically introduced. These can be partitioned into three classes. At the most basic level there are three monotonicity assumptions.

A. **Monotonicity Assumptions**

A1: Monotonicity of Preferences

$U^i(w-x,s)$, $\bar{U}^i(w-x)$ are nondecreasing and strictly increasing in their first arguments.

A2: Absence of Risk Loving

$$\frac{\partial}{\partial w} U^i, \frac{\partial}{\partial w} \dot{U}^i \text{ are nonincreasing functions.}$$

A3: Conditional Stochastic Dominance

Let t be a permutation of the signal vector s. Then for every

$\alpha = (\alpha_0, \ldots, \alpha_n)$, β and m and every agent

$\text{Prob}(t_0 \leq \beta \mid t_0 \leq \alpha_0, t_1 \leq \alpha_1, \ldots, t_m \leq \alpha_m, t_{m+1} = \alpha_{m+1}, \ldots, t_n = \alpha_n)$

is nonincreasing in $\alpha_1, \ldots, \alpha_n$.

Assumption A1 is a minimal restriction which requires that an agent's preference ordering of the signals is not altered by a change in his wealth. A higher signal is one which yields a higher utility.

Assumption A2 says that the marginal utility of wealth declines (or remains constant) as wealth or any one of the signals increases. Given aversion to wealth risk it seems entirely plausible to impose this stronger assumption. Indeed, for the oil field example, aversion to wealth risk implies A2.

Assumption A3 is the generalization of the Conditional Stochastic Dominance assumption discussed in Section III. Essentially it says that if an agent obtains more favorable information about one signal, he will tend to believe that other signals are more likely to be favorable as well. In the limiting case of **independent** signals more favorable information about one signal has no effect on an agent's beliefs about the other signals.

B. **Symmetry Assumptions**

B1: Symmetry of Preferences

$U^i = U(w-x, s_i, s_{-i})$, $s_i \in S$ and U is symmetric in s_{-i}.

B2: Symmetry of Beliefs

$H^i(s_0, \ldots, s_n) = H(s_0, s_1, \ldots, s_n)$ is symmetric in $(s_1, \ldots, s_n)$.

B3: Symmetry of Equilibrium

The equilibrium mapping from signal to resource cost incurred in the contest is the same for each agent, that is $x_i = B(s_i)$, $s_i \in S$.

While Assumptions B1 and B2 are certainly strong, they are the natural first approximation. Assumption B1 makes agents identical prior to the receipt of their private information. Assumption B2 requires that if two different agents in fact receive the same private signal, they will have the same beliefs about other agents' signals.

Given Assumptions B1 and B2, the assumption of symmetric bidding behavior also seems reasonable. However, as the analysis at the end of Section II shows, there are examples of symmetric contests which have a continuum of asymmetric equilibria.

The third set of assumptions are also natural first approximations from a technical viewpoint. However, their relevance depends very much on the nature of the application.

C. Separability Assumptions

C1: Risk Neutrality

$$U^i = v^i(s_i, s_{-i}) - b$$

C2: Independence

$$H^i(s_0, \ldots, s_n) = F_0^i(s_0)F_1^i(s_1)\ldots F_n^i(s_n)$$

C3: Private Values

$$U^i = U^i(w-b, s_i).$$

As long as relative risk aversion does not exceed unity, an increase in

wealth of 1% has no more than a 1% effect on marginal utility. Therefore, the assumption of risk neutrality (constant marginal utility of wealth) is a reasonable approximation whenever the value of the prize is small relative to each agent's total wealth.

The independence assumption is a strong one. Clearly it is inappropriate in "common value" contests such as the oil field auction described above. However, for other applications such as the war of attrition, or in the auctioning of art and antiques, it is justifiable, again as a first approximation.

Assumption C3, private values, is a further strong restriction. Even in an art auction, the winner's utility will be influenced by others' bids if there is a chance that he may later wish to resell the painting he has just purchased.

We conclude this section with some observations about the nature of the probability functions $p^i(x)$, $i = 1,\ldots,n$.

In all the contests examined so far, the winner is always the agent who incurs the greatest cost. That is,

$$
P_i(x) = \begin{cases} 0, & x_i < \max\{x_1,\ldots,x_n\} \\ \dfrac{1}{m}, & x_i \text{ is one of } m \text{ maximal elements of } \{x_1,\ldots,x_m\} \end{cases}
$$

is a discontinuous function of each agent's cost.

As noted in Section I, there are applications for which a continuous probability function seems more natural. Consider again the political competition model. Plausibly the political impact, z_i, of an agent's promotional activities is less than perfectly correlated with his expenditures x_i. Suppose, then that

$$z_i = A(x_i, \epsilon_i), \quad i = 1, \ldots, n$$

where A is an increasing function and ϵ_i is independently and identically distributed. If it is the agent with the largest political impact, rather than the one who spends the most, who is the winner, the implied win probability is continuous in x_i.

For example, suppose $n = 2$,

$$z_i = x_i \epsilon_i$$

and that ϵ_i has cumulative distribution function

$$F(\epsilon_i) = 1 - e^{-a\epsilon_i}.$$

Given spending levels of x_1 and x_2 agent 1 wins if

(4.7)
$$x_2 \epsilon_2 < x_1 \epsilon_1.$$

For any given ϵ_2 the probability that (4.7) holds is

$$\Pr\left\{\epsilon_1 > \left(\frac{x_2}{x_1}\right)\epsilon_2\right\} = e^{-ax_2\epsilon_2/x_1}.$$

Taking the expectation over ϵ_2

$$P_1(x_1, x_2) = a \int_0^\infty e^{-ax_2\epsilon_2/x_1} e^{-a\epsilon_2} d\epsilon_2 = \frac{x_1}{x_1 + x_2}.$$

The win probability is thus equals to the agent's share of the total amount spent in pursuit of the political payoff.

V. Auction Design

In this final section we summarize some of the central results in auction theory. This literature has focused primarily on the choice of an auction by a risk neutral seller. The starting point is the <u>Revenue Equivalence Theorem</u>. In its original form (Vickrey (1961)) this states that if

Assumptions A, B and C of Section IV are satisfied, then expected revenue from the open ascending bid auction and the sealed high bid auction is the same.

Given risk neutrality and private values, the expected payoff to buyer i becomes

$$U^i = v^i - b + w$$

where

$$v^i = v(s_i).$$

Given independence of the signals, we also have independence of valuations. The resulting model is therefore exactly that which we considered in Section III, with the additional independence restriction.

Using the discrete version of this model it is not difficult to see why the revenue equivalence holds. First we note that, in each auction, the expected payoff to type v_t is

$$U_t = v_t Pr\{v_t \text{ is a winner}\} - \text{Expected payment by type } t.$$

In each auction type t wins against lower types and wins half the time against his own type. Given Assumption C2, independence, we can write the probability that an opponent is of type j as p_j. Then, in equilibrium,

$$(5.1) \qquad U_t = v_t \left(\sum_{j=0}^{t-1} p_j + \frac{1}{2} p_t \right) - \text{Expected payment by type } t.$$

We shall now argue that the expected payoffs are the same in the two auctions. It will then follow directly from (5.1) that expected payments are also identical, hence, revenue equivalence.

We begin with the open auction. Type t and type $t+1$ always win if up against types $0, 1, \ldots, t-1$. In each case the selling price is the same so the only difference is the difference in the buyers' valuations ($v_{t+1} - v_t$). If the opponent is of type t there is a zero profit for type t and since the selling price is v_t, a profit of ($v_{t+1} - v_t$) for type $t+1$. Thus again the difference in payoffs is the difference in valuations.

When up against types higher than $t+1$, neither buyer type can make a profit. Therefore the difference in equilibrium expected payoffs is

$$(5.2) \qquad U_{t+1} - U_t = (v_{t+1} - v_t) \left(\sum_{j=0}^{t} P_j \right).$$

Now consider the sealed high bid auction. We saw in Section III that a type t bidder adopts a mixed strategy over some interval $[\bar{b}_{t-1}, \bar{b}_t]$. In particular, if he bids the upper support $\bar{b}_t$, a type t bidder wins with probability

$$\sum_{j=0}^{t} P_j$$

and so

$$(5.3) \qquad U_t = \left(\sum_{j=0}^{t} P_j \right) (v_t - \bar{b}_t).$$

But $\bar{b}_t$ is also the minimum bid by type $t+1$. The probability of winning is the same and so

$$(5.4) \qquad U_{t+1} = \left(\sum_{j=0}^{t} P_j \right) (v_{t+1} - \bar{b}_t).$$

Subtracting (5.3) from (5.4) we obtain (5.2). Finally we note that, in both

auctions $U_0 = 0$. Therefore expected payoffs are identical as claimed.

For the continuous version of the model, there is a helpful geometrical argument which also illustrates the result. Regardless of the details of the auction, the expected payoff to a buyer with valuation v can be written as

(5.5) $\qquad U = vP - R = v\{\text{probability of winning}\} - \{\text{expected payment}\}$

Moreover, given the rules of the auction, if all other buyers are bidding their equilibrium strategies, buyer 1 can compute the win probability P and expected payment R associated with each of his bids. This implicitly defines an equilibrium mapping $R(P)$. If buyer 1 has valuation v he then solves the following problem

(5.6) $\qquad \max_{P} \{U^1 = vP-R \mid R = R(P)\}.$

This maximization is depicted in Figure 3. Note that the indifference contours for buyer 1 are lines of slope v. Therefore buyer 1 chooses $P*$ so that

(5.7) $\qquad \dfrac{dR}{dP}\bigg|_{U} = v = R'(p)$

With valuations distributed continuously according to the c.d.f. $F(v)$, and $(n-1)$ other buyers all bidding according to the symmetric equilibrium, buyer 1's probability of winning is

(5.8) $\qquad P* = F^{n-1}(v).$

Differentiating (5.8) and making use of (5.7) we obtain

(5.9) $\qquad \dfrac{dR}{dv} = \dfrac{dR}{dP*}\dfrac{dP*}{dv} = v(n-1)F(v)^{n-2}F'(v).$

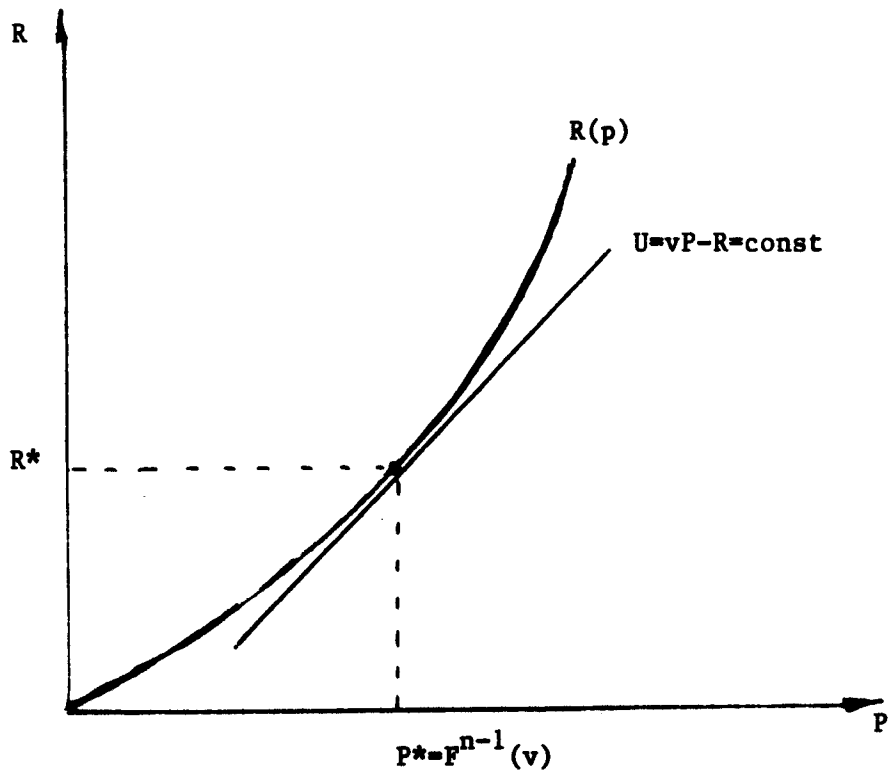As long as the minimum valuation is zero so that $R(0) = 0$, we can

Figure 3: Equilibrium best reply

reintegrate (5.9) to obtain

(5.10) $\qquad R(v) = \int_0^v x dF^{n-1}(x).$

That is, the expected payment by a buyer with valuation v is <u>independent</u> of the form of the auction.

Actually we have not simply proved Vickrey's original revenue equivalence result, but something far stronger. For <u>any</u> auction in which equilibrium bidding is strictly monotonic and the prize is awarded to the winner, the above argument goes through unchanged. Suppose, for example, that the high bidder wins the auction but all the bids are paid to the seller. Extending the analysis of Section I, it can be shown that higher valued buyers will indeed bid more in such an auction. Therefore expected revenue is the same as in the open ascending bid auction.

Using (5.10) we can immediately write down the equilibrium bidding strategy. Since R(v) is the equilibrium expected payment we have, with n bidders

$$b*(v) = \int_0^v x dF^{n-1}(x).$$

By way of contrast, in the sealed high bid auction the expected payment is the bid b(v) times the probability of winning $F^{n-1}(v)$. Hence

$$b(v) = \int_0^v \frac{x dF^{n-1}(x)}{F^{n-1}(v)}.$$

We now consider some of the implications of relaxing the assumptions. First of all we note that the private values assumption (C3) is not crucial. To see this, consider the two buyer case. Buyer 1, with signal s_1 wins if buyer 2 has a lower signal s_2. Buyer 1's expected value, conditional on

winning is therefore

$$(5.11) \qquad \bar{v}(s_1) = \frac{\int_0^{s_1} v(s_1,s_2)F'(s_2)ds_2}{F(s_1)}.$$

His equilibrium expected payoff

$$U^1(s_1) = \int_0^{s_1} v(s_1,s_2)F'(s_2) - b(s_1)F(s_1)$$

can therefore be rewritten as

$$(5.12) \qquad U^1(s_1) = \bar{v}(s_1)F(s_1) - b(s_1)F(s_1)$$

$$= \bar{v}(s_1) \text{ (probability of winning)} - \text{expected payment.}$$

Since (5.12) has exactly the form of equation (5.6) the analysis carries over. That is, revenue equivalence continues to hold.

Returning to the private values model, suppose Assumption C1 is replaced by the assumption that buyers are risk averse. That is

$$U^i = u(w-b,s_i), \quad \text{where} \quad u_{11} < 0.$$

For simplicity, suppose that the object for sale has a monetary valuation $v(s_i)$ so utility can be rewritten as

$$U^i = u(w+v(s_i)-b).$$

Consider the sealed high bid auction. If p is the probability of winning with a bid of b

$$EU^i = pu(w+v(s_i)-b) + (1-p)u(w).$$

Each buyer, in deciding whether to raise his bid considers the tradeoff between the greater win probability Δp and the increased cost Δb. From (5.12)

(5.13) $\qquad \Delta EU^i = \Delta p[u(w+v(s_i)-b) - u(w)] - pu'(w+v(s_i)-b)\Delta b.$

For a risk neutral buyer, with constant marginal utility of wealth, this can be rewritten as

(5.14) $\qquad \Delta EU^i = [\Delta p(v(s_i)-b) - p\Delta b]u'(w+v(s_i)-b).$

Buyer i then chooses a bid such that the marginal incentive to bid higher (the bracketed expression) is zero.

For a risk averse buyer, declining marginal utility implies that

$$u(w+v(s_i)-b) - u(w) > u'(w+v(s_i)-b)(v(s_i)-b)).$$

Comparing (5.13) and (5.14) it follows that a risk averse buyer has a greater incentive to bid higher. Intuitively, a risk averse buyer bids higher because his disutility of losing is greater.

In the open ascending bid auction, the introduction of risk aversion has no effect on equilibrium bids. Just as in the risk neutral case, each buyer has an incentive to stay in the auction until the asking price reaches his valuation $v(s_i)$. Risk aversion therefore leaves expected seller revenue unchanged in the open auction and increases expected revenue in the sealed high bid auction. It follows that expected seller revenue is higher in the latter.[3]

We next consider the implications of retaining assumption C1 (risk neutrality) but dropping the assumption of independence and private values. There are two central results both due to Milgrom and Weber (1982). First, in such an environment, public information _increases_ expected seller revenue. This public information may be information provided by the seller

---

[3]In fact, as Matthews (1983) and Maskin and Riley (1984) have shown, a risk neutral seller can further exploit buyers' fear of loss to raise expected revenue. Losing bidders must pay an entry "fee" $c(b)$, which is positive for low bids and negative for high bids.

prior to the auction or even information which arrives after the auction. In the latter case, the seller exploits the information by introducing a royalty payment, contingent upon the ex post signal.

The key insight which explains this result is that public information makes each buyer's own private signal less important. As a result differences in buyers' willingness to pay are reduced. It is precisely these differences which allow a buyer to make a positive expected profit and which therefore lower expected seller revenue. The point is most readily made when the public information is conclusive. Then all private signals are irrelevant and so bidders will compete away the entire surplus.

The second result concerns the comparison of the two auctions. Milgrom and Weber show that with risk neutrality, and with beliefs satisfying the assumption of conditional stochastic dominance, expected revenue is _lower_ in the sealed high bid auction.

To understand this result we return to the discrete 2 buyer example of Section III. Buyers valuations are either 0, v_1 or v_2 and beliefs given by the joint probability matrix

$$\begin{bmatrix} P_{00} & P_{01} & P_{02} \\ P_{10} & P_{11} & P_{12} \\ P_{20} & P_{21} & P_{22} \end{bmatrix}.$$

Crucial to the discussion is that, given conditional stochastic dominance

$$(5.15) \qquad \frac{P_{10}}{P_{10}+P_{11}} \geq \frac{P_{20}}{P_{20}+P_{21}}.$$

A simple example of a joint distribution satisfying conditional stochastic dominance and hence this property is depicted in Figure 4. Each buyer draws from the same urn with replacement. Each buyer knows the

distribution in the two urns but **neither** knows from which he is drawing. However observing a ball does reveal information. For if a buyer draws a zero he knows that the urn is more likely to be urn B and so the probability that his opponent will also have drawn a zero is relatively high. If it is urn B the joint probability of two zeros is $(0.6) \times (0.6) = 0.36$. If it is urn A the joint probability is $(0.2) \times (0.2) = 0.04$. Since the two urns are equally likely the joint probability of both drawing a zero,

$$P_{00} = \frac{1}{2}(0.04) + \frac{1}{2}(0.36) = 0.20$$

The rest of the probability matrix is similarly computed. We have

$$[P_{ij}] = \begin{bmatrix} .20 & .08 & .12 \\ .08 & .04 & .08 \\ .12 & .08 & .20 \end{bmatrix}$$

For the open ascending bid auction, a type 1 bidder makes a profit only against a type 0 bidder (who does not bid). His expected payoff is therefore

$$U_1 = \frac{P_{10}v_1}{P_{10}+P_{11}+P_{12}}.$$

A type 2 buyer also makes a profit against a type 1 bidder by bidding just more than v_1. His expected payoff is therefore

$$(5.16) \qquad U_2 = \frac{P_{20}v_2 + P_{21}(v_2-v_1)}{P_{20}+P_{21}+P_{22}} = \frac{(P_{20}+P_{21})(v_2-v_1)+P_{20}v_1}{P_{20}+P_{21}+P_{22}}$$

We now turn to the sealed high bid auction. As argued in Section 3, a type 1 bidder will adopt a mixed strategy over some interval $[0,\bar{b}_1]$ while a type 2 bidder will adopt a mixed strategy over some interval $[\bar{b}_1,\bar{b}_2]$. Since one of a type 1 buyers' equilibrium bids is zero, his equilibrium expected payoff is

(5.17) $$\bar{U}_1 = \frac{v_1 P_{10}}{P_{10} + P_{11} + P_{12}}$$

that is, exactly as in the open auction.

Moreover, if a type 1 buyer makes his maximum bid, $\bar{b}_1$, he bids against an opponent of type 0 or 1. Therefore,

(5.18) $$\bar{U}_1 = \frac{(v_1 - \bar{b}_1)(P_{10} + P_{11})}{P_{10} + P_{11} + P_{12}}.$$

Combining (5.17) and (5.18) we obtain

(5.19) $$v_1 - \bar{b}_1 = \left(\frac{P_{10}}{P_{10} + P_{11}}\right) v_1 \geq \left(\frac{P_{20}}{P_{20} + P_{21}}\right) v_1, \quad \text{by (5.15).}$$

Thus a type 1 buyer is less aggressive and makes a smaller maximum bid than he would if his beliefs were the same as those of a type 2 buyer.

Now consider a type 2 buyer. Since his minimum bid is $\bar{b}_1$ and, when he bids this, he wins only against type 0 and type 1 buyers, his expected payoff is

$$\bar{U}_2 = \frac{(v_2 - \bar{b}_1)(P_{20} + P_{21})}{P_{20} + P_{21} + P_{22}} = \frac{(v_2 - v_1)(P_{20} + P_{21}) + (v_1 - \bar{b}_1)(P_{20} + P_{21})}{P_{20} + P_{21} + P_{22}}.$$

Substituting from (5.19), we obtain

(5.20) $$\bar{U}_2 \geq \frac{(v_2 - v_1)(P_{20} + P_{21}) + v_1 P_{20}}{P_{20} + P_{21} + P_{22}}.$$

Moreover, this inequality is strict whenever inequality (5.15) is strict.

Comparing (5.16) and (5.20) it follows that the expected payoff of type 2 buyers is higher in the sealed high bid auction. But then expected revenue to the seller must be lower.

The key insight is that, in the sealed high bid auction, a type 2 buyer's minimum bid and hence his expected payoff, is determined by the maximum bid of type 1 buyers. This in turn is determined by the beliefs of type 1. As we have seen, the higher the probability a type 1 buyer assigns to the valuation v_0, the less aggressively he bids. A type 2 buyer is then able to benefit from the less aggressive bidding.

In contrast, when buyers bid in the open auction, the connection between the payoffs to type 3 and the beliefs of type 2 is severed. Expected revenue is then greater than in the sealed high bid auction.

Thus far we have considered only relaxing one or more of the class C assumptions. In particular the assumption of symmetric preferences, beliefs and bidding behavior has been maintained throughout. Given a symmetric environment, Maskin and Riley (1986) have shown that there is a strong presumption that equilibrium is symmetric in the sealed high bid auction. However, with open bidding and <u>common values</u>, there can be a continuum of equilibria. To see this, consider the two buyer case and suppose buyer 2 adopts the monotonic strategy

$$b_2 = B_2(s_2).$$

Let $\phi_2(b) = B_2^{-1}(b)$ be the inverse of this strategy. If the bidding reaches b and then buyer 2 drops out, buyer 1's net payoff is

$$v(s_1, \phi_2(b)) - b.$$

If this is positive, buyer 1 is happy to win the auction. However, if this is negative, buyer 1 has already waited too long. We conclude, therefore, that buyer 1 should drop out when the asking price reaches b_1^* satisfying

(5.21) $$v(s_1, \phi_2(b_1^*)) - b_1^* = 0.$$

This implicitly defines buyer 1's best reply $B_1(s_1)$. Writing the inverse of this function as $\phi_1(b) = B_1^{-1}(b)$ we can substitute into (5.21) to obtain the necessary condition

(5.22) $v(\phi_1(b),\phi_2(b)) - b = 0.$

Since exactly the same argument holds for buyer 2, condition (5.22) is the only necessary condition. Therefore any pair of strictly increasing inverse bid functions satisfying (5.22) are equilibrium inverse bid functions.

Finally, Maskin and Riley (1987) have compared the two common auctions when all the symmetry assumptions are dropped. To take the simplest possible case, suppose there are two buyers. The first's valuation is a draw from the distribution $F_1(v)$ and the second from the distribution $F_2(v)$. Loosely speaking, the main conclusion of this research is that the sealed high bid auction is likely to yield higher revenue when the upper support of one distribution is significantly larger than the upper support of the other. On the other hand, if the two distributions have the same support and F_1 and F_2 are sufficiently different, the ranking is likely to be reversed.

Two simple examples illustrate these results. Suppose first that buyer 1's valuation is continuously distributed according to the c.d.f. $F(v)$ on $[0,1]$ while buyer 2's valuation is some number L with certainty, where $L > 1$. Let $[\underline{b},\bar{b}]$ be the support of the two buyers' bid distributions. Buyer 1 will submit a bid if and only if his valuation, v_1, exceeds $\underline{b}$. Then if buyer 2 bids $\underline{b}$ he wins if $v_1 < \underline{b}$, that is, with probability $F(\underline{b})$. His equilibrium expected payoff is therefore $F(\underline{b})(L-b)$. If buyer 2 bids 1 he wins with probability 1. Therefore

$$F(\underline{b})(L-\underline{b}) \geq L-1.$$

Hence

$$F(\underline{b})L > L-1,$$

that is,

$$(5.23) \qquad F(\underline{b}) > \frac{L-1}{L}$$

As L becomes large the right hand side approaches 1. Therefore

$$\lim_{L \to \infty} F(\underline{b}) = \lim_{L \to \infty} \underline{b} = 1$$

It follows that, as L becomes large, the support of the equilibrium bid distribution collapses towards a bid of 1. Expected seller revenue therefore approaches 1.

For the open ascending bid auction, buyer 2 will always outbid buyer 1, who stays in the bidding until his valuation is reached. Expected seller revenue is therefore

$$E\{v_1\} = \int_0^1 vdF(v).$$

It follows that, for all L sufficiently large, the sealed high bid auction generates greater expected revenue.[4]

For the second example, suppose buyer 1's valuation is again equally likely to be 0 or 1, but buyer 2 has a valuation of 1 for sure. In the open auction expected revenue is then 1/2. Turning to the sealed high bid auction, arguing as in Section 3, equilibrium bidding strategies will be mixed strategies over some interval $[0,\bar{b}]$, for a buyer with a valuation of 1. But buyer 2 has an expected gain of 1/2 if he bids ϵ, since, with

---

[4]For the uniform case, with $F(v) = v$, it follows from (5.23) that, in the sealed high bid auction the minimum bid $\underline{b}$ exceeds 1/2 for all $L \geq 2$. Therefore expected revenue is higher in the sealed high bid auction for all $L \geq 2$.

probability 1/2 his opponent has a zero valuation. Therefore buyer 2 will never bid more than 1/2 and so the mixed strategies are over an interval $[0, \bar{b}]$ where $\bar{b} \leq 1/2$. It follows that expected revenue in the sealed high bid auction is less than 1/2.

VI. Concluding Remarks

While the theory of auctions and contests has advanced dramatically over the last few years, important questions remain unresolved. First of all, exploration of the asymmetric model has only just begun. Second, almost all of the theory on auctions has concentrated on the sale of a single unit. There is, as yet, almost no theory on sealed bid auctions in which each buyer has a downward sloping demand curve for multiple units.[5]

Third, the auction literature has focused almost exclusively on a single sale. In practice, major participants enter auctions repeatedly and so have an opportunity to develop bidding reputations. Moreover, in the sale of related items, the price of one item affects expectations about prices at later auctions. This not only affects the way buyers will bid but also has implications for the seller. Specifically, the seller must decide how high a minimum bid to set and whether to announce it. It would be very nice indeed to be able to draw firm theoretical predictions for repeated auctions of this type (vintage wine, old masters, etc.)

Finally, research is only just beginning to produce useful insights on the role of bidding in government contracting. Here an additional problem is that often the product cannot be completely specified at the time of the

----

[5]If all the items for sale are sold at the highest unsuccessful bid it is easy to see that each buyer has an incentive to submit his true demand curve. Modelling strategies when buyers must pay their schedule of bids is, however, a difficult challenge.

initial competition for the contract. Final terms are only determined after the contract winner completes further R&D. Knowing this, and anticipating later renegotiations, agents make initial bids that do not reflect anticipated costs.

One recent proposal is that there should be a further round of bidding for the right to produce after R&D is completed (some such contracts have recently been negotiated). Except under the very simplest of assumptions, the theoretical implications of this and other proposed incentive schemes remain to be worked out.

REFERENCES

Bishop, D.T. and Cannings, C., "A Generalized War of Attrition," <u>J. Theor. Biol</u>, 70 (1978), 85.

Crémer, Jacques and McLean, Richard P., "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist When Demands are Interdependent," <u>Econometrica</u>, 53 (Mar. 1985), 345-61.

Fudenberg, Drew and Tirole, Jean, "A Theory of Exit in Oligopoly," UC-Berkeley Working Paper (1983).

Graham, Daniel A. and Marshall, Robert C., "Collusive Behaviour at a Single Object English Auction," Working Paper 85-01, Duke University (Feb. 1985).

Hillman, Arye and Riley, John G., "Politically Contestable Rents and Transfers," mimeo, UCLA (May 1987).

Maskin, Eric and Riley, John, "Auctions with Asymmetric Beliefs," Working Paper No. 254, UCLA, (Apr. 1983).

_____, "Optimal Auctions with Risk Averse Buyers," <u>Econometrica</u>, 52(6) (Nov. 1984), 1473-518.

_____, "Existence and Uniqueness of Equilibrium in Sealed High Bid Auctions," mimeo, UCLA, (Mar. 1986).

Matthews, Steven A., "Selling to risk Averse Buyers with Unobservable Tastes," <u>Journal of Economic Theory</u>, 30(2) (Aug. 1983), 370-400.

Maynard Smith, John, "The Theory of Games and the Evolution of Animal Conflicts," <u>J. Theor. Biol.</u>, 47 (1974), 209.

McMillan, John and McAfee, R. Preston, "Auctions and Bidding," <u>Journal of Economic Literature</u>, 25 (June 1987), 691-754.

Milgrom, Paul R., "Auction Theory," in Truman Bewley (ed.), <u>Advances in Economic Theory 1985</u>, Cambridge: Cambridge University Press, 1986.

_____ and Weber, Robert J., "A Theory of Auctions and Competitive Bidding," <u>Econometrica</u>, 50(5) (Sept. 1982), 1089-122.

Myerson, Roger B., "Optimal Auction Design," <u>Math. Operations Res.</u>, 6(1) (Feb. 1981) 58-73.

Nalebuff, Barry and Riley, John G., "Asymmetric Equilibrium in the War of Attrition," <u>J. Theor. Biol.</u>, 113 (1985), 517-27.

Posner, Richard, "The Social Costs of Monopoly and Regulation," <u>Journal of Political Economy</u>, 83 (Aug. 1975), 807-27.

Riley, John G., "Ex Post Information in Auctions," Working Paper No. 367, UCLA (Mar. 1986).

_____ and Samuelson, William F., "Optimal Auctions," <u>American Economic Review</u>, 71(3) (June 1981), 381-92.

Tullock, Gordon, "Efficient Rent Seeking," in J.M. Buchanan, R. Tollison and G. Tullock (eds.), <u>Toward a Theory of the Rent Seeking Society</u>, Texas A&M Press, 1980, 269-282.

Vickrey, William, "Counterspeculation, Auctions, and Competitive Sealed Tenders," <u>Journal of Finance</u>, 16(1) (Mar. 1961), 8-37.