

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

AT YALE UNIVERSITY

Box 2125, Yale Station
New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 496

Note: Cowles Foundation Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. Requests for single copies of a Paper will be filled by the Cowles Foundation within the limits of the supply. References in publications to Discussion Papers (other than mere acknowledgment by a writer that he has access to such unpublished material) should be cleared with the author to protect the tentative character of these papers.

AUCTIONS AND BIDDING MODELS: A SURVEY

Richard Engelbrecht-Wiggans

September 11, 1978

AUCTIONS AND BIDDING MODELS; A SURVEY*

by

Richard Engelbrecht-Wiggans

Abstract

Auctions and bidding models are attracting an ever increasing amount of attention. The Stark and Rothkopf (1977) bibliography includes approximately 350 papers on the subject; additional work has been reported since the bibliography was compiled. This paper presents a general framework for classifying and describing various auctions and bidding models, and surveys the major results of the literature in terms of this framework.

Introduction

Auctions and bidding have long been used as methods for allocating and procuring goods and services. Although they were a subject seldom analyzed formally several decades ago, a substantial body of literature has developed more recently. The recent Stark and Rothkopf (1977) bibliography includes approximately 350 papers on the subject and the number is continually increasing, at what appears to be an accelerating rate.

This paper presents a unifying framework for classifying and describing auctions and bidding models and for discussing the related theory.

*This work relates to Department of the Navy Contract N00014-77-C09518 issued by the Office of Naval Research under Contract Authority NR 047-006. However, the content does not necessarily reflect the position or the policy of the Department of the Navy or the Government, and no official endorsement should be inferred.

The United States Government has at least a royalty-free, nonexclusive and irrevocable license throughout the world for Government purposes to publish, translate, reproduce, deliver, perform, dispose of, and to authorize others so to do, all or any portion of this work.

A variety of major results, concepts, and controversies are surveyed within this framework. Hopefully, such a survey will indicate not only what has been done, but also give a feel for what areas of possible interest have been neglected.

The objective of this paper is to survey the major ideas related to auctions and bidding. In a number of cases, specific results have been interpreted, reworded, or slightly extended to be consistent with the overall framework of the survey; responsibility for any incorrect or misleading comments and interpretations remains mine. Although an attempt has been made to cite a representative sample of relevant papers, the objective is to survey the knowledge on the subject of auctions and bidding models. No attempt has been made to mention every paper. For further and more complete references, the reader should refer to the Stark and Rothkopf bibliography and to the bibliographies of the cited works.

General Auction Model

Auctions may be viewed as a case of games with incomplete information as defined by Harsanyi (1967a, 1967b, 1967c). In such a model, there is an underlying true state of nature. The true state of nature prescribes the relevant characteristics and number of objects being auctioned, the von Neumann and Morgenstern (1953) utility functions and number of participating strategic players (typically some or all of the bidders), and the behavior of any non-strategic players (typically the auctioneer and perhaps some of the players).

While it is assumed that all players know precisely what true states of nature are possible and the probability distribution over these possible states, the players do not know precisely what true state of nature prevails

(or "has been chosen by Nature") in any particular situation. Each player, however, may receive some information about the true state of nature. A player receives information by observing the value (or outcome) of a, perhaps vector valued, random variable; the value is not revealed to any of the other players. An example of such information is any seismic data obtained by a firm participating in offshore oil lease sales. The probability distribution of the possible observations by each player depends on the true state of nature; the conditional distributions are known precisely by all the players.

Each player must specify a bidding strategy. A bidding strategy specifies how a player will use any information he might observe to determine his actual bid; bidding strategies are chosen before any information is observed, but determine a player's bid upon observing the information. Bidding strategies may be unrestricted and allow players to bid arbitrary functions of the observed information. Alternatively, the model may restrict players; for example by restricting each player's bidding strategies to specifying a single number (i.e., a multiplier) and having the player's bid be this number times some estimate of the true value of the object based whatever information is eventually observed. Other restrictions are possible and will be discussed later.

The game has a payoff function, again known precisely to each player. The payoff function determines, on the basis of players' bids, to whom each object is awarded and how much is paid by or to each player. Although charges to players for objects received or by players for services rendered are the usual cause for such payments, there may be payments associated with the cost of participating in the auction, the cost of preparing each bid, or in the case of fair division schemes, the division of the revenue generated by the auction.

There is an important distinction between the probability distribution of a random variable (e.g., the true state of nature or information to be observed by a player) and the value or outcome of the random variable. All the probability distributions are known precisely by all players, whereas the outcome of each random variable is known to only one or none of the players. Likewise, there is a distinction between bidding strategies and bids; the strategy is a function of the (usually) random valued information a player might observe and is used to determine the player's bid only after he has observed any information available to him.

Specific auction methods include sealed bid, progressive, and Dutch. The methods differ from one another in the degree to which the auctioneer has any active role and in the effect of sequential bidding. The model presented above does not capture the finer distinctions of these various alternatives; a variety of such very specific models and examples are discussed and compared by Greisner and Shubik (1963).

Most of the existing auction models, although often not specifically defined as games with imperfect information, are consistent with this game theoretic view. A scheme is presented below for classifying such auction models. Of the models which are clearly not game theoretic, for instance those which do not view the true state of nature as a random variable, most still fit comfortably into the classification scheme; in a few cases the classification scheme, however, might be awkward. The classification scheme, however imperfect, is presented in the hope that a succinct scheme will not only help identify the models, whether previously studied or not, most appropriate in a particular situation, but also that such a scheme will aid in unifying this survey of the literature.

Throughout this survey, auctions will be described according to four

major components, to wit: players, objects, payoff functions, and strategies. Some of these components in turn consist of several more specific sub-components. The various components and their possible descriptions are listed below.

1. Players:
 - a) Number participating: n , 1, 2, 3, ..., Random;
 - b) Utility functions: Linear, Non-linear, Other;
2. Objects:
 - a) Number: m , 1, 2, 3, ..., Random;
 - b) Information on object's value: Known, Symmetric, Identical, Other;
 - c) Physical characteristics: Identical, Symmetric, Other;
3. Payoff function:
 - a) Award Mechanism: Highest/Lowest bidder(s), Shared, Other;
 - b) Price: Incentive, Bonus, Profit share, Royalty, First rejected, Lowest accepted, Other;
 - c) Reservation price: None, Zero, Known (non-zero), Set (randomly) by auctioneer and non-strategic players, Other;
 - d) Other transfers: Auction participation costs, Bid preparation costs, Information costs, Redistribution of revenues (fair division schemes);
4. Strategies: Unrestricted, Multiplicative/Additive factor, Linear function, Additive across objects; Other.

In this classification, the terms "known," "identical," and "symmetric" have the following special meanings. The values of objects (for example)

are "known" if there is no uncertainty about the value (i.e., this component of the random variable describing the true state of nature is degenerate). The values would be "identical" if they are all equal to a single outcome of the random variable and "symmetric" if they are equal to the outcomes of independent identically distributed random variables.

A player, or strategic bidder, is anyone whose bidding strategy is unspecified by the model. Thus, traditional decision theoretic models correspond to games with one player; the one player being the lone strategic bidder. The behavior of any non-strategic bidders is incorporated into the true state of nature, one component of which is the reservation price of each object. Bidding models may specify a fixed number of players, most often two. Some models are in terms of n players, where n can be any integer strictly larger than one. Occasionally, the number of players is random. Although game theoretic usually have a known (fixed) number of players, the case of random numbers of players may be modelled by having only a random number of players receiving information which can result in competitive bids.

Players are often assumed to have linear utility functions; occasionally the assumption is explicit, but more often it is implicit in a statement such as "bidders are assumed to maximize expected profits." Non-linear utility functions are sometimes considered for single player models. Very occasionally considerations of risk aversion or capacity and budget constraints result in models with more than one player each with the same non-linear utility function. There are apparently no results for models with more general situations (unless there is no uncertainty in the physical characteristics of the objects).

A large number of bidding models involve auctions with only one object.

If there is more than one object, the objects are usually assumed to be identical. Models with more than one symmetric, or more generally valued, objects are rarely studied explicitly. Commonly, multi-object auctions are treated as if they were a number of independent single object auctions.

The physical characteristics (e.g., the number of barrels of oil under a given offshore tract) of an object may either be known to all players or uncertain. Note that even if an objects' characteristics are known to all players, different utility functions may result in players having different values for the objects. When the characteristics of the objects are not known, players' strategies will in general be non-trivial functions of any information they observe about the true state of nature.

In auctions where the players are bidding in order to obtain an object such as an offshore oil lease, the object is almost invariably awarded to a highest bidder. When players are bidding on a contract, the contract is usually awarded to a lowest bidder. Actually, when one uses the convention (as this survey will) that positive prices indicate money being paid by the player (to the auctioneer, or more generally, to Nature), then "high bid wins" and "low bid wins" auctions are actually the same. While there may be some conventions by which "high bid wins" auctions have been implicit reservation price equal to zero (whereas "low bid wins" auctions have no implicit reservation price) the classification scheme used in this survey specifies the reservation price explicitly. Since any implications on reservation price are thus of no consequence, all such auctions will be described as awarding the object to a "high" bid.

Occasionally, auction models are studied in which objects may be "shared"; players may be awarded fractional shares of objects. In such cases, and in models with more than one object, the awards are usually made

so to maximize the sum of the bids submitted by the players on the sets of objects they are awarded. The sum would be minimized in situations analogous to "low bid wins" auctions. Such award mechanisms are natural extensions of those awarding the object to an extreme bid in single object auctions.

Cases in which awards are made in part on considerations apart from the monetary bid may often be modelled as having players specifying multi-component bid functions. The non-monetary components of the "bid" may include product delivery dates, and quality or performance guarantees. Gilbert (undated) models bidding on cable television franchises as an auction with multi-component bids in which the players are uncertain how the components will be combined to determine the awards.

Occasionally there are "auctions" which might more appropriately be described as more general games. An example is the dollar bill auction in which the bill is given to the highest "reasonable" bidder; a bid is "reasonable" if less than zero or less than twice some other "reasonable" bid. This survey focuses on auctions with single component bids and awards being made to a high bidder.

The reservation price in an auction may be known to be zero, or it may be known that there is none (i.e., a reservation price of negative infinity). In decision theoretic models the reservation price is the lowest bid which will result in an object being awarded to a strategic bidder. Thus, the reservation price may be determined by the (random) "bids" of any non-strategic bidders.

In single item auctions, bidding strategies may be arbitrary functions of any information observed, or may be restricted to special forms. The most common restriction is that a player's bid is a multiple of some (usually

unbiased) estimate, based on any information observed, of the true value of the object to this player. As mentioned before, bids could possibly consist of several components; however we will not consider such cases. If more than one object is being sold, then the strategy may specify a bid for each possible subset of objects or be restricted to specifying a bid on each individual object and the bid on a set of objects being assumed to be the sum of the bids on the individual objects in that set. Other possible strategies include allowing players to submit a bid on each possible fractional share (any real number from zero to one) of an object.

The discussion to this point has been in terms of "one shot" auctions, any information obtained about the true state of nature or other players' bidding strategies is of no use in subsequent auctions. Such assumptions are at best an approximation to practical situations. The assumptions become especially suspect when certain parameters of a one shot auction are estimated using historical data on similar auctions.

While it is possible to model an entire sequence of games as a single big game, modelling a sequence of auctions as one big auction obscures much of the underlying structure and results in a model very difficult to analyze. A practical alternative used by Oren and Rothkopf (1975) is to model a sequence of auctions with a single player as a control problem. Agnew (1972) uses a similar approach and presents an algorithm for determining the lone strategic bidder's optimal markup over his known true value.

A slightly different approach is to try to determine the general structure of optimal (or at least good) strategies. Kortanek, Soden and Sodaro (1973) determine that in order to maximize the total of the awarded objects' contribution over direct cost in a variety of models, the general

form of the bid on an object should be the sum of the direct cost, the opportunity costs, and a competitive advantage fee; each of these terms may have slightly different definitions for different models. Attanasi (1974) and Attanasi and Johnson (1975b) use a similar approach but arrive at a slightly different interpretation of the form of optimal bidding strategies.

Sequences of auctions with more than one player are much more difficult to analyze; there are essentially no results in this area. Brams and Straffin (1977) examine the athletic drafting system and show that if a team knows precisely its own preference (ranking) of athletes and if teams are allowed to "bid" strategically (teams need not always choose the highest ranked of the remaining athletes), there are simple examples in which non-Pareto optimal allocations of athletes occur. Indeed, the resulting allocations are "far" from Pareto optimal in the sense that no sequence of bilateral trades (between one athlete from each of two teams) can result in a Pareto optimal allocation. Although, athletic drafts are not auctions in the traditional sense, the possible inefficiency in them indicates that it might be appropriate to examine more traditional sequential auctions for similar inefficiencies.

The order in which objects are auctioned affects the final allocation. Schotter (1974) considers a model of a horse auction. Each of a number of sellers has a horse to be sold and has a reservation price below which he will not sell. Each of a number of buyers wants exactly one horse and has a maximum price which he is willing to pay for a horse; as far as each buyer is concerned, all horses are of equal value. Each horse is sold in a progressive auction or, equivalently, sold to a high bidder at the second highest price, and it is assumed that bidding is "sincere"; buyers bid their true values. Depending on whether the horses are sold in decreasing or increasing

order of reservation price, the number of horses and the prices at which they are sold varies. Selling the horse with highest reservation price first results in the greatest number of horses sold; the richest buyers buy the most expensive horses, leaving the less expensive horses to the less wealthy buyers. The reverse order results in fewer horses being sold, but with a greater total profit to the sellers. Engelbrecht-Wiggans (1977) gives additional examples of sequential auctions where, if players are restrained to bid their true values, no order of auctioning the objects results in more than a small fraction of the profits and revenue of a Pareto optimal allocation. Although the above mentioned examples assume sincere bidding, the results suggest the order in which objects are auctioned plays strong role in the outcome of sequential auctions.

Players

The case of one player (the one strategic bidder) corresponds to a decision theoretic approach. The non-strategic, though often still random, bidding behavior of the remaining bidders is incorporated into the true state of nature via the reservation price. Thus, whether or not the player is awarded an object in response to a particular bid depends solely on the random reservation price.

Several models have been proposed for calculating the probability distribution of the reservation price. Friedman (1956), in his pioneering work on auctions, suggested that the probability distribution of the reservation price be determined by multiplying together the distributions of the various non-strategic bidders' bids. The non-strategic bidders' distributions might be obtained by considering historic data on related auctions. The implicit assumption of this approach is that the probability

of winning an object is equal to the probability of independently outbidding each of the non-strategic bidders.

Subsequent work in auctions has cast doubt on the appropriateness of such an independence assumption. If the player is uncertain about the true value of the object, a high bid may result from his observing information suggesting an overly optimistic true state of nature; a low bid may result conversely. Thus, chances are that if the player outbids a particular non-strategic bidder, the player has submitted "too high" a bid and will most likely also have outbid a number of the remaining non-strategic bidders. Conversely, a bid which is beaten by a particular non-strategic bidder tends to be a "low bid" which will likely be less than a number of the remaining non-strategic bids. Thus, the probability of outbidding a particular non-strategic bidder would not be independent of outbidding another non-strategic bidder. The independence assumption is inaccurate when the player is uncertain about the true characteristics of the object; alternative explanations, including limited collusion among non-strategic bidders, have been used to argue against the independence assumption when the player knows the characteristics of the object precisely.

The value of any model depends not so much on if it is absolutely correct or not, but rather on how good of an approximation it is. Thus, in some cases, the Friedman approach may be appropriate. An alternative approach, apparently without mathematical justification, but based purely on empirical goodness of fit is a formula of Gates (1967) for the probability distribution of the reservation price. A considerable controversy has ensued over the relative merits of the "Gates model" versus the "Friedman model." Rosenshine (1972) succinctly reviews and discusses the controversy to that point, but did not manage to end it. Apparently the source of the problem lies

in an attempt to prove one or the other model as THE correct approach; probably, neither is absolutely correct, but each may be an appropriate, simple, model in different situations. However, as observed by Clerckx and Naert (1972) the Gates and Friedman models may result in quite different optimal bidding strategies for the player, and thus the choice of Gates or Friedman ...or neither...model should be of some concern.

An alternative approach is to ignore the individual components giving rise to the distribution and simply use the distribution of resulting reservation price. LaValle (1976b) uses such an approach to analyze one player models. The approach would however be impractical if one must specify the actual distribution of the reservation price as part of the model. Fortunately, this difficulty is avoided if there is sufficient historical data; the data required may actually be less than that required by the Gates or Friedman approaches. As suggested by Hanssmann and Rivett (1959), the past data may be used to obtain an empirical distribution of the ratio of the players' estimate of the value of objects awarded in past auctions to the highest non-strategic bid in each corresponding auction. This distribution approximates the probability distribution of the reservation price.

When all the bidders are strategic bidders, and thus considered players, the reservation price (if any) is usually assumed to be fixed by the auctioneer and known to all the players. Situations in which the reservation price is not known can be handled, at least theoretically, by the game theoretic model. However, practically speaking, such an approach would require some means for estimating an appropriate distribution on the reservation price.

Auctions with a variable number of players have received considerably less attention than auctions with a fixed, known, number of bidders.

Models with uncertain numbers of bidders have the additional requirement that a probability distribution must be specified. Casey and Shafter (1964), Dean (1964) and Friedman (1956) have each suggested that the number of bidders is Poisson distributed. These suggestions, however, are often based more on theoretical rather than empirical arguments.

In auctions with a number of similar players, each independently of the others deciding whether or not to bid on a particular object, it might be assumed that each player has the same probability of bidding on an object. Under such conditions, the number of bids an object receives will be binomially distributed. As the number of players becomes large, and the probability any one of the players actually bids becomes small in such a way that the average number of bids on the object remains constant, then the binomial distribution approaches a Poisson.

When data can be observed from a number of auctions with similar objects, the observed distribution of the number of bids may be used to estimate the actual distribution. Keller and Bor (1978) study bidding data for a collection of approximately similar construction contracts in the United Kingdom. The observed data is consistent with the Poisson model; an alternate distribution for the data is a Gamma distribution.

The distribution of the number of bids in federal offshore oil lease sales appears not to be Poisson. Indeed, the distribution is occasionally strongly bimodal. Engelbrecht-Wiggans (1978b) proposes a simple model to show that assumptions similar to those for the common Poisson model may also hold here. When the objects being sold differ in value or other characteristics, there may be different distributions for the number of bids on the different objects. The observed distribution of bids would be a composite of these different distributions and need not be Poisson even

if the underlying distributions were.

In the model, there are three categories of objects. While players would like to bid only on objects in the most valuable category, they occasionally mistake a valuable object as one not worth bidding on and vice versa. The resulting distribution of number of bids is a mixture of binomial (or, under the appropriate conditions, approximately Poisson) distributions and fits the observed data much more closely than a simple Poisson distribution could. Thus, a player considering bidding on an object may assume that the number of competing bids will be Poisson distributed; however, the mean number of competing bids will depend on whether the object is indeed worth bidding on or not.

While the above suggests that players behave differently with regard to different objects in multi-object auctions such as federal offshore oil lease sales, work of Dougherty and Lohrenz (1977) suggests that bidders may also differ in how seriously they bid. In studying the distribution of the size of bids, Dougherty and Lohrenz (1976) developed the "30-30 deletion algorithm" to identify "non-serious" bids; bids either much less than the next higher bid or only a small fraction of the mean of the remaining bids. Of the 170 non-serious bids (1.9% of the total) identified by the algorithm, 89% were submitted by three firms; furthermore, 90% of these three firms' bids were identified as non-serious. The evidence suggests that there may be both serious and non-serious bidders; this observation is, however, only tangentially related to the main subject of the paper, and the statistical significance of the data (if any) is not explored.

Although most bidding models assume linear utility functions, the effects of non-linear utilities has received some attention in auctions with only one player where the player knows the physical characteristics

of the single object being auctioned. Hanson and Menezes (1968) prove the existence of an expected utility maximizing strategy under quite general continuity assumptions on the utility functions and the distribution function of the reservation price (or, equivalently, highest non-strategic bid). Uniqueness of the optimal strategy is proven under somewhat more restrictive conditions. Explicit bounds are established for the effect on the optimal strategy by changes in the true value; the change in bid is in the same direction, and of magnitude not exceeding, that of the change in the value of the object.

Hanson and Menezes discuss three different measures of risk aversion. One of these, the risk index of Arrow (1963, 1965) and Pratt (1964), is used by Baron (1972) to characterize and compare utility functions. An increase in the risk index (signifying increased risk aversion) results in an increase in the optimal bid in single object auctions where the lone player knows the true value of the object. The change in the optimal bid resulting from a change in true value of the object is greater for utility functions with constant risk aversion than for decreasing risk aversion, and even greater for increasing risk aversion.

Blaydon and Marshall (1974) note that the above results depend on the assumption that the player knows the value of the object. An example is provided to show that if the value is uncertain, the optimal bid may vary in either direction with changes in the risk index. Baron (1974) elaborates further on this point.

Attanasi and Johnson (1975b) consider models sequential auction models with non-linear utility functions. The models consider optimal bidding strategies in a sequence of markets, and thus assume there is only one strategic bidder. The effects of risk aversion in sequential auctions

are similar to those in one-shot auctions.

In auctions with more than one object, either the analysis must be in terms of expected dollar value of various subsets of items or else the utility function must have several components. If the concern is expected total value or if the utility functions are additive across components then multi-object auctions may be treated as independent simultaneous single object auctions. However, Engelbrecht (1977) and Scott (1975) have shown that additive utility functions are equivalent to multi-attribute risk neutrality.

Multi-object auctions have received almost no attention; the apparent implicit assumption being that it is appropriate to treat such auctions as a number of independent simultaneous auctions. Such an approach is inappropriate in at least some situations where bidders face capacity constraints, are subject to budget restrictions, or have risk averse utility functions. The few, very specialized, models will be described later.

Objects

By far the most commonly studied auction is that of a single object; if there is more than one object, the number is usually known and the objects are usually identical. Sometimes the true characteristics of an object are assumed to be known to all players and either different players have different utility functions (and thus, different true values for an object) or the players must decide how much and on which objects to bid subject to some capacity or budget constraint. More often, the true characteristics of an object are not known. Different players may observe different information and form different estimates of the object's true value; as noted by Brown (1975), such a situation gives rise to imperfect competition

among the bidders.

When the true characteristics of an object are not known, each player gains information about the true state by observing the value of a random variable whose distribution depends on the true state. A common assumption, especially in the literature on oil lease bidding, is that the information random variable is a random multiple of the true value of the object; sometimes the information random variable is a random error term added to the true value. The distribution of the multiplier (or additive error) is often assumed to be independent of the true state of nature; statistical work of Crawford (1970) and Dougherty and Lohrenz (1976) for bidding on oil leases may be responsible for the common assumption that in such auctions the information is a lognormal random variable (with a distribution independent of the true state of nature) times the actual true value. Winkler and Brooks (1977) study models in which the observed information consists of an error random variable added to the true value. Rarely considered are more complicated dependencies of the information random variable on the true state of nature; perhaps in part because in practical applications it may be difficult to estimate such general conditional distributions.

The dependence of the observed random variable on the true state of nature is not always clearly specified in the literature. Indeed, it appears that "observed value" and "unbiased estimate of the true value" are often (implicitly) assumed to mean the same thing. However, for additive or multiplicative errors, these two numbers are in general equal only if the distribution of the true state of nature (or, more precisely, the marginal distribution of the true value of the object) has mean one or zero respectively.

Payoff Functions

The payoff function of a game determines who gets what on the basis of the strategies chosen by the players and the true state of nature. In auctions, the payoff function determines the prices of the objects and to whom each object is awarded. Occasionally, the payoff function also specifies fees for preparing or submitting bids, or a cost for participating in the auction.

Almost invariably, single object auctions with single component monetary bids award the object (if awarded at all) to a high (or low) bidder; if all the bids are less (respectively, greater) than the reservation price, the object is not awarded to any player. Perhaps the only seriously studied exception to this rule is the share auction in which fractional shares of the object may be awarded; such an auction however may be viewed as a limiting case of the multi-object auctions discussed next.

When more than one object is being auctioned, a common extension of the high bid wins rule is to determine a partition of the objects into subsets, one for each player, which maximizes the sum of the amounts bid by each player on the subset of objects actually awarded to him. In cases where different players have different utility functions, know their respective true values for each object precisely and bids are equal to true values, this award mechanism assures a Pareto optimal allocation of the objects. When bidding strategies are restricted to being additive across objects, then this award mechanism is identical to awarding each object to a high bidder on that object.

In addition to determining to whom each item is awarded, the payoff function sets prices on the objects. While the price paid for an object is often set equal to the amount bid by the player to whom the object (or

set of objects) is awarded, many variations are possible and actually used. There may, of course, also be charges for things other than objects; for example, bid preparation costs.

The price may be a function solely of the amount bid by the player to whom the object is awarded and on the true state of nature. Examples of such pricing mechanisms include incentive, bonus, royalty, and profit sharing. Additional variations are possible.

Under incentive pricing, the price is the true value of the object (perhaps including a standard profit) less a fraction of the amount by which this true value exceeds the amount bid. In the extreme case that the fraction is zero, the mechanism becomes "cost plus fixed fee." The other extreme case, the fraction equal to one, results in a price equal to the amount bid.

One goal of the intermediate forms of incentive pricing is to encourage efficiency. If bids are unbiased estimates of the true value of the contract (recall, that by our convention, the value of a contract is the negative of its cost) then the contracted firm is paid a positive "incentive" whenever the cost is less than the estimate (the negative of the bid). Fisher (1969) observed that such apparently desirable "under-runs" tend to be larger for incentive contracts (with an intermediate incentive rate) than for cost plus fixed fee contracts, but suggests that this may be due to players strategically over-estimating the cost rather than increased efficiency. However, neither Fisher nor Deavers and McCall (1966) found any conclusive relationship between the incentive rate and the under-runs.

Variations on an extreme case of incentive pricing include the bonus pricing currently used in most offshore oil lease auctions. Under bonus bidding, the price is equal to the bid amount plus a fraction of the value of any oil recovered. Under such a scheme, the actual price of an object (i.e.,

lease) will not be known immediately upon conclusion of the auction; the price is not known until the site has been developed and its true value becomes known. Other variations on incentive pricing are possible. Many of which, including the above, result in a price which is a linear combination of the amount bid, the true value, and any (positive) profit.

A slightly different form of variation is when players bid on the incentive rate or the fraction of profits to be included in the price. Scherer (1964) studies setting the incentive rate through negotiations independent of setting the target cost. The United States Government experimented with selling a small number of oil leases under royalty bidding; under which the players' "bids" are what fraction of the recovered oil they will pay in exchange for the lease.

The amount of revenue generated under various pricing schemes has been compared. Reese (1978a) concludes that for oil lease sales, on the average profit share pricing should result in more revenue than royalty pricing, and royalty pricing in more revenue than the currently used bonus pricing. Different price mechanisms may also have effects on how leases will be developed; the expected revenue should not be the only criterion for comparing different price mechanisms. Attanasi and Johnson (1975a) and Kalter, Tyner and Hughes (1975) have also compared various pricing mechanisms for oil lease auctions; all authors reach basically similar conclusions. Wilson (1977b) concludes that, on the average, a share auction would result in even less total revenue than an auction with bonus bidding.

The Office of OCS Program Coordination (undated) analysis of the auction in which oil leases were sold under royalty bidding concludes that there tend to be more bids on sites sold under royalty bidding than under bonus bidding; under royalty bidding a site is also less likely to receive

no bids. However, the relative merits of the two schemes remain inconclusive due to a variety of reasons. The amount of revenue eventually generated may be more relevant than the number of bids submitted. The size and number of bids may have been affected by the fact that this auction was an experiment and that bidders found themselves in an unusual auction. Finally, the number of sites sold under royalty bidding was very small.

In a number of "second price" auctions, the price of an object depends on bids other than just the highest bid. The common progressive auction in which bidders raise the price of an object until no one desires to raise it further and the object is then awarded to the last bidder to raise the price at the price to which he raised it has been modelled, by Vickrey (1961a), as an auction in which the object is awarded to a highest bidder at a price equal to the second highest bid (or at some small increment above the second highest bid). If several identical objects are being auctioned, and players may bid on how much they would pay for the first, second, etc., objects awarded to them, then an appropriate extension of the model is to set the price of all objects equal to the highest unsuccessful (rejected) bid. As will be discussed later, second price auctions tend to result in bidding strategies which are simpler to calculate (or at least estimate) than for other auctions.

Friedman (1959, 1963) proposed that treasury bills should be auctioned according to the second price mechanism. With bidding strategies easier to determine, more bidders would presumably compete in the market. Vickrey (1961a, 1961b) has shown that it may be incorrect to expect that second price auctions would result in less total revenue than auctions in which the price is set equal to the amount bid. If different players have different true values for each of a number of identical objects and each player knows

his value, then setting prices either at the amount bid or setting prices uniformly at the highest rejected bid (or even, uniformly at the lowest accepted bid) will result in the same expected revenue even though the different price mechanisms give rise to different sets of equilibrium strategies. If a player must pay the price he bids, he will on the average hedge more than if he will typically pay something less than his bid.

In treasury bill auctions, however, different players typically have close to the same true value for the bills; the difficulty is that the true value is unknown. Goldstein (1962) argues qualitatively that second price auctions may not result in as much revenue. Smith (1966) gives examples of distributions of the reservation prices such that a profit maximizing player in a second price auction will expect to pay more than when the price equals the price bid; first and second price auctions would give rise to different distributions of the reservation price, but is not clear that there is any case of a game with incomplete information giving rise simultaneously to the two distributions used in this example. Since in some special cases, second price auctions result in a higher revenue while in others such auctions result in reduced revenue, the appropriateness of second price auctions for treasury bills must in part depend on the details of the treasury market.

In addition to payments for objects received or services rendered and for cost of bidding or participating in the auction, another form of payments arise when fair division problems are viewed as slightly generalized auctions. Typical fair division schemes involve auctioning the estate according to some multi-commodity auction and then dividing the resulting revenue among the heirs. Such auctions are outside of the scope of this survey; Butler (1970) surveys a large number of traditional fair division schemes, while Dubins (1977) and Engelbrecht-Wiggans (1977) study several

schemes allowing players to express preferences as to how objects are allocated among the other players and permitting bids to be non-additive across objects.

Bidding Strategies

Players may be assumed to select their bidding strategy according to any one of a number of criteria. In "min-max" models, each player chooses a strategy which maximizes the minimum possible utility of the final outcome over all possible combinations of bidding strategies of the remaining players. Occasionally, especially in simulations of auctions with small stakes, the players will try to maximize the amount by which their profits exceed those of the remaining players. In most models, however, Nash (1950) equilibrium strategies are sought; strategies are in equilibrium if each player uses a strategy which, for the particular strategies used by the remaining players, maximizes the expected utility of the outcome. One slight variant of such equilibria, are "local" equilibria in which each player's strategy results in attaining a local maximum of his expected utility; such a concept is useful to eliminate the spurious equilibrium which requires all players to bid arbitrarily low whenever there is a positive probability of there being no other bids submitted. A second variant on equilibrium strategies is to consider sets of stable strategies; strategies are stable if they come "close" to satisfying the equilibrium conditions.

In order to determine optimal min-max bidding strategies, a player need not consider what strategies the remaining players will actually use. This considerably simplifies the calculation of optimal strategies and allows solution of some otherwise extremely difficult models. Engelbrecht-Wiggans (1977) considers multi-commodity auctions where each player knows his value

for each possible subset precisely, each player submits a bid on each possible subset of the objects and the price of a subset of objects is equal to the amount bid by its winner. In such auctions, a min-max strategy is for each player to bid his true value on each subset.

The practical value of min-max solutions is limited by the fact that it is independent of the probability that any of the low utility outcomes actually occur. In some cases, all the low utility outcomes may (collectively) occur with very low probability whereas the high utility outcomes may occur with high probability. In other cases the probabilities may be reversed. A min-max analysis compares only the worst possible cases and ignores the possibility that a low probability very bad worst case may be balanced by a high probability extremely good outcome.

More often, bidding strategies are analyzed in terms of expected utility of the outcome; or, if utility is linear in money, in terms of expected profits. For models with a single player, this results in a traditional Bayesian decision analysis. Models with more than one player invariably require players' strategies to satisfy one of equilibrium concept variants mentioned above.

Players may be unrestricted as to their choice of bidding strategies, or restricted to only using strategies from some prespecified class. In almost every bidding model studied in the literature, the bidding strategies are real-valued functions; very seldom are bids with more than one component considered. The single component is usually interpreted as a monetary amount, but may occasionally represent a royalty or profit sharing rate.

When all players are concerned expected profits (i.e., all players have single attribute, linear, utility functions) and receive symmetric information, then the symmetric pure equilibrium strategy (if one exists)

is the solution to a linear first order differential equation. Thus it is possible to write an explicit symbolic expression for the symmetric equilibrium strategy. However, since each of the constants of the differential equation is an integral (which, for typical probability distributions, can not be evaluated in closed form), the expression for the equilibrium strategies involves several sets of multiple integrals which would require quite sophisticated numerical analysis to evaluate; there is little hope for obtaining a closed form explicit expression for the equilibrium strategies.

The functional relation between information and bids may also be restricted. A common assumption, is that bid functions are "multiplicative." Under multiplicative bidding, a player's choice of strategies is limited to specifying a multiplier (before observing any information); his bid is this multiplier times an estimate of the object's true value based on whatever information is observed.

Perhaps the simplest multiplicative (and additive) strategy is for each player to use any information observed to calculate an estimate, typically unbiased, of the true value of an object and let his bid be equal to that estimate. Work of Beckmann (1974), LaValle (1967a) and Vickery (1961a) shows that all players bidding their expected values for an object in a single object auction is in equilibrium if the object is priced at the highest rejected bid and if each player's expected value of the object is independent of any information observed by the remaining players. An example of such a situation is when different players have symmetric true values for the object and each player knows his own true value precisely. Vickery (1961b) gives a similar result for auctions with several (e.g. M) identical objects when one object is awarded to each of the M highest

bidders and each object is uniformly priced equal to the highest rejected bid (i.e., the $M+1^{\text{st}}$ bid).

For general single object auctions (or auctions with several identical objects where at most one object is awarded to any player) with prices of all objects uniformly equal to the highest rejected bid and with all players maximizing expected profits and receiving symmetric information, it is relatively easy to calculate the symmetric equilibrium strategy. The equilibrium strategy is the ratio of two single integrals involving the distributions of the true state of nature and the information observed by a player. Even though it is often impossible to evaluate the integrals in closed form, numerical approximations of single integrals are relatively accurate and easy to calculate. In general, the equilibrium strategy is not simply the expected value of the object based on any information observed by a player.

In auctions with the price of an object equal to the amount bid by the winning player, each player simply bidding his expected value of an object is not in equilibrium. The disequilibrium of any particular strategy may be verified by showing that it does not satisfy the desired differential equations. Such a verification is particularly simple in the case of players maximizing expected profit and receiving symmetric information.

One intuitive explanation why bidding expected values is not in equilibrium is a phenomenon known as the "bidder's curse"; the individual to whom an object is awarded tends to be the one who most overestimated the true value of the object. This phenomenon was originally analyzed in single player auctions by Capen, Clapp and Campbell (1971). Oren and Williams (1975) study the bidder's curse in general auctions with more than one player where the price of an object is equal to the amount bid.

If each player bids an unbiased estimate of his own true value for an object, then the maximum bid will in general be biased (upward) with respect to each of the player's true values; the maximum of an unbiased estimate and any second random variable will be biased upwards unless the second random variable is never greater than the unbiased estimate (in which case, the maximum would always be equal to the unbiased estimate, and thus also unbiased). Thus, regardless of any asymmetries or information dependencies, if a player bids an unbiased estimate of his true value then he will in general be awarded the object only at a price which averages higher than his true value.

If players have identical utilities but receive information of varying degrees of accuracy, then the bidders' curse will have the strongest affect on those players with the worst information. A player with good information will never overbid by very much, and will usually win an object because the less well informed players underbid the object. A player with poor information, on the other hand, will tend to win only when he overestimates the value; the existence of a player with good information limits the chances of obtaining the object at a bargain price.

Wilson (1967) shows that if there are only two players, one with perfect information and the other with less than perfect information about the true value of an object, then the poorly informed player may still obtain the object at a bargain price. With only one well informed player, the perfectly informed player's expected profit maximizing strategy will take advantage of the other players' poor information and try to obtain the object at a bargain price by bidding relatively low. However, it seems likely that if there are more than one well informed players, these players must now bid against each other. It is then no longer possible to obtain

a bargain by merely outbidding the poorly informed bidders. Dougherty and Nozaki (1975) verify that the aggressiveness which one bids should increase as the accuracy of information available to competing bidders improves.

A poorly informed player competing with several well informed players will rarely obtain a bargain and will usually win an object only when he overbid. Indeed, Wilson (1977a) and Milgram (1977) show that under a variety of assumptions, the price of the object under equilibrium bidding in auctions with the price equal to a highest bid tends (in probability) to the true value of the object as the number of symmetric best informed players tends to infinity. Thus, a poorly informed player competing with a large number of well informed players will win an object only by overbidding the true value.

Hughart (1975) presents a simple model in which the poorly informed players cannot expect any positive profit even though they are competing with only one well informed player; the poorly informed players must either obtain better information or have no incentive to participate in the auction. If the poorly informed players choose not to participate, there is no competition for the well informed player and the object may be sold at a very low price. If, however, the poorly informed players choose to participate and desire to profit by doing so, they must obtain additional information, presumably at some cost, and possibly duplicating some of the information already observed by the well informed player.

Hughart claims that either alternative is socially undesirable for offshore oil lease auctions. Either using a second price auction or having the government obtain accurate information about the lease sites and provide this information (without cost) to all potential players is advocated. In analyzing various alternative auction mechanisms for offshore oil lease

auctions, Reese (1978b) concludes that if the government is as efficient at obtaining information as private oil companies, then it should obtain such information and distribute it to all potential bidders. In doing so, the government can expect not only an increase in revenue from the auction in excess of the cost of the information distributed, but also that the revenue revenue increases by more than the players' profits decrease. This suggests that such dispersement of information will result in an allocation of oil leases closer to a Pareto optimal allocation.

An alternative to simply bidding the expected value of an object, is to bid some fraction of it. If the fraction is sufficiently less than one, then the actual price need not average higher than the true value. Such multiplicative strategies are often considered in the literature on offshore oil lease auctions.

Engelbrecht-Wiggans (1978a) shows that restricting strategies to being multiplicative results in several mathematical simplifications when all players are concerned with expected profits (i.e., all players have linear single-attribute utility functions) and all players receive symmetric information. First, the symmetric equilibrium strategy (often referred to as the "optimal bid fraction") is equal to the ratio of single integrals; this expression is much simpler than that for equilibrium unrestricted strategies, and is very easy to evaluate numerically whenever closed form expressions for the integrals fail to exist. Second, and perhaps most important in light of the vagueness with which most models are described, is that if the information is a random multiple of the true value and the probability distribution of the multiple is independent of the true state, then the optimal bid fraction is independent of the probability distribution of the true state.

The optimal bid fraction depends on the number of players and the accuracy of the information they receive. The bid fraction has a maximum for a finite number, typically less than a dozen, players. For many players, there is likely to be at least one player who grossly overestimated the true value of the object; thus in order to avoid the bidders' curse, one bids more conservatively against very large numbers of players. If there are only a very few players, then there is a chance of obtaining the object at a bargain price; thus one should also bid more conservatively against very small numbers of players. The more accurate the information observed by the players, the less of an effect the bidders' curse has. Thus, as the variance of the ratio of the information to the true value decreases, the optimal bid fraction increases. Capen, Clapp and Campbell (1971) obtain similar results for the optimal bid fraction in auctions with only one player.

Equilibrium multiplicative strategies are not necessarily in equilibrium when strategies are unrestricted. Rothkopf (1969, 1971) proves that if the information is a random multiple of the true value and the posterior distribution (after observing any information) of the ratio between the observed information and the true value is independent of the observed information, then equilibrium multiplicative strategies are also in equilibrium when the strategies are not restricted. Rothkopf (1977b) also observes that these conditions are in general only satisfied if the true value has a diffuse uniform distribution prior to observing the information. Winkler and Brooks (1977) prove a corresponding result for models with additive information; if the true value has a diffuse uniform distribution prior to observing any information and the players receive symmetric information consisting of a random error added to the true value (where the distribution

of the error is independent of the true value) then the symmetric equilibrium strategy is additive, or more precisely, consists of adding a constant "hedge" to the observed information.

If the true state of nature does not have a diffuse uniform distribution, then equilibrium multiplicative bids are in general not equilibrium unrestricted bids. An intuitive understanding of the cause may be gained by considering the expected value of the object after having observed any information. If the posterior distribution of the random multiple is independent of the observed information (or, equivalently, the true value has a prior diffuse uniform distribution) then this expected posterior value (or Bayes estimate) is an unbiased estimate of the true value. However, whenever the prior distribution is not diffuse, the Bayes estimate will not in general be unbiased.

A non-diffuse prior distribution on the true state establishes (in a crude, non-quantitative way) an idea of what size numbers are typical for the observed information. Although the observed information is an unbiased estimate if the random multiplier has mean one, a "large" number for the observed information is likely to arise from the true value having been multiplied by a "large" value of the multiplier rather than from the true value actually being "large." Conversely, if the observed information is a "small" number, one suspects that this is because one just happened to observe information equal to a "small" multiple of the true value. Thus, the Bayes estimate of the true value is smaller than an unbiased estimate if the observed information is a "large" number, and conversely for "small" numbers.

The appealing simplicity of multiplicative bids suggests determining how close equilibrium multiplicative bids are to equilibrium unrestricted

bids. Rothkopf (1977c) calculates equilibrium linear strategies; linear strategies require bids to be a prespecified constant added to a prespecified multiple of the observed information. As the variance of the prior distribution becomes large compared to the variance of the random multiplier, the additive term of the equilibrium linear strategy becomes small compared with the multiplicative term; the strategy goes toward being multiplicative.

The convergence of equilibrium linear strategies to multiplicative strategies suggests that perhaps equilibrium unrestricted strategies might also converge to multiplicative strategies. Indeed, since multiplicative strategies are in equilibrium for diffuse uniform prior distributions on the true value, such convergence seems highly likely. There, however, appear to be two possible weaknesses in this approach; first, a strategy being predominantly multiplicative does not imply that any multiplicative strategy is close to being in equilibrium, and second, there is little indication of how close equilibrium unrestricted strategies are to being multiplicative short of the limit.

Engelbrecht-Wiggans (1978a) studies the disequilibrium of multiplicative strategies by numerically analyzing asymmetric example of oil lease bidding; the distributions and parameters are chosen to be consistent with bidding data on OCS sale #40 and with assumptions commonly stated in the literature. It is shown that equilibrium multiplicative strategies are quite far from being in equilibrium. Indeed, when all players use the equilibrium multiplicative strategy, any individual player can increase his expected profits by approximately one third (which translates into about \$500,000 per lease in this example) if no longer restricted to bid multiplicatively. If strategies are restricted to be multiples of a player's Bayes estimate, then the equilibrium restricted strategies are much closer

to being in equilibrium for the unrestricted case; the corresponding increase in expected profits is less than ten percent.

The example does suggest that equilibrium multiplicative strategies may rapidly converge to equilibrium unrestricted strategies. If, in each of a sequence of auctions, each player performs a decision theoretic analysis and uses the bidding strategy which would have been optimal in the previous auction, then after the third iteration, the optimal strategy has an expected value within one percent of that actually used; the strategies converge in one less iteration if the players are initially restricted to multiples of Bayes estimates. Thus, under moderately stable market conditions, one might expect that each player performing a Bayesian analysis after each auction would result in bidding strategies close to equilibrium after a small number of "learning experiences." It should be noted however, that in this example, the expected profit under the equilibrium unrestricted strategies is only about three fourths that under the equilibrium multiplicative strategies; thus all the players would profit if they all believed that multiplicative strategies were in equilibrium and therefore voluntarily restricted themselves to using only multiplicative strategies!

Smith and Case (1975) consider sets of strategies which are almost in equilibrium. In particular applications (e.g., in repeated auctions with converging strategies as in the above example) any set of strategies close enough to equilibrium will tend to be stable; no player will find it beneficial to deviate from any strategy which is sufficiently close to giving the maximum possible expected utility. Since stable strategies have similar self policing characteristics to equilibrium strategies, they can be used as alternative solutions to bidding games. Smith and Case examine a collection of stable sets of strategies and give an example where there

is a stable set which results in substantially more profit to all players than the equilibrium strategies. They suggest that the players might strive (e.g., through signalling in repeated auctions) to converge to such a set of strategies. Once the players are using such strategies, there appears to be little incentive (especially for any far sighted player) to deviate from them in future auctions; thus players may repeatedly use the same strategies even though they are not strictly in equilibrium.

An alternative source of apparent non-equilibrium behavior is if it is assumed that bidders do not share information yet there is actually some collusion (and the strategies are in equilibrium when taking into account the collusion). Schilling and Gallo (undated) have developed a collection of computer programs which attempts to discover collusion among bidders when none is allowed. Occasionally, as with jointly prepared offshore oil lease bids, some collusion among bidders is explicitly allowed. Dougherty and Lohrenz (1977) study the effect of permitting joint bids in offshore oil leases and conclude that it does not decrease the number of bids submitted. Since joint bidders presumably have at least an accurate information as solo bidders, allowing joint preparation of some of the bids should result in more competitive bidding and thus a higher expected revenue to the government.

When more than one object is being auctioned, it is possible to restrict the relationship between bids on various subsets of objects. In particular, bids may be restricted to being additive across objects; the bid on a subset of objects is equal to the sum of the bids on the individual objects in that subset. Occasionally, the players' utility functions are such that bids will be additive even if not restricted to be so. However, Raiffa (1970) shows that if the objects are statistically independent,

monetary valued lotteries then players' true value functions are additive in general if and only if their utility for money is either linear or exponential. This suggests that bids are not likely to be additive unless so restricted.

Although at least a hundred oil lease sites are typically auctioned simultaneously, most analyses assume bids to be additive across objects and have been in terms of multiple simultaneous independent single object auctions. While treating multi-commodity auctions as a number of simultaneous independent single commodity auctions simplifies the analysis, such an approach may ignore the effects of budget or capacity limitations, constraints on exposure, or risk aversion. These aspects of multi-commodity auctions have received relatively little attention.

Goodman and Baurmeister (1976) and Stark and Mayer (1971) consider multi-commodity auctions with one player. The single player faces a decision theoretic problem of how much to bid and in which auctions to participate. Stark and Mayer model the optimization problem as a linear program. Goodman and Baurmeister develop an optimization algorithm which they claim is computationally efficient for up to six or seven simultaneous auctions. In a related vein; Rothkopf (1977a) derives a linear programming model to decide how much and in which auctions a player should bid if he faces a constraint on exposure; there is limit on the total value of the bids which he may have outstanding at any time.

The problem appears to become substantially more difficult if there is more than one strategic player. Engelbrecht-Wiggans (1978a) examines a very simple example of a number of non-independent simultaneous single commodity auctions. The objects are all identical and each auction uses the first rejected price mechanism; each player values a single object at

some fixed (known) amount, considers additional objects worthless, and is allowed to independently submit bids in up to two randomly selected auctions. In this example, the mixed equilibrium strategy (no pure strategy exists) is for each player to submit one "high" bid and one "low" bid. Although all the objects are identical and all players have the same known value for an object, the players' capacity constraint on the number of objects they can use results in two distinct levels of bidding. Although this is only a very simple example, it suggests that some of the variation among bids in auctions with pure strategy equilibria is due to capacity constraints rather than just different estimates of an object's true value.

Cook, Kirby and Mehndiratta (1975) consider a class of two player multi-object auctions where the players have limited resources. A player may not bid so as to win, on the average, contracts requiring more than his available resources. The model requires that players submit bids equal to one of a finite set of numbers; although such discretization may be appropriate in some situations, it has grave consequences in this analysis.

The results are considerably simplified and illuminated by reinterpreting them for the case in which there is only one constrained commodity. The equilibrium strategies are for both players to bid as high as possible (more precisely, the maximum allowed bid level not exceeding the true value) whenever such bidding is consistent with the capacity constraints. If a player (by the assumptions of the model, there will be at most one) would violate his capacity constraint by uniformly bidding the maximum allowed, this player should make his strategy feasible by reducing an appropriate number of bids to the next lower permitted level; under these equilibrium strategies, a player will never be awarded an object on which he bid at the second level. Although there does not appear to be an equally transparent

explanation of the results for more than one constrained resource, the equilibrium strategies have a similar "bid as high as you can on as many contracts as you can without violating your resource constraints, and bid somewhat lower on everything else" flavor. Again, the existence of equilibrium strategies in general depends on the fact that possible bid levels are discrete; such a model has limited practical value. (The published work also contains a number of serious errors, including an incorrect theorem with a proof which correctly proves a result different from that stated.)

Finally, there may be payments to and from players in addition to payments for objects. Typical examples include the costs of preparing bids, obtaining information, and participating in an auction. Auction models typically do not include such costs; such costs may be negligibly small in some situations.

Occasionally such costs have been considered explicitly. Most of the results however are for two player constant sum games, commonly referred to as "Colonel Blotto" games. Typically, the major concern is how players should allocate their limited resources among the various objects. Such games are however only a very specialized case of multi-commodity auctions; the interested reader is referred to the surveys and discussions of Beal and Heselden (1962) and Shubik and Weber (1978).

Further Research

This survey reveals at least a few areas which might warrant further research. Included among these are the effects of auctioning more than one object simultaneously, the equilibrium nature of bidding strategies actually used, and the effects of asymmetries in players' information or utility functions. Each of these areas is discussed briefly below.

The simple second price auction example surveyed indicates that multiple independent simultaneous auctions may result in allocations which are far from Pareto optimal. The question remains of how inefficient independent simultaneous auctions are in more typical situations, e.g., situations with less severe capacity constraints. Are there any alternative auction mechanisms which result in allocations sufficiently closer to Pareto optimal allocations so as to justify any additional costs from using such schemes?

Some of the difficulties associated with capacity constraints in simultaneous independent auctions are alleviated if the objects are auctioned sequentially. However, the order in which objects are auctioned may affect the final allocation. Most existing sequential auction models assume that players know precisely the true value of each object and that players will use multiplicative bidding strategies. In what situations, if any, is a sequential auction to be preferred over independent simultaneous auctions?

When several similar objects are being sold, a player may rationally submit different bids on the objects' even if he estimates them all to have the same value. Thus, some of the variance among bids on an object arises from sources other than players' uncertainty about the objects true value. In practical situations, how much of the variance in bids is due to such strategic considerations; how should one use data on the distribution of

bids observed on an object to determine what uncertainty players face?

Multiplicative strategies are not necessarily equilibrium strategies. Indeed, rather restrictive conditions must be satisfied before multiplicative strategies are precisely in equilibrium. Under what conditions are multiplicative strategies in equilibrium, and how sensitive is this equilibrium to slight changes in the model parameters? Even if multiplicative strategies are not exactly in equilibrium, are they close enough to be considered stable?

In the example surveyed, multiplicative strategies can be thought of as an initial strategy which will be modified and converge to a stable strategy under repeated play. In this example, strategies based on Bayes estimates resulted in faster convergence than multiples of unbiased estimates. Is this true in general? Are there other simple forms of strategies which converge rapidly to stable strategies. Is it possible to neatly characterize what initial strategies will result in convergence to a stable strategy or an equilibrium strategy, can one estimate the rate of convergence, and how is the convergence affected if the nature of the auctions changes slightly from one auction to the next (for example, what happens if, at some point, the number of players changes)?

Most of the formal analysis of auctions has been for models with identical players. The symmetry of such models greatly facilitates calculating equilibrium strategies. However, actual auctions, at best, only approximately satisfy such symmetry conditions. While it may be difficult to analyze general asymmetric models, perhaps slightly less general models may be analyzed to determine what is the value of additional (or more accurate) information, what effect a number of amateurs (who use simple bidding strategies and/or receive less accurate information) has on the auction, and

how the outcomes are affected if the players have symmetric information but are allowed to form coalitions which pool members' information and submit joint bids.

Any model is only an approximation of the real situation. The usefulness of any results from analyzing the model depends on the appropriateness of the model. Seldom have auction models been analyzed for their robustness; the question of how much small changes in the model affect the analysis and resulting conclusions has received very little attention. Many of the models require that certain parameters or probability distributions be determined empirically. Inaccuracies in determining these parameters, together with any approximations built into the model, could lead to results and conclusion of little or no relevance to the situation being studied.

A number of questions may be asked related to the robustness of a model. How close to equilibrium are strategies if players are slightly mistaken about the true underlying probability distributions and how does this affect their expected profit and the auctioneers' expected revenue? Is a player's optimal strategy or expected profit strongly affected by any mistakes made by opponents in determining their optimal strategies; in how global of a sense are equilibrium strategies close to equilibrium?

Finally, what are the affects of some of the finer details of actual auctions? Do players obtain useful information in a common progressive auction (as opposed to conducting it as a sealed second bid price auction)? What are the effects of slight collusion (e.g., signalling among players)? What is to be done about the possibility that the rules of the game may be changed after the bidding has started (e.g., a Federal judge may declare an oil lease auction null and void and prevent any immediate development of

the leases)?

Many additional questions can be raised. While some are of a quite broad theoretical nature, others are more specific and should be of concern to anyone currently bidding in auctions. Of course, as long as auctions remain important market mechanisms, the more questions which are answered, the more new questions will arise.

REFERENCES

- R. A. Agnew (1972), "Sequential Bid Selection by Stochastic Approximation," Naval Research Logistical Quarterly, Vol. 19, pp. 137-143.
- K. J. Arrow (1963), "Comment on James S. Dusenbury, 'The Portfolio Approach to the Demand for Money and Other Assets,'" Review of Economics and Statistics, Vol. 45.
- _____ (1965), "Aspects of the Theory of Risk Bearing," The Academic Book Store, Helsinki, Finland.
- E. D. Attanasi (1974), "Some Interpretations of Sequential Bid Pricing Strategies," Management Science, Vol. 20, No. 11, pp. 1424-1427.
- _____ and S. R. Johnson (1975a), "Leasing Policies for the Extractive Industries," Annals of Regional Science, Vol. 10, No. 2, pp. 36-49.
- _____ (1975b), "Sequential Bidding Models: A Decision Theoretic Approach," Industrial Organization Review, Vol. 3, No. 1, pp. 43-55.
- D. P. Baron (1972), "Incentive Contracts and Competitive Bidding," American Economic Review, Vol. 62, No. 3, pp. 384-394.
- _____ (1974), "Incentive Contracts and Competitive Bidding: Reply," American Economic Review, Vol. 64, No. 6, pp. 1072-1073.
- E. M. L. Beale and G. P. M. Heselden (1962), "An Approximate Method of Solving Blotto Games," Naval Research Logistical Quarterly, Vol. 9, No. 2, pp. 65-80.
- M. J. Beckmann (1974), "A Note on Cost Estimation and the Optimal Bidding Strategy," Operations Research, Vol. 22, No. 3, pp. 510-513.
- C. C. Blaydon and P. W. Marshall (1974), "Incentive Contracts and Competitive Bidding: Comment," American Economic Review, Vol. 64, No. 6, pp. 1070-1071.
- S. J. Brams and P. O. Straffin, Jr. (1977), "The Paradox of Player Selection in Sports Drafts," Working Paper, Department of Mathematics, Beloit College.
- K. J. Brown (1975), "A Note on Optimal Fixed-Price Bidding with Uncertain Production Cost," Bell Journal of Economics, Vol. 6, No. 2, pp. 695-697.
- D. Butler (1970), "Topics in the Theory of Fair Division," unpublished masters thesis, Cornell University.
- E. C. Capen, R. V. Clapp and W. M. Campbell (1971), "Competitive Bidding in High Risk Situations," Journal of Petroleum Technology, Vol. 23, pp. 641-651.

- B. J. Casey and L. R. Shaffer (1964), "An Evaluation of Some Competitive Bid Strategies for Contractors," Report #4, Department of Civil Engineering, University of Illinois.
- M. Clerckx and P. A. Naert (1972), "Alternative Bidding Models: Application to a Bidding Problem in the Construction Industry," Working Paper 73-37, European Institute for Advanced Studies in Management, Brussels.
- W. Cook, M. Kirby, and S. Mehndiratta (1975), "A Game Theoretic Approach to a Two Firm Bidding Problem," Naval Research Logistical Quarterly, Vol. 22, No. 4, pp. 721-739.
- P. B. Crawford (1970), "Texas Offshore Bidding Patterns," Journal of Petroleum Technology, March, pp. 283-289.
- E. V. Dean (1964), "Contract Award and Bidding Strategies," Technical Memo 20, Case Institute of Technology.
- K. L. Deavers and J. J. McCall (1966), "Notes on Incentive Contracting," The RAND Corporation, RM-5019-PR.
- E. L. Dougherty and J. Lohrenz (1976), "Statistical Analysis for Bids for Federal Offshore Leases," Journal of Petroleum Technology, November, pp. 1377-1390.
- _____ (1977a), "Money Left on the Table in Sealed, Competitive Bidding: Federal Offshore Oil and Gas Lease Bids," Paper Presented at Society of Petroleum Engineers Meeting, Dallas, Texas, February 21-22, 1977.
- _____ (1977b), "Statistical Analysis of Solo and Joint Bids for Federal Offshore Oil and Gas Leases," Paper Presented at Society of Petroleum Engineers Meeting, April 1977.
- _____ and M. Nozaki (1975), "Determining Optimal Bid Fraction," Journal of Petroleum Technology, March, pp. 349-356.
- L. E. Dubins (1977), "Group Decision Devices," American Mathematical Monthly, Vol. 84, No. 5, pp. 350-355.
- R. Engelbrecht (1977), "A Note on Multivariate Risk and Separable Utility Functions," Management Science, Vol. 23, No. 10, pp. 1143-1144.
- R. Engelbrecht-Wiggans (1977), "On the Fair and Efficient Allocation of Indivisible Commodities," School of Operations Research and Industrial Engineering Technical Report #356, Cornell University.
- _____ (1978a), "An Example of a Multi-Object Auction Game," Cowles Foundation Discussion Paper #494, Yale University.
- _____ (1978b), "A Model for the Distribution of the Number of Bidders in an Auction," Cowles Foundation Discussion Paper #495, Yale University.

- _____ (1978c), "Multiplicative Bidding and Convergence to Equilibrium Strategies," Cowles Foundation Discussion Paper #500, Yale University.
- I. N. Fisher (1969), "An Evaluation of Incentive Contracting Experience," Naval Research Logistical Quarterly, Vol. 16, pp. 63-83.
- L. Friedman (1956), "A Competitive Bidding Strategy," Operations Research, Vol. 4, pp. 104-112.
- M. Friedman (1959), in U.S. Congress, Joint Economic Committee, "Hearing; Employment, Growth, and Price Levels," 86th Congress, First Session.
- _____ (1963), "Price Determination in the United States Treasury Bill Market, A Comment," Review of Economics and Statistics, pp. 318-320.
- M. Gates (1967), "Bidding Strategies and Probabilities," Journal of the Construction Division, ASCE, Paper 5159, pp. 75-107.
- R. J. Gilbert (undated), "Valuation Uncertainty and Competitive Bidding," undated manuscript.
- H. Goldstein (1962), "The Friedman Proposal for Auctioning Treasury Bills," Journal of Political Economy, Vol. 70, No. 4, pp. 386-392.
- D. Goodman and H. Baurmeister (1976), "A Computational Algorithm for Multi-Contract Bidding under Constraints," Management Science, Vol. 22, No. 7, pp. 788-798.
- J. H. Griesmer and M. Shubik (1963), "Toward a Study of Bidding Processes, Part Three. Some Special Models," Naval Research Logistical Quarterly, Vol. 10, No. 3, pp. 199-217.
- F. Hanssmann and B. H. P. Rivett (1959), "Competitive Bidding," Operational Research Quarterly, Vol. 10, No. 1, pp. 49-55.
- D. L. Hanson and C. F. Menezes (1968), "Risk Aversion and Bidding Theory," in J. P. Quirk and A. M. Zarley (eds.), Papers in Quantitative Economics, Lawrence, pp. 521-542.
- J. C. Harsanyi (1967a), "Games with Incomplete Information Played by 'Bayesian' Players, Part I. The Basic Model," Management Science.
- _____ (1967b), "Games with Incomplete Information Played by 'Bayesian' Players, Part II. Bayesian Equilibrium Points," Management Science, Vol. 14, No. 5, pp. 320-334.
- _____ (1967c), "Games with Incomplete Information Played by 'Bayesian' Players, Part III. The Basic Probability Distribution of the Game," Management Science, Vol. 14, No. 7, pp. 486-502.
- D. Hughart (1975), "Informational Asymmetry, Bidding Strategies, and the Marketing of Offshore Petroleum Leases," Journal of Political Economy, Vol. 83, No. 5, pp. 969-986.

- R. J. Kalter, W. E. Tyner, and D. W. Hughes (1975), "Alternative Energy Leasing Strategies for the Outer Continental Shelf," Department of Agricultural Economics Working Paper, Cornell University.
- A. Z. Keller and R. H. Bor (1978), "Strategic Aspects of Bidding Against an Unknown Number of Bidders," paper presented at TIMS-ORSA Meeting, New York, N.Y., May.
- K. O. Kortanek, J. V. Soden, and D. Sodaro (1973), "Profit Analysis and Sequential Bid Pricing Models," Management Science, Vol. 20, No. 3, pp. 397-417.
- I. H. LaValle (1967a), "A Note on the Vickrey Auction," Research paper, School of Business Administration, Tulane University.
- _____ (1967b), "A Bayesian Approach to an Individual Player's Choice of Bid in Competitive Sealed Auctions," Management Science, Vol. 13, No. 7, pp. A584-597.
- Milgrom (1977), "A Convergence Theory for Competitive Bidding with Differential Information," Research Paper #400, Graduate School of Business, Stanford University.
- J. Nash (1950), "The Bargaining Problem," Econometrica, Vol. 18, pp. 155-162.
- Office of OCS Program Coordination (undated), "An Analysis of the Royalty Bidding Experiment in OCS Sale #36."
- M. E. Oren and A. C. Williams (1975), "On Competitive Bidding," Operations Research, Vol. 23, pp. 1072-1079.
- S. S. Oren and M. H. Rothkopf (1975), "Optimal Bidding in Sequential Auctions," Operations Research, Vol. 23, pp. 1080-1090.
- J. W. Pratt (1964), "Risk Aversion in the Small and in the Large," Econometrica, Vol. 32, pp. 122-136.
- H. Raiffa (1970), Decision Analysis, Addison Wesley, Reading, Massachusetts.
- D. K. Reese (1978a), "An Analysis of Alternative Bidding Systems for Leasing Offshore Oil," Working Paper, Graduate School of Business and Public Administration, Cornell University.
- _____ (1978b), "Competitive Bidding for Offshore Petroleum Leases," Working Paper, Graduate School of Business and Public Administration, Cornell University.
- M. Rosenshine (1972), "Bidding Models: Resolution of a Controversy," Journal of the Construction Division, ASCE, Vol. 98, pp. 143-148.
- M. H. Rothkopf (1969), "A Model of Rational Competitive Bidding," Management Science, Vol. 15, No. 7, pp. 362-373.

- _____ (1971), "An Addendum to 'A Model of Rational Competitive Bidding,'" Management Science, Vol. 17, No. 11, pp. 774-777.
- _____ (1977a), "Bidding in Simultaneous Auctions with a Constraint on Exposure," Operations Research, Vol. 25, No. 4, pp. 620-629.
- _____ (1977b), "On Multiplicative Bidding," Working Paper, Xerox, Palo Alto Research Center.
- _____ (1977c), "Equilibrium Linear Bidding Strategies," Working Paper, Xerox, Palo Alto Research Center.
- F. M. Scherer (1964), "The Theory of Contractual Incentives for Cost Reduction," Quarterly Journal of Economics, Vol. 78, No. 2, pp. 257-280.
- D. J. Schilling and J. C. Gallo (undated), "Computer Programs to Detect Collusive Bidding on Public Contracts," Working Paper, Department of Economics, University of Missouri, Columbia, Missouri.
- A. Schotter (1974), "Auctioning Böhm-Bawerks Horses," International Journal of Game Theory, Vol. 3, No. 4, pp. 195-215.
- M. Shubik and R. J. Weber (1978), "Systems Defense Games: Colonel Blotto, Command and Control," Cowles Foundation Discussion Paper #489, Yale University.
- R. F. Scott (1975), "Multivariate Risk Aversion, Utility Independence and Separable Utility Functions," Management Science, Vol. 22, No. 1, pp. 12-21.
- B. Smith and J. Case (1975), "Nash Equilibria in a Sealed Bid Auction," Management Science, Vol. 22, No. 4, pp. 487-497.
- V. Smith (1966), "Bidding Theory and the Treasury Bill Auction: Does Price Discrimination Increase Bill Prices?" Review of Economics and Statistics, Vol. 48, No. 2, pp. 141-146.
- R. M. Stark and R. H. Mayer, Jr. (1971), "Some Multi-Contract Decision Theoretic Competitive Bidding Models," Operations Research, Vol. 19, No. 2, pp. 469-483.
- R. M. Stark and M. H. Rothkopf (1977), "Competitive Bidding: A Comprehensive Bibliography," manuscript, April, 1977.
- W. Vickrey (1961a), "Counterspeculation, Auctions and Competitive Sealed Tenders," Journal of Finance, Vol. 41, No. 1, pp. 8-37.
- _____ (1961b), "Auction and Bidding Games," in Recent Advances in Game Theory, Mathematical Annals #29.
- J. von Neumann and O. Morgenstern (1953), Theory of Games and Economic Behavior, Princeton University Press, Princeton, New Jersey, 1944, 2nd edition 1947, 3rd edition 1953.

R. B. Wilson (1967), "Competitive Bidding with Asymmetrical Information," Management Science, Vol. 13, No. 11, pp. A816-A820.

_____ (1969), "Competitive Bidding with Disparate Information," Management Science, Vol. 15, pp. A446-A448.

_____ (1977a), "A Bidding Model of Perfect Competition," The Review of Economic Studies, Vol. 44, No. 3, pp. 511-518.

_____ (1977b), "The Loss of Revenue in a Share Auction," Research Paper #388, Graduate School of Business, Stanford University.

R. L. Winkler and D. G. Brooks (1977), "Competitive Bidding with Dependent Value Estimates," Working Paper, Indiana University.