Putting Auction Theory to Work: Ascending Auctions with Package Bidding

By Paul Milgrom[*]
First draft: November 22, 2000
This draft: December 15, 2000

Abstract. Ascending auctions with package bidding and their economic uses are explained and the main results of recent FCC-sponsored experiments with such auctions are summarized. A benchmark model is formulated that accounts for the experimental results. In the benchmark, if each bidder bids "straightforwardly" at each round for its potentially most profitable package, then the total payoff is approximately maximized by the final allocation—all payoff approximations here have error bounds proportional to the bid increment. With just two bidders, straightforward bidding strategies constitute an approximate equilibrium, but there can also be other equilibria. A bidder whose competitors all bid straightforwardly has a best reply that entails bidding straightforwardly but with maximal delays, suggesting a concern for the auctions time-to-completion. A new class of simple "bid improvement rules" is introduced that prevents such delays and accelerates the auction without unnecessarily degrading auction performance relative to the unmodified benchmark. The results are used to formulate recommendations for the FCC auction design.

1. Introduction

Large asset sales frequently involve several potentially complementary assets, which might be sold individually or as a package. Since the packaging decision can affect both the seller's receipts and the efficiency of the outcome and since the best decision by a revenue or efficiency criterion[1] generally depends on the bidders' own preferences, it is common for a seller to consult with potential bidders before making its decision.

Sometimes, auction designs allow bidders a choice of packages on which to bid. For example, Cassady (1967) describes a sale of five buildings of a bankrupt real estate firm in which three buildings were defined to constitute a "complex." An ascending auction was used with bids taken for the individual buildings as well as for the complex.[2]

[*] Stanford University and Harvard University, paul@milgrom.net. I thank Evan Kwerel for the probing questions that inspired this research, John Asker for his helpful research assistance in evaluating the Cybernomics experiment, Eva Meyersson-Milgrom for surgical jargon-removal, and Lixin Ye for his comments on a previous draft.

[1] The revenue and efficiency criteria can lead to quite different choices Palfrey (1983). Milgrom (2000a) reports examples in which the sum of total value and auction revenue is constant across packaging decisions, so that there is a dollar for dollar trade-off between creating value and raising revenue.

[2] Sometimes, bidders for large packages are required to bid on certain smaller packages as well. An example is an auction I designed for selling the power portfolio of the Portland General Electric Company (PGE), which was adopted by the company and the Oregon Public Utility Commission. The auction design

Keywords: Auction theory, FCC auctions, Package bidding, Combinatorial bidding, Activity rule, Bid Improvement rule, e-commerce, electronic commerce

In the last few years, there has been growing interest in allowing bidders nearly complete freedom to name the packages they bid on during the auction and to change their choices as the auction unfolds. Unlike the real estate example described above, bidders would even be allowed to bid on non-nested, overlapping packages, as for example when one bidder bids on a pair of items AB and that same bidder or another bids on a pair of items BC. The FCC spectrum auction currently planned for the 700 MHz band is such a design. Other examples include proposed auctions for industrial procurement on the Internet in which individual sellers may offer all or part of a bill of materials and services sought by buyers (Milgrom, 2000b).

The present study is an exploration of the theory of ascending auctions in which bidders are permitted to bid on any packages they wish and to change those packages during the course of the auction. Section 2 describes the background for the theory, including the developments that have led to the increasing interest in the subject. It includes subsections discussing technology, spectrum regulation, generalized Vickrey auctions and their advantages and disadvantages for practical use, the theory of the simultaneous ascending auction (SAA), and experimental evidence regarding the effective limits of the SAA design.

The analysis begins in Section 3 with the formulation of a benchmark package auction design. The analysis in this section employs the assumption that bidders behave myopically according to what I call "straightforward" bidding. The environment is modeled with the same generality as the generalized Vickrey model and in particular covers both substitutes and complements. A theorem establishes the efficiency of the outcome when bidding is straightforward and bid increments are sufficiently small. The analysis shows that this favorable possibility depends critically on the mutual exclusivity of bids made by the same bidder in different rounds and on the possibility of that each bidder's minimum bid on any package depends on that bidder's own past bids.

Section 4 investigates bidder incentives to engage in straightforward bidding in the benchmark auction design. A theorem shows that with just two bidders, straightforward bidding strategies constitute an approximate equilibrium of the benchmark auction game. (In this approximation and the ones that follow, the "error" is proportional to the bid increments, which are specified by the auctioneer.) Another theorem shows that for a bidder interested only in acquiring the whole package of items, straightforward bidding is an approximate best reply strategy when others bid straightforwardly. A third theorem shows that regardless of the preceding conditions, one strategy that is always an approximate best reply to straightforward bidding is "slow straightforward bidding," which is the strategy of passing until bidding by others ceases and then bidding straightforwardly until the auction ends.

If we confine attention to strategies that are eventually straightforward, these findings generally coincide with two important results found in laboratory experiments. Compared to the SAA, the ascending auction with package bidding (1) leads to more nearly efficient equilibrium outcomes and (2) takes many more rounds to reach completion.

---

requires that bidders for the whole package of plants and contracts must also name "decrements" for individual power supply contracts on which there are competing individual bids.

The need for rules to accelerate ascending auctions is a standard topic in discussions of the FCC auctions. The usual approach is to set minimum bids for each package with respect to which activity is measured and to apply an "activity rule."[3] The only similar rule in the benchmark auction restricts bidders to new bids that exceed their own previous high bids on each package.

Section 5 explores an alternative way to accelerate the auction, based on enhancing the bid improvement requirement of the benchmark auction. Traditional minimum bid requirements implicitly ascribe a "value" to each package. When such value information is available, it can be used instead to assess how each bidder's bids stand relative to value and to specify that bids by some buyer at a new round are acceptable only if the set of bids contains one that is better, by that standard, than the best previous bid by that bidder. The proposed rule still accelerates the auction, but it entails fewer distortions than a minimum bid rule and supports the value-maximizing allocation in a wider range of circumstances.

Section 6 considers more implications of the analysis for the FCC auction design and Section 7 concludes by summarizing the reasons to expect more problems with disruptive strategic behavior in the FCC auction than in the initial laboratory experiments.

2. Background

A variety of developments have contributed to the present drive to develop and implement package auctions. These can be grouped into three general categories: rapid advances in technology, favorable developments in regulated spectrum markets and unregulated Internet exchange markets, and research that highlights the potential benefits of package auctions.

Changing Technology and Markets

The most important group factors contributing to the new package bidding designs is associated with the rapid advance of technology, which enables certain new auction designs. To understand the technical challenge, suppose that bidders submit bids for overlapping packages. Given these bids, the first step of finding the sets of "consistent" bids in which each individual item is included in just one package ("sold just once") is a hard computational problem. Then, the total bid associated with each such package must be computed and the revenue-maximizing set of "consistent" bids must be found. All this must be done quickly, while bidders sit in front of their individual computer terminals.

To get an idea of the size of the problem, consider the proposed auction for licenses in the 700 Mhz spectrum in the US, which is presently scheduled to be run in the spring of 2001. The twelve licenses on offer allow for 4095 distinct combinations involving between one and twelve licenses. A decade ago, this number of combinations might have overwhelmed users and posed serious computational problems. Now, however, there are processors, interfaces, algorithms, and communications systems that make it practical for

[3] This is a rule that makes a bidders eligibility for bidding in round *t* of an auction depend on its level of activity, including its standing high bids, in round *t*-1. See Milgrom (1996) for additional description and discussion.

users to identify and bid for many combinations, for auctioneers to compute optimal bid combinations, and for all to verify and track the progress of the auction, even from remote locations.

Even as technology was advancing, markets were changing in ways that facilitate the adoption of sophisticated auction designs. The adoption of US legislation authorizing spectrum auctions in 1993 and the bold decision by the Federal Communications Commission (FCC) the following year to adopt the computerized simultaneous ascending auction (SAA) gave an important boost to advocates of more sophisticated auction designs. The perceived successes of spectrum auctions have led some to propose even more ambitious designs.

In Australia, spectrum regulators eager to "let the market decide" all details of the allocation initiated a serious discussion about the sale of "postage stamp" sized licenses. These would entail very small geographic areas and narrow slivers of bandwidth to be licensed and ultimately recombined as desired by spectrum buyers. Those proposals were shelved because of concerns that the complementarity among the licenses might make the auction and subsequent resale markets perform poorly.

Shortly afterward, another area of hi-tech applications began to develop as Internet-based businesses raced to develop electronic markets that could serve the needs of business customers. Often, industrial buyers seek to purchase not just single components but all the materials and services for a large project, such as a construction project. Multiple suppliers may each supply part of the buyer's needs on terms that may involve quantity discounts, which make the buyer's procurement optimization problem a non-convex one. If such procurements are to be managed by competitive bidding, then some form of package auction will be needed.

These developments and others have inspired new research by economists, operations researchers and computer scientists into the theory and practice of package bidding.

Vickrey Auctions: Advantages and Disadvantages

The theory of package bidding, like so much of auction theory, began with the seminal paper by William Vickrey (1961). Vickrey devised an auction mechanism for the purchase or sale of homogeneous goods. Focusing on the case of a sale, Vickrey's mechanism can be described as follows. Each bidder is asked to report to the auctioneer its entire demand or supply schedule for all possible quantities. The auctioneer uses that information to select the allocation that maximizes the total value. It then charges each buyer a price equal to the lowest bid the buyer could have made to win its part of the final allocation, given the other bids. Vickrey showed that it is in each bidder's interest to report or "bid" its actual demand schedule truthfully, regardless of the bids made by others. Subsequent work by Clarke (1971) and Groves (1973) demonstrated that a generalization of the Vickrey mechanism leads to the same "*dominant strategy property*" in a much wider range of applications. In particular, Vickrey's conclusion holds when the goods are not homogeneous, provided the bids on "all possible quantities" are replaced by bids on "all possible packages." This extension has come to be known as the "generalized Vickrey auction."

These discoveries had profound ramifications. For some operations researchers, they seemed to reduce the economic problem of auction design to a computational problem. If only one could describe and compute values and allocations quickly, it seemed, then the generalized Vickrey auction would become a practical solution to a wide range of resource allocation problems.

For economists, Vickrey's research innovated techniques of analysis and raised expectations about the possibilities for designing effective auctions, but also aroused certain doubts. How well would the generalized Vickrey auction perform if the simplifying assumptions Vickrey used were altered?[4] As it has turned out, a closer look shows that the generalized Vickrey auction has characteristics that limit its usefulness for many applications. Since the objections to the Vickrey design constitute an important part of the motivation for devising ascending auctions with package bidding, let us review the most important ones.[5]

One potential drawback of the Vickrey auction is the sheer number of combinations that a bidder may have to evaluate to bid in the auction. For some applications, the special structure of the problem restores computational tractability. For others, computational issues are becoming less important as technology advances.

In many cases, however, the valuation of large assets is a people-intensive process, the cost of which is not much reduced by advances in technology.[6] Potential buyers who find it too expensive to investigate every packaging alternative will instead choose a few packages to evaluate fully. In these cases, good auction design requires accounting for such costs and choices.

When package evaluation is costly, the choice of packages to evaluate is itself an equilibrium problem, because the profitability of a bidder's choices is affected by the other bidders' choices. For example, suppose the items offered in an auction are {ABCD}. It does little good for a bidder to bid for package AB unless someone else is bidding either on CD or on C and D separately. In a Vickrey auction, since all valuation decisions must be made in advance, bidders must guess about which packages are most

[4] Vickrey's own work expresses doubt about the usefulness of his invention, based on the idea that it would be too costly, but that doubt appears to be misplaced. Williams (1999) finds that all Bayesian mechanisms that yield efficient equilibrium outcomes lead to the same expected equilibrium payments as the generalized Vickrey auctions. This establishes that any tradeoff between payments and efficiency is inherent in the problem and not a special consequence of Vickrey's design.

[5] The following discussion of disadvantages of the Vickrey auction draws heavily from a report to the FCC by Charles River Associates and Market Design Inc (1997). The reports to the FCC and related papers were presented at a conference sponsored by the FCC, the National Science Foundation, and the Stanford Institute for Economic Policy Research. See http://www.fcc.gov/wtb/auctions/combin/papers.html.

[6] Valuing significant business assets involves both investigating the asset itself and creating business plans showing how the assets will be used. For example, a bidder hoping to purchase parts of an electrical generating portfolio might investigate the physical condition of each plant, the availability of land and water for cooling to allow plant expansion, actual and potential transmission capacity, and other physical variables. In addition, it will consider labor and contractual constraints, zoning and other regulatory constraints, the condition of markets in which power might be sold, partnerships that might enhance the asset value, and so on. The final valuation is the result of an optimization over business plans using all this information, and tempered by human judgment. While some of these evaluation activities are shared costs among packages, others are quite idiosyncratic.

relevant and how to allocate their limited evaluation resources. In comparison, a multi-round ascending auction economizes on the need to guess because at least some bidders will be able to adapt their plans based on observations made during the auction.

Vickrey auctions can promote inefficiency by providing perverse incentives for mergers and joint bidding. A simple calculation shows that mergers among bidders in Vickrey auctions results in lower prices for the merged firm or joint bidders without affecting the prices paid by others. This contrasts sharply with incentives in markets where the same price is paid for each unit. In such markets, it is usually the non-participants who benefit most from any non-efficiency-enhancing merger, which can make such mergers hard to arrange. That difference means that Vickrey rules promote inefficient mergers relative to alternative market rules. Even if mergers and joint bidding can somehow be prohibited, the Vickrey design is still vulnerable to other kinds of agreements bidders may make outside the auction, in which bidders agree to adopt Nash equilibrium strategies that reduce both the seller's revenue and the efficiency of the outcome.[7]

A related characteristic drawback of the Vickrey auction that is often considered a drawback is its use of price discrimination: two bidders may pay different prices for identical allocations, even when both have made the same bids for those allocations.[8] Such discriminatory prices are sometimes illegal and often regarded as "unfair."

The theoretical performance of the Vickrey auction is quite sensitive to some of the assumptions of the Vickrey model. Other auction mechanisms have better theoretical performance in a variety of circumstances, such as when there are effective limits on bidder budgets,[9] "common value" uncertainty (Milgrom (1981), Klemperer (1998), or endogenous entry decisions (Bulow, Huang and Klemperer (1999)).

---

[7] Indeed, suppose that A's values are (5,5,5), B's are (5,5,5) and C's are (0,0,20). In a Vickrey auction, if the bidders play their weakly dominant strategies, C will win the package XY and pay a price of 5+5=10, while A and B win nothing and pay nothing. If, however, A and B discuss the matter beforehand without C's knowledge, they could agree to play the (weakly dominated) Nash equilibrium in which A and B each bid 100. With those strategies, each wins one item at a price of zero! (I am indebted to Jeremy Bulow for suggesting this particularly striking example.)

[8] To illustrate the price discrimination problem, suppose there are three bidders—A, B and C—and two items—X and Y. A valuation for a bidder is a triple (x,y,z), specifying how much the bidder would be willing to pay for item X alone, item Y alone, and the package XY. If the parties report valuations of (12,12,12), (11,11,11), and (0,0,20). The result is that A and B will each be awarded an item (at an efficient allocation, either bidder may get either item) at prices of 9 and 8 respectively, even though the items are perfect substitutes.

When the items are not identical, the price discrimination is not so obvious, but the auction outcome is not generally "envy free." Bidders may then prefer the price and allocation assigned to others and may complain on that basis.

[9] Che and Gale (1998) compare first-price and second-price auctions in the face of budget constraints, but the comparison here is somewhat different in this case. To illustrate, suppose there are two identical items. A bidder X has value of 5 for one item and 10 for the pair, but has a budget of just 6, which limits its bids. If X has a single competitor with values of 3 for one item and 7 for a pair and bids accordingly, then X must bid at least 4 to acquire a single item, as required for efficiency. If, however, X's competitor has a value of 3 for one item or 5 for the pair, then the same bid would cause X to lose one item. It should instead bid no more than 2 for one item and 6 for the pair in order to acquire both. Notice that in each case the

Finally, the revelation of bidders’ maximum willingness to pay during the auction can be problematic (Rothkopf, Teisberg and Kahn (1990)), at least for non-computerized auctions in which secure encryption technologies are not available. Winning bidders may fear that information revealed by their bids will be used by auctioneers to cheat them or by third parties to disadvantage them in some negotiation. Similarly, the public has sometimes been outraged when bidders for government assets are permitted to pay significantly less than their announced maximum prices in a Vickrey auction (McMillan (1994)). A bidder’s motive to conceal its information can destroy the dominant strategy property that accounts for much of the appeal of the Vickrey auction.[10]

Simultaneous Ascending Auctions

These drawbacks of the Vickrey auction have created interest in exploring multiple round designs in which bidders must pay the amounts of their own winning bids. The *multiple rounds* feature provides feedback to bidders about relevant packages, economizes on bidder evaluation efforts, conceals the winning bidder’s maximum willingness to pay, and may lead to better performance when common value issues are significant. The “*pay-your-bid*” feature has several possible beneficial effects: it may forestall accusations of price discrimination, alleviate problems associated with budget constraints, and discourage the type of “collusive” strategies possible in a Vickrey auction, in which a bidder increases its own bid solely to reduce a collaborator’s price.

The simultaneous ascending auction (SAA), which has been employed by the FCC in the US for most of its radio spectrum auctions, has the multiple round and pay-your-bid characteristics. At the end of each round, there is a standing high bid and bidder on each item offered for sale. Initially, the standing high bid for each item is zero[11] and the standing high bidder is the seller. At each round, bidders may raise the bid by an integer number of increments on any items that they wish and the process repeats itself until there is a round with no new bids on any item. At that point, bidding on all items is closed and the standing high bids determine the prices.

Although early experimental testing of the SAA demonstrated that it performed well in some environments possibly resembling the FCC environment (Plott (1997)), it has a variety of theoretical limitations. Perhaps the most important of these is its degraded performance in experiments when the items for sale are mutually complementary (Ledyard, Porter and Rangel (1997)), a condition that may have applied to the radio spectrum auctions (Ausubel, Cramton, McAfee and McMillan (1997)). To explain the role of complements in theoretical terms, we compare two different situations.

Vickrey price is less than X’s budget, but sincere bidding is nevertheless not optimal for X and indeed X has no dominant strategy.

[10] Notice that a similar case can be made against ordinary sealed tenders, since the theoretical bid functions are invertible to reveal bidders’ values. In this respect, ascending auctions are theoretically superior to both kinds of sealed bid auctions because they better conceal the winning bidder’s valuation.

[11] A reserve price may also be used.

In the first situation, the items for sale are mutual substitutes[12] for all the bidders. In addition, the bid increment is "small" and the initial prices are low enough to attract at least one bid during the auction for every item. In such cases, suppose bidders bid "straightforwardly" at each round for the items in a package they most prefer at the current prices. Then, the final allocation is efficient and the final prices are competitive equilibrium prices for an economy with "almost" the same values as the actual economy, differing by at most the relevant bid increment (Milgrom, 2000a).[13]

The preceding result demonstrates several things. First, despite the non-convexity associated with the indivisibility of the items, market-clearing prices do exist when goods are substitutes for all bidders. Second, when goods are substitutes, the information communicated during the course of the SAA is rich enough to allow the auction algorithm to discover equilibrium prices and allocations. Third, the auction algorithm has a certain robustness property that allows it to recover from early anomalous bidding behavior. Starting from *any* prices that are sufficiently low, such as the ones that arise early in a long auction when bidders may still be exploring how to bid, the sequence of prices and allocations under straightforward bidding from that point onward still converges to equilibrium prices and an efficient allocation.

In the second situation, some items are sometimes complements.[14] In that case, the conclusions change drastically. Indeed, let S denote the set of valuations in which the bidders regard the items as substitutes. If T is *any strict superset* of S and provided that there are at least two bidders, there exists a profile of valuations drawn from T such that no competitive equilibrium exists.[15]

Intuition for this result is provided by Table 1, which tabulates bidder values. In the table, bidder 1's values are an arbitrary set of values in which the two licenses are complements. Bidder 2's values are then constructed so that (1) the items are substitutes for bidder 2, (2) the unique efficient outcome is for bidder 1 to win both licenses, (3) the total of the prices necessary so that bidder 2 will not want any item exceeds bidder 1's valuation. Together, properties (2) and (3) ensure than no competitive equilibrium exists.

[12] The common terms, "net substitutes" and "gross substitutes," emphasize the distinction between compensated and uncompensated demand. Since models of corporate bidders in the FCC auctions invariably abstract from wealth effects, compensation is irrelevant for them. The important point here is "mutuality"—each good is a substitute for each other good. This mutual substitutes property may be defined by supermodularity of the expenditure function, as in Milgrom and Roberts (1991). For an alternative formulation that treats preferences as primitives, see Gul and Stacchetti (1999).

[13] To facilitate comparison between this result and the analogous result established below for the package auction, it is useful to record the hypothetical bidder values more precisely. For that purpose, let the bid increment for item m in the SAA be I_m. Let l be a bidder and let A_l denote the set of items l acquires in the auction. Then, in the hypothetical economy, for any package T, l's value is $\hat{v}_l(T) = v_l(T) - \sum_{m \notin A_l} I_m$.

[14] The idea that price formation processes behave drastically differently in the cases of substitutes and complements has a long history in economics. Arrow, Bloch and Hurwicz (1959) first established the stability of *tatonnement* in the case of gross substitutes. Milgrom and Roberts (1991) showed that the same sort of stability holds over a vast set of discrete and continuous time, synchronous and asynchronous, backward- and forward-looking price-setting processes. Scarf (1960) provided examples of global instability in the case when the goods are complements, sharply contrasting with the case for substitutes.

[15] This result was first obtained in an early draft of Milgrom (2000a). Shortly afterwards, a draft of Gul and Stacchetti (1999) announced a similar result using a different but equivalent definition of substitutes.

A bidder in the SAA who perceives complementarities among items in the auction that are not perceived by others faces a difficult bidding decision. For example, suppose that the bidder's values of A and B add up to less than its value of the package AB. If the bidder decides that it will bid for both A and B until the prices add up to the value of package AB, it exposes itself to the risk that it will eventually win just one of the items at a price that exceeds its stand-alone value. The result would be a loss for the bidder and possibly an inefficient allocation as well. Yet, if the bidder stops bidding any sooner, it could be passing up a profitable and efficient outcome. This "exposure problem" is inherent to the SAA and makes achieving efficient outcomes problematic. The 1994 decision by the FCC to adopt a rule permitting bid withdrawals subject to a penalty aimed to mitigate this problem, but does not solve it completely.

Experimental Evaluation of Ascending Auction Designs

Besides the theory of package bidding, contributions by economic experimenters played a crucial role. Particularly influential was a study sponsored by the FCC and conducted by Cybernomics (2000) comparing the experimental performance of the SAA to that of a particular combinatorial auction called the simultaneous ascending auction with package bidding (SAAPB). The major findings of that study are summarized in Table 2 below.

The study was conducted under four experimental conditions. In the first, a bidder's value for any package was equal to the sum of its values for the individual items in the package. This condition involves no complementarities. The remaining three conditions involved increasing amounts of complementarity, labeled low, medium and high. Bidder values were drawn at random for each experimental condition and were used twice, once for a group of subjects participating in the non-package auction—the SAA—and once for a group participating in the package auction—the SAAPB. Efficiency in the study was measured by the ratio of the total value of the allocation resulting from the auction to the maximum of that total over all possible allocations.

The experimental results show several prominent features. First, the measured efficiency of the SAA falls off markedly as complementarities increase, but the efficiency of the package auction is largely unaffected by complementarity.[16] Second, the SAAPB took roughly three times as many rounds to reach completion, compared to the SAA. In addition, revenues are higher in all conditions for the SAA compared to the SAAPB.

All experiments require making implementation choices that may affect the experimental outcome. For that reason, experimental results are most convincing when similar results are obtained under a variety of relevant conditions. For example, the Cybernomics experiment involved complementarities only over non-overlapping sets of items, contrary to what some have expected to be the case in the 700 Mhz auction. This absence of overlaps greatly simplifies the problem of identifying efficient allocations.

Most importantly for our purposes, that are aspects of the Cybernomics experiment that make them likely to under-represent the strategic interactions among bidders that

[16] It is interesting that the ascending package auction appears to generate higher efficiency than the SAA even when there are no complementarities.

might be expected in the FCC auctions. First, the experimental subjects’ lack of information about other bidders’ values is not typical of FCC spectrum auctions and may make it much harder for them to exploit the strategic opportunities that the auction affords. Compounding this is the fact that rounds were relatively short, affording subjects little opportunity to evaluate others’ bids and assess the strategic opportunities. Third, the relatively long training sessions that subjects required seemed to highlight their difficulty in understanding the rules, further limiting their ability to exploit gaps in the rules. Long as these sessions were, they fall far short of the preparation undertaken by bidders in the FCC auctions, where the stakes are also very much higher. Finally, unlike bidders in the FCC auction, subjects in the experiments had no access to expert assistance or to analyses that could pinpoint opportunities for strategic bidding.

Despite these limitations, the long history of successes of various “combinatorial auctions” in laboratory settings, beginning with the experiments by Rassenti, Smith and Bulfin (1982), makes it important to take the Cybernomics results seriously. In the next two sections, we provide a theoretical analysis that seeks to account for the results of the Cybernomics experiments and to explore the strategic opportunities that such auctions create.

3. “Straightforward Bidding” in a Package Auction

Let there be finite sets of items to be sold $M=\{1,\ldots,|M|\}$ and of parties $L=\{0,1,\ldots,|L|\}$, with $l=0$ designating the “seller” and $l=1,\ldots,|L|$ designating “bidders.” Each individual l has a valuation vector $v_l=(v_{lA}:A\subset M, A\neq\varnothing)\in\mathbb{R}_+^{2^{|M|}-1}$ the components of which specify the net expected profit it could earn from a business plan using each possible subset of M—these subsets are also called “*packages*.” We also limit and simplify our analysis by the following assumptions:

(i) *Private values*: each bidder l knows its own value vector v_l,

(ii) *No externalities*: A bidder’s payoff does not depend on what its competitors acquire.

a. A bidder l that acquires package A and pays price b_{lA} earns a net payoff of $v_{lA}-b_{lA}$.

b. A bidder l that acquires nothing and pays nothing earns a net payoff of zero.

(iii) *Non-negative marginal values*: $v_{lA}-v_{lA-\{m\}}\geq 0$ for all l, A, and $m\in A$

(iv) *Zero seller values*: $v_{0A}\equiv 0$

In theory, given this framework, the most general relevant kind of package bidding is the kind allowed in the generalized Vickrey auction, in which bidders are free to make mutually exclusive bids on as many packages of items as they may wish. Such a rule imposes no restrictions on what value a bidder may bid for any package and no restrictions on what packages the bidder may name.

Various other rules governing package bidding have been proposed, in which “bids” from a single bidder are not required to be mutually exclusive. Allowing such bids in the analysis below would have no consequences, for allowing such bids merely enriches the language in which the complete bid vector can be expressed.[17]

Within this framework, auctions that permit *only* bids that are not mutually exclusive can be regarded as imposing restrictions on the bidder’s freedom. For example, the SAA can be regarded as a package auction in which a bidder who bids on any package is required also to bid on every subset of that package using a bid function b that satisfies the additivity condition: $b(A \cup B) = b(A) + b(B)$ for every two disjoint packages A and B. Adopting this perspective, the bids on individual items are just a convenient way of parameterizing bids on all the packages of which the individual items are elements.

Many more details, including ones that are normally left unspecified in game theoretic analyses,[18] are needed to complete the rules of the benchmark model. Here are the ones that play a role in the analysis of straightforward bidding.

First, all bids are firm offers. A bidder can never reduce or withdraw a bid it has made on any package. Any new bids a bidder makes on package A must exceed zero and must also be equal to the bidder’s best old bid on A plus some positive integer number of package-specific (and possibly history-dependent) bid increments.

Second, after each round in the benchmark model, the auctioneer identifies a set of “provisional winners,” which is the *consistent* set of bids (that is, bids on non-overlapping packages) that entails the highest total price. This calculation enforces the mutual exclusivity of bids, that is, each bidder may have only one provisional winner. The full history of winning and losing bids are made public after each round.

Third, the auction continues round by round until there are *two consecutive rounds* with no new bids. The auction then ends and the provisionally winning bids at that time become the winning bids in the auction.

In contrast to the SAA, no “activity rules” are included in the benchmark auction analysis. In the SAA, these rules imposed minimum bidding activity levels in each round of the auction. Bidders who fail to meet these bidding levels lose some or all of their eligibility to continue placing new bids in the auction. Excluding such rules from the benchmark auction prepares us to analyze their significance later in the paper.

Several other differences between the benchmark auction and the SAA merit special emphasis. First, the minimum bids can differ *among bidders* on any item or package. Many proposals for package auctions lack this feature, which plays an important role in our analysis. Second, a bid that was a losing bid at round t can become a provisional winner at *later* round, such as round t+1. This is illustrated in Table 3 by the bid of 5 by bidder Y, which becomes a provisional winner in round R+1 even though it was not one in round R. Third, the price of an item or package can *decrease* from round to round. This

[17] Nissan (1999) investigates the expressive power of various “languages” for package bidding, supposing that the objective of a bidding language is to express the richest possible set of plausible preferences as succinctly as possible.

[18] The details omitted in conventional game theoretic analyses include how long each bidder has to submit its bid, the design of the user interface, how much discretion the auctioneer has to make exceptions, and many more.

is illustrated in the table by the fall in the price of Item A from 5 in round R+1 to 4 in round R+2. In the SAA, prices for individual items can never fall.

These are complicating features that make the auction less transparent for onlookers and that create certain new strategic bidding issues. Nevertheless, without these features, straightforward bidding would not generally lead to such nearly efficient outcomes, as described in the next sub-section.

Package Auctions with Straightforward Bidding

Let H_t denote the list, or "history" of bids made by all bidders up to and including round t. Let $B_{lA}^t = B_l(H_t, A)$ denote the highest bid made by bidder l for package A up to time t and let $B_{lA}^0 \equiv 0$. We assume that the seller sets reserves of zero.[19] If a bidder fails to make a new bid on a package A in any round, we represent that formally by a repetition of the bidder's previous highest bid on A. Thus, a bidding strategy b_l for any bidder l is a map from histories to new bids that satisfies the minimum bid restriction that, for every package A, $b_{lA}^t > B_{lA}^{t-1} \Rightarrow b_{lA}^t \geq \min(r_A, B_{lA}^{t-1} + I_A^{t-1})$, where $I_{lA}^{t-1} \equiv I_l(H_{t-1}, A)$ is the bid increment applicable to bidder l for package A at round t. With the convention that failures to bid on a package are represented by repetitions of the old bid, we have $B_l^t \equiv b_l^t$ for all t>0.

The doubly indexed vector $x = (x_{lA}; l = 1,...,L, A \subset M, A \neq \varnothing)$ designates a "*package assignment*" or "*allocation*." The allocation is "*feasible*" if:

$$\begin{aligned}
&\sum\nolimits_{l\in L} \sum\nolimits_{\{A: m\in A\}} x_{lA} \leq 1 \text{ for all } m \in M \\
&\sum\nolimits_{A\subset M} x_{lA} \leq 1 \qquad \text{for all } l \in L \qquad (1)\\
&x_{lA} \in \{0,1\} \qquad \text{for all } l \in L, A \subset M
\end{aligned}$$

The first two sets of constraints say that each item must be allocated at most once (possibly to the seller) and each bidder can be allocated only one package. The last allows us to interpret x_{lA}=1 as assignment of package A to bidder l and x_{lA}=0 as non-assignment.

The provisional winning allocation for round t, x^{*t} maximizes the sum of the provisionally accepted bids. This sum can be written explicitly or, equivalently, using dot-product notation:

$$x^{*t} \in \underset{x \text{ feasible}}{\operatorname{argmax}} \sum\nolimits_{l\in L} \sum\nolimits_{A\subset M} B_{lA}^t x_{lA}^t = \underset{x \text{ feasible}}{\operatorname{argmax}}\, B^t \bullet x^t \,. \qquad (2)$$

Let us assume that in case there are multiple optima in (2), there is some fixed tie-breaking rule that depends only on the vector of best bids $B^t = \left(B_{lA}^t; l \in L, \varnothing \neq A \subset M\right)$.

We now investigate the "straightforward" strategy of bidding at each round on the package that has the highest profit potential. Professor Charles Plott has called this

[19] Any reserve prices $r = \{r_A : A \subset M, A \neq \varnothing\}$ for the seller could be captured by the "seller's bid function" $B_0^t \equiv r$. For simplicity and focus, no analysis of reserve prices is made in this paper.

strategy "bidding the gradient" and reported that it is consistent with the behavior of at least some subjects in package auction experiments.

The lowest price that l can bid for any package A at round t is l's highest bid from the previous round if l was the provisional winner, or otherwise that bid plus one increment:

$$\underline{B}_{lA}^{t}=\begin{cases}B_{lA}^{t-1} & \text{if } x_{lA}^{*t-1}=1\\ B_{lA}^{t-1}+I_{lA}^{t-1} & \text{otherwise}\end{cases} \tag{3}$$

Let $v_{lA}-\underline{B}_{lA}^{t}$ be called l's "potential profit" from a bid on A at round t and let A_l^t be the set of packages that maximizes that potential profit and on which the potential profit is positive. Then, the straightforward bidding strategy $\hat{b}$ is the one in which l bids only on the packages in A_l^t and makes the minimum bid on each of those.[20] As defined here, "making the minimum bid" formally includes the possibility of making no new bid if either (i) no bid entails positive profits or (ii) the bidder is already the high bidder for the unique items in A_l^t. In mathematical notation, this may be written as follows:

$$\hat{b}_{lA}^{t}=\begin{cases}\underline{B}_{lA}^{t} & \text{if } A\in A_l^t\\ B_{lA}^{t-1} & \text{otherwise}\end{cases} \tag{4}$$

Let T denote the final round of the auction. Let $x^*=x^{*T}$ denote the final allocation. Finally, designate each bidder's profit by $p_l=\sum_{A\subset M}\left(v_{lA}-B_{lA}^{T}\right)x_{lA}^{*}$.

<u>Theorem 1.</u> Suppose that the seller's reserve is zero (or "non-binding"). Let S denote the set of bidders that adopt the straightforward strategies. Then for any feasible allocation x:

$$\begin{aligned}&\sum_{l\in S}\sum_{A\subset M}v_{lA}x_{lA}^{*}+\sum_{l\notin S}\sum_{A\subset M}B_{lA}^{T}x_{lA}^{*}\\ &\geq\sum_{l\in S}\sum_{A\subset M}\left[v_{lA}-I_{lA}^{T}(1-x_{lA}^{*})\right]x_{lA}+\sum_{l\notin S}\sum_{A\subset M}B_{lA}^{T}x_{lA}\end{aligned}. \tag{5}$$

In particular, if all bidders adopt straightforward strategies,

$$\sum_{l\in L}\sum_{A\subset M}v_{lA}x_{lA}^{*}\geq\max_{x\text{ feasible}}\sum_{l\in L}\sum_{A\subset M}\left[v_{lA}-I_{lA}^{T}(1-x_{lA}^{*})\right]x_{lA}. \tag{6}$$

<u>Proof.</u> Pick a bidder $l\in S$ and a package A. Since bidder l does not improve the bid on package A at round T, the following must hold:

$$v_{lA}-B_{lA}^{T}-I_{lA}^{T}\left(1-x_{lA}^{*}\right)\leq p_l. \tag{7}$$

Beginning with (2), for any feasible allocation x, we may calculate as follows:

$$\begin{aligned}0\geq B^{T}\bullet(x-x^{*})&=\sum_{l\notin S}\sum_{A\subset M}B_{lA}^{T}(x_{lA}-x_{lA}^{*})+\sum_{l\in S}\sum_{A\subset M}B_{lA}^{T}(x_{lA}-x_{lA}^{*})\\ &=\sum_{l\notin S}\sum_{A\subset M}B_{lA}^{T}(x_{lA}-x_{lA}^{*})+\sum_{l\in S}\sum_{A\subset M}v_{lA}(x_{lA}-x_{lA}^{*})+\text{Remainder}\end{aligned}$$

[20] This specifies that bidders improve their non-winning bids on any maximally profitable item. A similar analysis applies when l bids on only one item at a round.

Using (7),

$$\text{Remainder} = \sum_{l\in S}\sum_{A\subset M}(v_{lA} - B^T_{lA})(x^*_{lA} - x_{lA})$$
$$= \sum_{l\in S}\sum_{A\subset M}(p_l - v_{lA} + B^T_{lA})x_{lA} \geq -\sum_{l\in S}\sum_{A\subset M} I^T_{lA}\left(1 - x^*_{lA}\right)x^t_{lA}$$

Combining these expressions yields (5). Then, (6) follows instantly. ■

The second part of Theorem 1 has an interpretation in terms of competitive equilibrium for a hypothetical economy in which the participants trade in "packages tailored to individuals." Let $B^T = \left(B^T_{lA}; l = 1,..., L, \varnothing \neq A \subset M\right)$ denote the vector of bidder and package-specific bids at the end of the auction when all bidders bid straightforwardly. Assume that each agent in the hypothetical economy has no value for any package that is not tailored for its use. For packages that are tailored for *l*, agent *l*'s values are given by $\hat{v}_{lA} = v_{lA} - I^T_{lA}(1 - x^*_{lA})$, that is, they coincide with the actual values for the packages actually assigned to the bidder and are one increment less for other packages. We complete the hypothetical economy by introducing a firm that can produce any feasible allocation of the goods at a cost of zero.

For this economy, we can verify that the given prices and allocation constitute a competitive equilibrium. Indeed, by (2), the auction allocation maximizes the hypothetical firm's net profits at the given profits. Also, by construction of the preferences and the definition of straightforward bidding, the hypothetical agents prefer their assigned packages at the prices. Hence, the auction outcome is a competitive equilibrium of the artificial economy described above, and therefore efficient for it. That efficiency conclusion implies inequality (6).

It is important to note that at this competitive equilibrium, the supporting prices are *individual specific* prices. This fact highlights the difficulty of constructing a package auction that is non-discriminatory and yet still supports efficient allocations. We return to this issue in section 5.

Using the discreteness of the goods, Theorem 1 can be used to derive an exact optimality conclusion.

Corollary 2. Suppose each bidder *l* adopts its straightforward strategy $\hat{b}_l$. For any given set of valuations $\{v_l\}$, there exists $\bar{e} > 0$ such that if $\max_{x \text{ feasible}} \sum_{l\in L}\sum_{A\subset M} I^T_{lA}x_{lA} < \bar{e}$, then $\sum_{l\in L}\sum_{A\subset M} v_{lA}x^*_{lA} = \max_{x \text{ feasible}} \sum_{l\in L}\sum_{A\subset M} v_{lA}x_{lA}$.

Proof. Let v^* denote the maximum total value. The finiteness of the set of feasible allocations implies that there exists a highest sub-optimal total value, which we denote by $\hat{v}$. Choosing $\bar{e} = v^* - \hat{v}$, condition (6) implies that $\sum_{l\in L}\sum_{A\subset M} v_{lA}x^*_{lA} > \hat{v}$, which is possible only if the expression is equal to v^*. ■

According to Theorem 1, if bidders did bid straightforwardly in the Cybernomics experiment, levels of efficiency close to 100% would be expected. There are easy extensions of these results to the case in which bidders bid straightforwardly but limit their highest bids to achieve some target profit. If the target profits are expressed as a fraction *p* of the value that is equal across bidders, then efficient outcomes will again be achieved to "within an adjusted bid increment," where the adjusted is to divide the total bid increment by (1-*p*). Other nearby variations yield similar results.

We turn next to an investigation of the bidders' incentives to adopt straightforward strategies or close relatives of straightforward strategies.

4. Bidders' Incentives

One appeal of the straightforward bidding strategy is that a bidder can compute it with no information about how many other bidders there may be, what their values are, and what strategies they are likely to adopt. To the extent that straightforward strategies come close to maximizing expected profits, it is likely that some bidders may adopt them, or something similar to them. The calculation of expected profits, however, depends on what the bidders know or believe about each other's values or strategy, so evaluating the optimality of straightforward bidding from a bidder's point of view necessarily involves assumptions beyond those employed in the previous section.

We begin our investigation by testing the set of environments in which straightforward bidding by a bidder l is an approximate best reply to straightforward bidding by other bidders. If the bidder's uncertainty is concentrated on environments in this class, then straightforward bidding will be at least an approximate best reply for the bidder.

<u>Notation</u>. Let x^{**} denote the total value maximizing allocation and let the following symbols denote (1) a measure of the total bid increment, (2) l's "Vickrey price" for package A, and (3) l's Vickrey profit.

$$\begin{aligned} e^* &= \max_{x \text{ feasible}} \sum_{l \in L} \sum_{A \subset M} I_{lA}^T x_{lA} \\ P_{lA}^V &= \max_{x \text{ feasible}} \sum_{k \neq l} \sum_{B \subset M} v_{kB} x_{kB} - \max_{\substack{x \text{ feasible} \\ x_{lA}=1}} \sum_{k \neq l} \sum_{B \subset M} v_{kB} x_{kB} \\ p_l^V &= \max_A v_{lA} - P_{lA} \end{aligned} \tag{8}$$

<u>Lemma 3</u>. Suppose all bidders besides l bid straightforwardly. Then the profits earned by bidder l using any strategy are bounded above by $p_l^V + e^*$.

<u>Proof</u>. Assume that all bidders except l adopt their straightforward strategies, while l adopts some strategy $\hat{b}_l$. Suppose that the result is that l's final bids are given by the vector B_l^T. If l acquires package A, then by Theorem 1, maximizing the right-hand side of (5) over the feasible allocations x in which l acquires no package leads to:

$$B_{lA}^T + \max_{\substack{x \text{ feasible} \\ x_{lA}=1}} \sum_{k \neq l} \sum_{B \subset M} v_{kB} x_{kB} \geq \max_{x \text{ feasible}} \sum_{k \neq l} \sum_{B \subset M} v_{kB} x_{kB} - e^* \tag{9}$$

and hence to $B_{lA}^T \geq P_{lA}^V - e^*$. Hence, $v_{lA} - B_{lA}^T \leq v_{lA} - P_{lA}^V + e^* \leq p_l^V + e^*$. ■

Theorem 4. Suppose all bidders besides l bid straightforwardly. Suppose in addition that $e^* < \min(d, \bar{e})$ (where $\bar{e}$ is as defined in Corollary 2). Then, the "pseudo-straightforward" bidding strategy—in which l bids as if it were straightforward and its values were equal to $\bar{v}_l(A) = \max(0, v_l(A) + d - p_l^V)$—is a strategy that maximizes l's profit to within d.

Proof. If l's values were given by $\bar{v}_l(A)$ and l bid straightforwardly, then by Corollary 2, the outcome would be efficient: l would be assigned its efficient package at a price entailing a non-negative profit. When l bids the same way with its actual values v, the allocation is unchanged but its profit that is higher by $p_l^V - d$. ■

Theorem 4 can be understood as follows. Bidder l's Vickrey price for any package A is defined to be the lowest value at which the efficient allocation would assign A to l. Hence, by theorem 1, if l bids less than the Vickrey price (minus e^*), it will not acquire A. By the same Theorem, if it bids the Vickrey price plus e^* and bids only for A, it can acquire package A for that price. That l prefers to acquire its efficient package on such terms follows either from the standard analysis of the Vickrey auction or from the First Welfare Theorem and the competitive equilibrium interpretation of theorem 1, as described earlier.

The two following corollaries of theorem 4 describe conditions under which straightforward bidding leads to the same path of bids during the auction as the strategy described in Theorem 4. Indeed, it is obvious that straightforward bidding and modified straightforward bidding begin in the same ways, so it is only necessary to check the ending conditions of the auction. In each of corollaries 5 and 6 below, verification of the ending condition is obvious, so the proofs are omitted.

Corollary 5. Suppose that there are two bidders and full information and that $e^* < \bar{e}$. Then, straightforward bidding is an e^*-equilibrium of the benchmark auction game and the ending prices are within e^* of the Vickrey prices.

The equilibrium described in corollary 5 cannot generally be the only equilibrium of the two-bidder auction. For example, suppose that there are two goods for sale and that the bidders' value triples for the two goods and the package are both given by (10,10,30). Suppose the bid increment for each good and the package is one. If straightforward bidding is an equilibrium (or near-equilibrium), then another equilibrium is described as follows. The first bidder opens by bidding 1 for item A and the second opens by bidding 1 for item B. Each bidder plans to stop bidding if these bids and only these bids are made; otherwise each plans to bid straightforwardly for the remainder of the auction. At these strategies, there is no point during the auction when any bidder could gain more than a single increment by deviating to a different strategy. Thus, this strategy profile is an e^*-equilibrium, yet it leads to a decidedly inefficient outcome and the low price of 1 for each item.

Corollary 6. Suppose that bidder l has positive value only for the package of the whole ($A \neq M \Rightarrow v_{1A} = 0$). If the other bidders bid straightforwardly, then bidding straightforwardly maximizes bidder l's payoff to within e^*.

The three preceding results stand in striking contrast to the corresponding conclusions about the SAA. In the SAA, bidders have systematic incentives to distort their bidding. For example, when the items are substitutes, bidders may engage in "demand reduction" as discussed by Ausubel and Cramton (1998) and Weber (1996), implicitly favoring smaller packages over larger ones. When items are complements, bidders in the SAA may have an incentive for demand exaggeration.[21]

In the package auction, there are no such incentives for package distortion.[22] If others bid straightforwardly using their actual valuation function or some artificial one, then, by theorem 4, a bidder's full information near-best reply coincides with straightforward bidding omitting bids that yield profits less than the Vickrey profit of p_l. This strategy sets the same profit target for all packages and bids only for the package that holds the highest potential profit.

In practice, bidders normally lack the information needed to compute their Vickrey profits, though they may learn something about it during the course of the auction. Put differently, the bidder's concern is that the price it needs to bid to become a provisional winner on any given package *A* may decline during the auction, so a premature high bid can result in missed profit opportunities. From this perspective, corollaries 5 and 6 are explained by the impossibility in those special cases of any reduction during the auction of the relevant prices when competitors bid straightforwardly.

This preceding paragraph suggests that bidders have an incentive to go slowly, in order to avoid making bids early in the auction that will turn out to be unnecessarily high. To investigate this possibility, consider the following strategy, which exploits the absence of activity rules in the benchmark auction game.

Definition. The "slow straightforward" bidding strategy for *l* is to raise *l*'s bid at round *t* only if both of two conditions hold: there were no new bids at *t-1* and *l* is *not* a provisional winner at *t-1*. When raising a bid at round *t*, apply the same formula—(4)—that defines straightforward bidding.

Theorem 7. Suppose all bidders except *l* adopt straightforward strategies and that $e^* < \bar{e}$. Then the "slow straightforward" strategy for *l* maximizes its profit to within e^*.

Proof. Suppose that *l* adopted a slow straightforward strategy, but using valuation function $\bar{v}_l$ as defined in the proof of Theorem 4. Repeating our earlier arguments, the

---

[21] Both kinds of incentives are illustrated by the following example. Suppose that there are two bidders and ten items. Bidder #1 values each item at $10. Bidder #2 values item 1 at $15 and, contingent upon acquiring item 1, values any two additional items at $8 each. The reserve price is $6. Assume that bidder #2 plans to bid ("myopically") up to $15 for item 1 and up to $8 each for a second and third item so long as the total current price for the three is less than $31. By bidding straightforwardly for all items, bidder #1 is likely to acquire 9 items at a price of approximately $8, earning a profit of $18. By withholding demand and bidding for just 7 items, it expects a price of $6 and a profit of $28. This strategy exploits the fact that the last nine items are substitutes for bidder 2. By jump bidding to $17+*e* for item 1 at the first round, it expects a price of $6 on the remaining items and a net profit of $29. This exploits the fact that item 1 is a complement for the remaining items for bidder 2.

[22] Ausubel (1997) makes a similar observation, arguing heuristically that the terminal point of an ascending package auction is likely to coincide with the Nash equilibrium of the corresponding "menu auction" (as analyzed by Bernheim and Whinston (1986)).

outcome would be efficient and would entail prices for *l* not exceeding its Vickrey prices by more than e^*. Apply lemma 3. ■

Not everyone can win at such a waiting game and, indeed, it is not generally possible for there to be an equilibrium in which all bidders earn their Vickrey profits. For example, suppose that there are two items and three bidders. Let bidder 1's value triple be (20,0,20), bidder 2's (0,20,20), and bidder 3's (0,0,10). Then the efficient allocation assigns the two items to bidders 1 and 2, respectively. However, the Vickrey prices are both zero and the Vickrey profits 20, which cannot be an equilibrium outcome, in view of the preferences of bidder 3.

In the example discussed just above, bidders 1 and 2 are implicitly collaborators in a common effort, and they are faced with a free rider problem. Each would like the other to bid up to 10, allowing it to win its item at a price close to zero. If the other bidder bids straightforwardly, it can achieve that outcome by adopting slow straightforward strategy, but both cannot do that. To achieve an efficient outcome, bidders 1 and 2 must somehow split the cost of defeating bidder 3, and the auction defines the bargaining protocol available to them to solve that problem.

5. Minimum Bids and Bid Improvement

Experimental evidence of long auctions combined with the analysis of the previous section revealing an incentive for bidders to "go slow" raise concerns about the amount of time that may be needed to complete an auction such as the one planned by the FCC. Even without strategically motivated slow bidding, straightforward bidding in the benchmark auction design can lead to a very slow pace, as bidders cycle through the packages one at a time to raise the total bid on each by a single increment.

To accelerate the pace of the actual auction, various rules have been proposed, including primarily activity rules and minimum bids. To be effective and to respect the bidders' computational constraints, any such rule should have two characteristics. First, it should ensure that the required minimum bids on all relevant packages rise at every round of the auction, at least near the end of the auction.

Suppose, hypothetically, that the auction were redesigned so that, for each bidder, the minimum bid at each round on each package on which it has previously bid is raised by the some increment I_l^t that does not depend on the package. Suppose bidders were required to bid on all or none of the packages on they had previously bid. The earlier analysis of straightforward bidding would apply to such an auction, and the hypothetical result of such bidding by all bidders would be approximately efficient. However, in addition to the strategic incentives that such an auction would create, it would also require that bidders actually evaluate all packages before the auction, rather than using the progress of the auction to determine which packages to assess and on which to focus the greatest evaluation efforts.

This discussion leads to our specification of the second desirable characteristic of a minimum bid rule. It is that a bidder who discovers, late in the auction, that a certain package *A* on which it has not previously bid may be part of the optimum should not be prevented from making a useful bid for that package. Traditional minimum bid rules cannot implement both desired properties. Minimum bids set too low fail to accelerate the auction; minimum bids set too high block "new entry" late in the auction.

To illustrate the problem, suppose that the following values apply to the subset {1,2} of a larger set of items for sale. Suppose the valuation sub-vectors $(v_{l1}, v_{l2}, v_{l12})$ for bidders l=1,2,3 are (8,8,8), (16,16,16) and (0,0,20), and that other bidders are uninterested in these items. Late in the auction, after losing out on bidding for what had seemed to be potentially more profitable items with indexes 3, 4 and higher, bidder 1 turns its attention to the licenses 1 and 2.

If the minimum bids for items 1 and 2 exceed 8 at the time that bidder 1 first decides to bid for these items, then bidder 1 will be blocked and an efficient allocation will not be achieved. If, however, the minimum bid is less than 8, then the minimum is not binding on bidder 2, who can switch back and forth between the items making new bids and extending the auction.

To solve this problem, I propose a class of alternatives to the minimum bid rule called "*bid improvement rules*," which work as follows. As in the case of minimum bid rules, one first needs a method to assign a tentative value to each package on which a bid might be made. To that end, I adopt a variation of the method that DeMartini, Kwasnica, Ledyard Porter (1999) had used to establish minimum bids for each package in their experiments. The method sets individual item prices to be the dual prices of a linear program such as that obtained by maximizing (2) subject to (1), but relaxing the constraint that the x_{lA} variables be integers. The dual program is the following one:

$$\min_{p_m^t, q_l^t} \sum_{m \in M} p_m^t + \sum_{l \in L} q_l^t \tag{10}$$

subject to

$$\begin{aligned} &\sum_{m \in A} p_m^t + q_l^t \geq B_{lA}^t \text{ for } l \in L, A \subset M \\ &p_m^t \geq 0, q_l^t \geq 0 \text{ for } m \in M, l \in L \end{aligned} \tag{11}$$

The dual variables comprise shadow prices (p^t, q^t), where $p^t = (p_m^t; m \in M)$ is the vector of prices for the constraints that each item is sold only once and $q^t = (q_l^t; l \in L)$ is the vector of prices for the constraints that each bidder is assigned just one package. In case the program has multiple solutions, those solutions form a convex set. The choice among the solutions is irrelevant for the analysis that follows.[23]

[23] One possibility is to choose among the dual optimal solutions as follows. Let V^* be the minimum value of the program. Add the constraint that $\sum_{m \in M} p_m^t + \sum_{l \in L} q_l^t = V^*$ and minimize $\sum_{m \in M} p_m^t$. Let P^* be the resulting minimum. Add the constraint that $\sum_{m \in M} p_m^t = P^*$ and minimize the strictly convex objective $\sum_m \left(p_m^t - p_m^{t-1}\right)^2$. This computationally easy process always leads to a unique solution for the price vector p^t.

Based on the prices p^t, one may compute the following relative value index for each package *A*: $R_A^t = \sum_{m\in A} p_m^t$. Then, the "quality" of bidder *l*'s best bid after any round *t* may be defined to by $Q_l^t = \max_A B_{lA}^t - R_A^t$.

The bid improvement requirement specifies that a bid b_l^t is acceptable if it improves on *l*'s best previous bid by at least one increment, that is, if it includes some package *A* such that $b_{lA}^t - R_A^{t-1} - Q_l^{t-1} \geq I_A^t$.[24] To ensure activity during the auction, this bid improvement rule can be used in conjunction with an activity rule that counts a bidder as active at a round when it either makes an acceptable new bid or when it is a provisional winner from the previous round. According to the activity rule, as soon as a bidder fails to be active in more than $N\geq 0$ rounds, it loses its eligibility to make new bids.

To analyze the impact of such a rule near the end of the auction, consider the following linear program, the primal corresponding to dual problem (10)-(11).

$$\max_x \sum_{l\in L} \sum_{A\in M} B_{lA}^t x_{lA} \tag{12}$$

subject to

$$\begin{aligned} &\sum_{l\in L} \sum_{\{A: m\in A\}} x_{lA} = 1 && \text{for all } m\in M \\ &\sum_{A\subset M} x_{lA} \leq 1 && \text{for all } l\in L \\ &x_{lA} \geq 0 && \text{for all } l\in L, A\subset M \end{aligned} \tag{13}$$

<u>Theorem 8</u>. Suppose that the bid maximization linear program (maximizing (12) subject to (13)) has an integer solution x^T and define $p_l^T \equiv \sum_{A\subset M} \left(v_{lA} - b_{lA}^T\right) x_{lA}^T$. Consider an auction with a bid improvement rule based on $R_A^t \equiv p^t \cdot 1_A$. Suppose that, at the final round *T*, (i) no unprofitable bids have been made ($p_l^T \geq 0$) and (ii) no bidder has foregone making an acceptable bid under the bid improvement rule that is more profitable than its x^T-assignment (for each bidder *l* and package *A*, $R_A^T + Q_l^T \geq v_{lA} - p_{lA}$). Then, the final auction allocation x^T maximizes the (unobserved) total value $\sum_{l\in L} \sum_{A\in M} v_{lA}^t x_{lA}$ over all feasible allocations. Also, setting $q_l = Q_l^T + p_l^T$ for each $l\in L$, (p^T,q) is an optimal solution of the dual.

<u>Proof</u>. Focus first on the total *bid* maximization program (12)-(13). By assumption, x^T is an optimal solution of that program and, for some $l = (l_l; l\in L) \geq 0$ (the prices on the one-package-per-bidder constraint), (p^T, l) is an optimal solution of the dual program (10)-(11). By construction, given p^T, $l = Q^T$ is the least vector such that (p^T, l) satisfies (11). Hence, by inspection of the objective (10), (p^T,Q^T) is an optimal solution of the dual. Hence, (x^T,p^T,Q^T) satisfy the feasibility and complementary slackness conditions in the total bid maximization program and its dual.

[24] There may also be additional requirements or alternative bid-quality measures. The original benchmark auction can be interpreted as incorporating multiple bid-quality measures and applying the rule that a bid is acceptable if it is an improvement according to at least one bid-quality measure.

Next turn attention to the (relaxed) total *value* maximization linear program and its dual. By construction, x^T is feasible in the program (the constraints are the same as in the bid maximization problem) and, by the assumption that $R_A^T + Q_l^T \geq v_{lA} - p_{lA}$, (p^T,q^T) is feasible in the dual. Since the dual prices have been increased relative to the total bid maximization problems only for bidders l who acquire items (for whom $\sum_{A\subset M} x_{lA} = 1$), complementary slackness of (x^T,p^T,Q^T) in the (relaxed) total bid maximization programs implies complementary slackness of (x^T,p^T,q^T) in the (relaxed) total value maximization programs. Hence, x^T and (p^T,q^T) are optimal solutions in the latter pair of programs.

By hypothesis, x^T is integer-valued (actually, {0,1}-valued), so that optimal solution of the two linear programs is also a solution of the two integer programs. ■

In general, the hypothesis of the theorem that the linear program has an integer solution will not be satisfied, and in those cases the theorem will not apply as an exact result. For practical purposes, the theorem's primary significance is that it is evidence that the bid improvement rule can sometimes preserve bidders' abilities to begin bidding on potentially valuable new packages until the very end of the auction. Because the rule requires only that *some* bid by a bidder at a round meet the acceptability requirement, a bidder's actual ability to begin bidding on such packages and to continue bidding while they remain profitable is considerable, even when the theorem itself does not apply. Moreover, the rule achieves this flexibility for bidders in a way that still assures that the auction will proceed at a brisk pace. As illustrated above, this combination of characteristics cannot be duplicated by (non-discriminatory) minimum bids on packages.

This preceding analysis does not, of course, take account of bidders' strategic motives when faced with a bid improvement rule. The incentive to slow the progress of the auction, which we have seen is strong in the benchmark auction, is clearly substantially suppressed by the bid improvement rule. Nevertheless, with full information and straightforward bidding by competitors, bidders necessarily have strategies that earn their Vickrey profits, and no rule can eliminate those strategic possibilities. For a bidder who knows its Vickrey profit level, one way to achieve when others bid straightforwardly is to bid straightforwardly oneself, but to stop making new bids once the Vickrey profit levels have been reached, as described in Theorem 4.

6. Implications for the FCC Auction Design

The analysis highlights several important advantages and concerns about package bidding. First, the relative absence of incentives to distort packages may help to explain the excellent efficiency performance of the auction in some laboratory experiments. Second, the incentive for bidders to wait-and-see to learn how much others are bidding offers a possible explanation of the relatively long times to completion of the auction in the lab.[25]

[25] Even without the "slow" version of straightforward bidding, the benchmark auction would take more time than the corresponding SAA, because the rules of the SAA, including especially the activity rule, implicitly force bids to be made on many packages at each round, rather than on just one as in the benchmark auction. However, both the FCC auction and the lab experiment included an activity rule that one might have been expected would lead to a speedy conclusion to the auction.

Several of the specific rules of the FCC auction mitigate the problem of delay. One is the package activity rule, which is absent in the benchmark model. Activity rules penalize bidders who fail to make sufficiently many bids by reducing or removing their eligibility to bid in future rounds. Such rules discourage the most obvious "slow" strategies, but allow bidders to dally by making other bids that they expect cannot win—a practice that has been called "bid parking" or "eligibility parking" in earlier FCC auctions.

Setting minimum bids more accurately can reduce parking by eliminating bids that have little or no chance of winning. However, theory indicates that setting appropriate package-specific minimum bids will be difficult, even when the integer problem is no concern and bidders behave straightforwardly. In contrast, the bid improvement rule does permit the setting of accurate bidder specific minimums in these circumstances.

The current FCC proposal employs a computationally simple minimum bid rule. The proposal also seeks to limit parking by employing an additional rule that sharply limits the number of packages on which a bidder may be active during the auction. If bid evaluation costs are high then, as discussed earlier, the direct cost of such a restriction is likely to be quite small.

7. Conclusion

Are the attractive efficiency properties found in experiments with package bidding likely to be repeated in a high stakes application like the 700 Mhz auction? Or did they result in large part from experimental conditions that favored straightforward bidding over more sophisticated strategies? Are bidders in the FCC auction, after spending substantial sum on analysis, likely to so bid differently from the experimental subjects that very different results should be expected?

There are good reasons to believe that, if the rules of the experimental setting were duplicated exactly, bidders in a real, high stakes auction would bid differently than the subjects in the Cybernomics experiment. Indeed, the serious strategic analysis that experimental subjects could not make in the allotted time but that some FCC bidders will make does reveal unexploited profit opportunities. The "slow" strategies of Theorem 7 illustrate this, and not just for the benchmark auction. The proof of theorem 7 uses only two properties of the auction and the bidders' strategies. The first is that, taken together, the rules and strategies comprise an algorithm to maximize the total value, given any history and whatever value is revealed by bidder l's maximum bids. Any proposed design that claims to achieve an efficient allocation under the bidders' likely strategies necessarily has this first property. The second is that the auction is "sufficiently dynamic," so that bidder l can watch the optimization as it progresses. This, too, is a feature of all the recent proposals to the FCC. Given these characteristics, the strategies that bidders would be inclined to adopt tend to disconfirm the first hypothesis and to lead to long auctions.

The actual FCC rules employ heuristics that aim to avoid the "long auctions" conclusion by limiting bidding strategies. These limits are damaging to efficiency in some possible environments, even if bidders bidder straightforwardly. Without specific information about the environment, it is impossible to know how much damage the heuristics may do.

In evaluating real designs, we must not forget that trade-offs such as these are often unavoidable. The generalized Vickrey auction is the unique auction that makes straightforward bidding by all a Nash equilibrium and still leads to efficient allocations across a wide range of environments. Any alternative auction that avoids the disadvantages of the Vickrey auction necessarily sacrifices that property.

References

Arrow, Kenneth, H.D. Bloch and Leonid Hurwicz (1959), "On the Stability of Competitive Equilibrium, II," *Econometrica* **27**: 82-109.

Ausubel, Larry (1997), "On Generalizing the English Auction," unpublished working paper. http://www.ausubel.com/auction-papers/97wp-generalizing-the-english-auction.pdf.

Ausubel, Larry and Peter Cramton (1998), "Demand Reduction and Inefficiency in Multi-Unit Auctions," working paper, http://www.cramton.umd.edu/papers1995-1999/98wp-demand-reduction.pdf.

Ausubel, Lawrence M., Peter Cramton, R. Preston McAfee, and John McMillan (1997), "Synergies in Wireless Telephony: Evidence from the Broadband PCS Auctions," *Journal of Economics and Management Strategy*, **6**:3, 497-527.

Bernheim, B. Douglas and Michael Whinston (1986), "Menu Auctions, Resource Allocation and Economic Influence," *Quarterly Journal of Economics*, **101**: 1-31.

Bulow, Jeremy, Ming Huang and Paul Klemperer (1999), "Toeholds and Takeovers," *Journal of Political Economy*, **107**(3): 427-454.

Cassady, Ralph (1967). *Auctions and Auctioneering*. Berkeley: University of California Press.

Charles River Associates and Market Design Inc. (1997), "Package Bidding for Spectrum Licenses," Report to the Federal Communications Commission. Available at http://www.fcc.gov/wtb/auctions/combin/cr_2.pdf.

Che, Yeon-Koo and Ian Gale (1998), "Standard Auctions with Financially Constrained Bidders," *Review of Economic Studies*, **65**: 1-21.

Clarke, E.H. (1971), "Multipart Pricing of Public Goods," *Public Choice*, **XI**, 17-33.

Cybernomics (2000), "An Experimental Comparison of the Simultaneous Multiple Round Auction and the CRA Combinatorial Auction." Report to the Federal Communications Commission. Available at http://www.fcc.gov/wtb/auctions/combin/98540191.pdf .

DeMartini, Christine, Anthony Kwasnica, John Ledyard, and David Porter (1999), "A New and Improved Design for Multi-Object Iterative Auctions," Cal Tech working paper.

Groves, Theodore (1973), "Incentives in Teams," *Econometrica* **61**: 617-31.

Gul, Faruk and Ennio Stacchetti (1999), "Walrasian Equilibrium with Gross Substitutes," *Journal of Economic Theory*, **87**, 9-124.

Klemperer, Paul (1998), "Auctions with Almost Common Values," *European Economic Review*, **42**: 757-769.

Keywords: Auction theory, FCC auctions, Package bidding, Combinatorial bidding, Activity rule, Bid Improvement rule, e-commerce, electronic commerce

Ledyard, John, David Porter and Antonio Rangel (1997), "Experiments Testing Multi-object Allocation Mechanisms," *Journal of Economics and Management Strategy*, **6**: 639-675.

McMillan, John (1994), "Selling Spectrum Rights," *Journal of Economics Perspectives*, **8**: 145-162.

Milgrom, Paul (1981), "Rational Expectations, Information Acquisition, and Competitive Bidding," *Econometrica*, **49**: 921-43.

Milgrom, Paul (1996), "Auctioning the Radio Spectrum," www.market-design.com.

Milgrom, Paul (2000a). "Putting Auction Theory to Work: The Simultaneous Ascending Auction," *Journal of Political Economy*, **108**:2, 245-272.

Milgrom, Paul (2000b). *An Economist's Vision of the B2B Marketplace*. Palo Alto: Perfect.

Milgrom, Paul and John Roberts (1991), "Adaptive and Sophisticated Learning in Repeated Normal Form Games," *Games and Economic Behavior*, **3**: 82-100.

Milgrom, Paul and Ilya Segal (2000), "Envelope Theorems for Arbitrary Choice Sets," forthcoming in *Econometrica*.

Nisan, Noam (1999), "Bidding and Allocation in Combinatorial Auctions," working paper, http://www.cs.huji.ac.il/~noam/mkts.html.

Palfrey, Thomas (1983), "Bundling Decisions by a Multiproduct Monopolist with Incomplete Information," *Econometrica*, **51**: 463-483.

Plott, Charles (1997), "Laboratory Experimental Testbeds: Application to the PCS Auction," *Journal of Economics and Management Strategy*, **6**: 605-638.

Rassenti, S.J., V.L. Smith, and R.L. Bulfin (1982), "A Combinatorial Auction Mechanism for Airport Time Slot Allocation," *Bell Journal of Economics* **XIII**: 402-417.

Rothkopf, Michael, Thomas Teisberg and Edward Kahn (1990), "Why Are Vickrey Auctions Rare?" *Journal of Political Economy*, **98**: 94-109.

Scarf, Herbert (1960), "Some Examples of Global Instability of the Competitive Equilibrium," *International Economic Review*, **1**: 157-172.

Vickrey, William (1961), "Counterspeculation, Auctions and Competitive Sealed Tenders," *Journal of Finance*, **XVI**, 8-37.

Weber, Robert (1997), "Making More from Less: Strategic Demand Reduction in the FCC Spectrum Auctions," *Journal of Economics and Management Strategy*, **6**:529-548.

Williams, Steven (1999), "A Characterization of Efficient, Bayesian Incentive Compatible Mechanisms," *Economic Theory* **XIV**: 155-180.

Tables

Table 1: Bidder Values

	Item A	Item B	Package AB
Bidder 1	a	b	a+b+c
Bidder 2	a+αc	b+αc	a+b

The example uses c>0 and 0<α<1.

Table 2: Findings of the Cybernomics Experiment

Complementarity Condition:	None	Low	Medium	High
Efficiency				
SAA (No packages)	97%	90%	82%	79%
SAAPB	99%	96%	98%	96%
Revenues				
SAA (No packages)	4631	8538	5333	5687
SAAPB	4205	8059	4603	4874
Rounds				
SAA (No packages)	8.3	10	10.5	9.5
SAAPB	25.9	28	32.5	31.8

Table 3: Sample Rounds in a Package Auction

	Item A	Item B	Package AB
Round R			
Bidder X	4	0	0
Bidder Y	5	0	0
Bidder Z	0	0	6
Provisional Winning Bids	-	-	6
Round R+1			
Bidder X	4	2	0
Provisional Winning Bids	5	2	-
Round R+2			
Bidder Y	5	6	0
Provisional Winning Bids	4	6	-