COMMENTS ON THE SECOND WYE RIVER PACKAGE BIDDING CONFERENCE

THE CONFERENCE

This conference, like its predecessor, was useful in both academic and practical terms. It fostered a fruitful exchange between computer scientists and economists, who have tended to emphasize different aspects of the auction design problem. It also provided a forum for some new ideas about combinatorial exchanges that we find very interesting and new analyses of computational and economic properties of the auction #31 rules and software, which seem ripe for application in the near future.

As in 1993-94, the FCC design proposals surpass both what theorists have thoroughly analyzed and what experimenters have thoroughly tested. Some of the package bidding research reported at the Wye River Conference, including our own work, breaks new ground and has not yet been digested by other scholars. Even with the new research, there remain significant gaps in our knowledge. Both scholars and practitioners are still in the process of exchanging ideas about (1) how to envision the package bidding problem for purposes of analysis, and what the important issues are from various perspectives, (2) what burdens long and complex auction mechanisms impose upon bidders and how to mitigate those, and even (3) how researchers can integrate the findings of theory, experiments and simulations to resolve open auction design issues.

One conception of package bidding that seemed implicit in the remarks of several conference participants is likely rooted in past package bidding experiments, which initially employed very flexible designs and then added structures, especially communications protocols, to improve the results. The conception is one of an auction with minimal rules in which bidders make and withdraw bids over time in complex environments and, somehow, this unrestrained process searches out an efficient allocation. We do not favor excessively flexible designs, which we think allow unnecessary opportunities for anticompetitive strategies and behavior and which waste too much energy in an unstructured search.

An opposite perspective, mentioned in the pre-conference messages by Mike Rothkopf and Vernon Smith, suggested that some kind of one-shot, sealed-bid auction could be a fruitful alternative, and could perhaps be run at much lower cost to the bidders. There are many ways to set rules for one-shot auctions and, until more details are specified, it would be premature to begin a debate. Here, we note only that the choice between one-shot and open auctions need not be absolute: auctions with a fixed final round such as the Ausubel-Milgrom ascending proxy auction design mentioned below blend the advantages of ascending and one-shot auctions.

A second conception, emphasized by the computer scientists, focuses on two issues: the bidder interface and the computations required by the auction. Karla

Hoffman's presentation at the conference explained the need to streamline computations with admirable clarity. With the existing rules, too much of the computational effort is being devoted to breaking ties and setting minimum bids. We agree that, in principle, streamlining these aspects need not do any economic harm.

As for the bidder interface, the computer scientist's conception differentiates among alternative package designs partly on the basis on the "expressiveness" of the "bidding language," that is, the ability of bidders to make bids that accurately reflect their own package preferences or values. The auction #31 interface necessarily connects bids within each round using the logical disjunction "OR." This means that if a bidder makes, say, two bids in a round, the auctioneer can accept one or the other or BOTH. A problem with these rules is that this OR language fails to be "fully expressive." A bidder who wants to acquire either A or B but NOT both cannot express this preference with a bid made in a single round. The FCC mitigated that problem for auction #31 by adopting Paul's suggestion that bids in different rounds should be treated as mutually exclusive alternatives (called an "XOR" relationship), while allowing a bidder who does not want bids to be mutually exclusive to renew them in the current round.

The actual auction #31 rules thus involves a hybrid of OR and XOR expressions that seems to have confused some commentators. Focusing on the bidding rules within a single round apparently led some to conceive of auction #31 as essentially using the "OR" language. If only OR bids were permitted, withdrawals would be needed to allow a bidder who has bid on A and now wants to bid on B to avoid the exposure problem. In the actual auction #31 rules, however, the mutual exclusivity of bids across rounds eliminates the inter-round exposure problem and with it the need to allow bid withdrawals as protection.

THEORY AND EXPERIMENT

The conference and particularly Vernon Smith's public letter to conference participants raised the issue of how to evaluate designs. Smith has argued that theory is of little value in these conditions and only experiments can credibly be used for policy recommendations.

These issues are hardly new ones: debates about the relationship between theory and experimental data have long been a part of science. When Lord Alfred North Whitehead was asked about which was more important, theory or facts, he answered famously, "Theory about facts." The post-modern view holds more provocatively that there are no "facts" independent of theory: any reporting schema already presumes something about the relevant concepts, which are necessarily based explicitly or implicitly on some theory.

There is no need, however, for us to couch these matters in abstract arguments

about the philosophy of science. For present purposes, it is sufficient to observe that because laboratory experiments can never literally duplicate the conditions of the FCC auctions in all respects, any attempt to extrapolate from laboratory results necessarily relies on some theory. Any recommendations that Smith or others may make uses theory, too. The scientific process must rely on both theory and experiment, and any theory must be stated explicitly enough so that it can compared to all the available evidence and checked for internal consistency.

Given the reported successes of ascending package auctions in the early Cybernomics experiments and the difficulties that subjects encountered in their more recent experiments, the role for theory looms especially large. Are either of the experimental findings generalizable? Which one? What accounts for the very different performances? How accurately do the laboratory conditions reflect the likely conditions of the FCC auctions? Was the short time allowed to bidders in the recent experimental auction important? Are real bidders with more time likely to be less confused by the auction rules? Does it matter that real bidders will invest much more time and analytical effort into bidding than the lab subjects did? Are the laboratory results colored by the particular values used in the experiment?

John Ledyard was the first to work a theory to explain the main facts about ascending package auctions. The two of us have continued the effort. Our findings are stated exactly in the papers we have released, but one can paraphrase what is probably the single most important conclusion of our work for present purposes as follows: "Straightforward" bidding in an ascending package auction with an XOR (or other fully expressive) bidding language leads to outcomes that (1) are efficient and (2) yield competitive total payments to the seller.

This theoretical conclusion is important because it may explain why the ascending package auctions have led to efficient outcomes in some laboratory experiments, as well as to lead to more specific predictions about the final outcome (which should satisfy the set of inequalities that describe the “core“ of the associated cooperative game). When the FCC releases the data we requested some months ago, we will test this prediction and other aspects of our explanation in detail using the experimental data. By identifying a mechanism by which the outcomes are achieved, our theory helps us to understand the likely scope of the experimental findings and increases our confidence in extrapolating from the auction experiments, based as they are on a particular set of value environments, to venture predictions about the likely performance of the FCC auctions.

Such a theoretical conclusion is important for several reasons. First, it guides the search for modifications of the rules that might make the auctions perform even better, both in the laboratory and in the actual circumstances of the FCC auctions. Second, by making the mechanism explicit, it sharpens our questions about whether the auction will perform well outside the laboratory, allowing us to

focus on the features that are responsible for the good performance of package auctions in some experiments.

IMPROVING THE FCC DESIGN

Paul's email messages to the FCC after the May 2000 Wye River auction conference were based on a version of our theory. His messages pointed out that the rules used in earlier laboratory experiments might lead to long delays in the real auction, amplifying concerns about the length of the auctions in the experiments. He argued that if the auction process converges reliably to an efficient outcome for the active bidders, as the experiments seemed to suggest, then it is necessarily in each individual bidder's interest to be as inactive as possible until late in the auction. Such behavior could slow the auction very substantially, leading to long delays. He recommended that the FCC incorporate carefully designed minimum bid and activity rules to mitigate that risk. Evidently, the FCC staff agreed with the main thrust of this analysis, as reflected in the minimum bid and activity rules for auction #31.

The Ausubel-Milgrom theory was designed to account for the relatively successful early auction experiments, but not all the package auction experiments have been so successful. Subjects in the recent FCC-sponsored experiments apparently failed to master the auction rules during the training period. This failure likely reflects some mismatch between the complexity of the rules and the brevity of the training period. While it is far from obvious that the bidders in the actual FCC auction will be as confused and poorly prepared as the laboratory subjects, the complexity of the current FCC design is still an important issue.

Our aim now is to help the FCC to use the theoretical and experimental evidence to make the auction process simpler, faster, more robust, and more resistant to collusion. These goals must be accomplished while still ensuring that the government receives a competitive price for its assets and without biasing the results in favor of particular technologies. (As the Ausubel-Milgrom analysis shows, these last two conditions are not satisfied by the Vickrey auction.)

Based on these objectives and the preceding analysis, we advocate rules and interfaces below that (1) make it easier for a bidder to bid "straightforwardly" for what it most prefers, subject to any financing constraints that it faces and (2) make it harder for bidders to engage in collusion by retaliating against bidders that bid too aggressively or deviate from some expected or promised standard.

USER INTERFACES, BIDDING LANGUAGES, AND VOLUNTARY BID WITHDRAWALS

With these considerations in mind, we turn attention to the bidding language.

There appeared to be a wide consensus at the recent Wye River conference on a primary weakness of the Auction #31 rules: the limited expressiveness of the within-round bidding interface. Much was said at the conference about the use of an XOR or of an OR-of-XOR structure, either of which would be fully expressive. A fully expressive language would make it feasible for bidders to bid straightforwardly. That would be a step in the right direction, but still not sufficient.

The OR-of-XOR structure is said to have two advantages over the XOR structure. First, to the extent that bidders wish to place bids with a corresponding additive structure, it enables bidders to do so easily and compactly. However, straightforward bids in an ascending package auction do not generally have such an additive structure. The goal should be to create a simple interface that makes straightforward bidding easy. This requirement is orthogonal to the OR-of-XOR structural issues.

The second advantage of OR-of-XOR is that reporting bids in this form simplifies the winner determination problem. This is useful, but secondary. Still, it is encouraging that computer scientist Kevin Leyton-Brown reported that it would be possible to keep these computational advantages of an OR-of-XOR interface even with the changes required to make straightforward bidding easy. [We note that he has since amplified his remarks; we make no effort here to incorporate his most recent comments.]

Since a straightforward bidder bids earliest for what is most profitable, it always prefers to have an earlier bid accepted. Consequently, provided the bidding language is fully expressive (and assuming a “private values environment”), straightforward bidders never wish to withdraw a bid. Bid withdrawals can impede the progress of an auction and have at times been utilized by bidders for seemingly anticompetitive purposes. At the same time, the withdrawal facility complicates the software and interface. Thus, provided that the bidding language is sufficiently or fully expressive, voluntary bid withdrawals should be eliminated. Moreover, to facilitate straightforward bidding, a good bidder interface should make it easy for a bidder to increase the bids on all its “active” packages by an equal amount at any round and, perhaps to impose some overall financing limit.

> *Generally, we advocate three changes in the bidder interface for auction #31 that make it easier to bid straightforwardly and harder to collude or delay the auction. These changes, in order of importance, are to eliminate voluntary bid withdrawals, adopt a fully expressive bidding language, and facilitate equal bid increases on all active packages.*

SPEEDING AND STREAMLINING THE PACKAGE AUCTIONS

There are various ways to speed up a package auction without constraining the behavior of straightforward bidders. The FCC already has activity and minimum

bid rules that accelerate the auction. However, these rules for auction #31 were hastily created last year based on the partial analyses available at that time. More is known now. The existing rules can sometimes impede straightforward bidders, blocking desirable bids especially on new packages introduced later in the auction.

The mandatory proxy bidder idea that we have suggested is an effective way to speed the auction and eliminate retaliatory strategies without impeding straightforward bidding. In each of its versions, it automates bidding at most rounds, eliminating the bid entry delays. Our mathematical analysis reveals that proxy bidding makes it possible--perhaps even easy--to reuse optimization information from round to round, reducing computational delays. In the one round version, this auction has the simplicity and speed of a sealed-bid auction, while avoiding the problems of the Vickrey auction design. In a multiple-round version, each bidder is given a fixed number of opportunities--say, three--to revise its proxy values and perhaps its financing limit. The revisions can enhance coordination among bidders who evaluate a limited number of packages. The fixed number eliminates opportunities for retaliation late in the auction, when they count most. This design blends the information-processing advantages of ascending auctions with the speed and simplicity of sealed-tenders.

We have also developed an alternative to activity rules for accelerating the auction: bid quality indexes and bid improvement rules. In a design like the FCC's auction #31 design, we would propose to use these to supplement or replace the minimum bids and activity rules by a requirement, consistent with straightforward bidding, that each non-provisional-winner must "improve" its bid list sufficiently in order to remain active in the auction (here "active" is an all-or-nothing condition, not an eligibility condition measured in MHz-pops).

There are several potentially workable ways to measure bid quality for a bid improvement rule. One measures the quality of a bidder's various packages by the excess (or shortfall) from the shadow prices of an approximating linear program. (Milgrom had advocated one such set of shadow prices last year, but any of the ones proposed at Wye River would also do. The main point is to create some reasonable bid quality index.) That bid, plus an increment, defines a minimum "quality" for the bidder. A bidder is active if it makes at least one bid of the minimum quality or better at a round, or if it is provisional winner from the previous round.

Given such a rule, a bidder can make any bids it likes, provided one meets the quality standard, and provided it satisfies whatever overall limits are imposed on packages. In particular, this rule never prevents an active bidder from introducing a new package that it finds promising. Since any quickly computable bid quality index is just an approximation, this flexibility to introduce new packages can be valuable for improving efficiency.

A more modest alternative is to apply the bid improvement rule to fix minimum

bids for new packages, while retaining the present activity rule for packages that have already received bids. This also has the effect of allowing new packages to be introduced late in the auction when their value becomes apparent, provided they are deemed to be sufficiently more valuable than a bidder's best active bid at the prior round, regardless of whether they meet the minimum bid requirement.

If voluntary bid withdrawals are eliminated, the auctioneer may still wish, at its own initiative, to prune the bid list to simplify the computations during the auction. (We are agnostic about the need for such pruning, limiting our recommendation to how to proceed if pruning is to be implemented.) In this case, the bid quality index serves a second purpose: to guide the auctioneer's decision about which bids to retain from round to round. One possibility is to retain each bidder's N "best bids", defined as any provisionally winning bids plus the highest quality bids among the others. In that way, bids with which other bidders might usefully combine are less likely to be prematurely deleted from the auction. The retained bids might also automatically include the most recent bids, in order to allow other bidders time to combine with new bids that are submitted.

> *The most promising change would be the introduction of the mandatory proxy rules. If this is to be done, some discussion of the rules should begin immediately. Failing that, rules that prevent voluntary withdrawals, retain valuable bids longer, and require bid improvements instead of minimum bids on all packages are all promising smaller changes that could improve the performance of auction #31.*

We hope that some of these newer ideas can still be considered for auction #31. Given the importance of this auction in setting a standard for future package auctions, it is worth intensive effort to get it right. As we have already disclosed, we have applied for patent protection relating to several aspects of the processes described above, including, without limitation, an XOR bidding structure, proxy bidding in package auctions, bid improvement rules, and bid trimming based on bid quality indexes (see U.S. Patent Nos. 5,905,975 and 6,021,398, and other applications pending).

Paul Milgrom and Larry Ausubel
December 14, 2001