An Experimental Basis for Public Policy Initiatives

Charles A. Holt, William M. Shobe, and Angela K. Smith*

Prepared for the Conference on "Promoting the General Welfare: American Democracy and the Political Economy of Government Performance," University of Virginia, January 2005

Abstract

Advances in communications technology and the impetus for deregulation have stimulated the development of new types of markets, some of which are implemented by public agencies on very large scales, e.g. national broadcasting spectrum auctions. In other cases, policy issues arise from the consideration of an economic reorganization of an existing market, e.g. a merger. This paper surveys the use of laboratory experiments in the development and testing of these new types of economic activity, and for evaluating possible effects of changes in existing markets. Specific applications to be considered include: auctions of landing slots, broadcasting bandwidth licenses, tradable emissions permits, and irrigation permits.

I. Introduction

The joining of experimental economics and the recent regulatory reform initiatives has resulted in a current of activity that promises to provide policy makers with a revolutionary tool for public policy design. Experimental techniques allow administrators to base policy on convincing empirical evidence and they provide an environment in which policies can be tested prior to implementation. The results of policy experiments serve both to solidify expert recommendations and persuade policy makers as to the behavior of actual individuals. This new policy design process contrasts sharply with existing efforts where policies are often based on casual empiricism and implicit economic theorizing. The costliness and ineffectiveness of many of these previous efforts have been well documented by advocates of regulatory reform.

With a few notable exceptions, experimental economics originated in efforts to test economic theories that were difficult or impossible to test empirically. Then in the 1990s, government efforts to auction off private rights to use the electromagnetic spectrum provided the single greatest incentive to shift some of the experimental focus to issues of public policy, which meant auction design in this case. Experimental procedures were also applied to test methods for measuring the economic value of goods such as environmental amenities, which are not generally traded in markets. Valuation of non-market commodities is often a crucial determinant of the benefit of proposals aimed to provide or protect environmental amenities. Validation and improvement of measurement techniques through experimentation increased markedly in the decade of the 1990s.

The regulatory reform movement also blossomed toward the latter part of the 20th century. Many policies were plagued by unintended consequences, which pointed to their obvious ineffectiveness and costliness. This factor, as well as the theoretical development of new market-based regulatory mechanisms, lead to modifications of existing policies and a desire to base future policy development on firmer evidentiary ground. Experimental results and insights can provide an important cross check for the analysis of proposed new public policies that might otherwise be evaluated solely in terms of theoretical models and general intuition. The usefulness of experimental methods has become established and some policy economists would no longer feel comfortable recommending a method of auctioning off billions of dollars of broadcasting licenses solely on the basis of a theoretical analysis of Nash equilibrium bidding strategies. Consequently, many laboratory experiments are designed to help economists provide useful advice to policy makers, or as Al Roth termed this activity, "whispering into the ears of princes."[1]

Some of the interest in controlled experiments is driven by a sense of frustration among public officials who fear a loss of control in unfamiliar situations. For example, a group of NASA officials traveled to Caltech several years ago to discuss results of experiments that simulated the budget procurement process associated with large "deep-space" exploration projects. After several failures of Mars expeditions, these officials were worried that the "faster, better, cheaper" approach was not going to work. Anyone

who has been involved in setting up laboratory experiments knows that there are many tough issues associated with devising experiments that have external validity for complex activities like a NASA procurement and project management process. In a candid discussion on the "red-eye" flight back East, one of the authors asked the Washington-based NASA officials what they really hoped to learn from such experiments. The reply was sensible, that even relatively minor insights could have large consequences with so much money on the line.

This paper will discuss a number of innovative public policy experiments that are driven by the need to develop and test new economic institutions, which are emerging as a result of deregulation, regulatory reform, and advances in computation and communications technology. Some of the applications to be discussed are auctions and web-based markets with large numbers of traders at diverse locations, thereby creating a competitive setting that defeats efforts to collude or manipulate market outcomes. The design of new market institutions is a risky business, rather like writing code for a computer program, in the sense that small errors and omissions may produce costly and embarrassing mistakes. Experiments can be used to spot problems with proposed market solutions, and they may provide useful information about likely performance comparisons for settings in which there is no prior evidence on which to base an empirical analysis.

Experiments can also provide an instructive function. Many policy makers may have trouble imagining the changes in behavior that can arise from the incentives created by a new policy regime. For example, market participants might adjust their behavior in response to a new rule to an extent that greatly reduces the net benefit of the rule. Although such changes can be predicted in theory, they may seem counter-intuitive to policy makers not trained in economics. Laboratory experiments can serve to convince policy makers of the likely effectiveness of regulatory reform proposals.

Laboratory experiments provide useful information to policy makers for a relatively modest expense. These studies can be designed and carried out in weeks or months, as long as extensive software development is not required. The results of such studies must be evaluated critically if the context, incentives, and subjects involved call into question the external validity of these experiments for the policy issue under

consideration. These validity concerns can be addressed in part with carefully designed field experiments, and laboratory experiments that have been enriched by including more relevant subject pools and more natural environments.

This paper is not a full survey of public policy experiments, but rather the effort is to develop several key themes on the basis of selected applications. Our perspective comes from two directions, from those who have designed and carried out such experiments, and from those who have developed and implemented policies on the basis of experiments. Indeed, the use of experimental evidence in the formulation of policy advice, in both the public and private sector situations, packs a dual punch. This evidence typically provides the advisor or "expert" with the confidence to make a clear recommendation that is less encumbered with qualifying conditions about theoretical assumptions or the continuation of past trends in data. The dual aspect involves persuasion; those being advised will have a clearer vision of the effects of advice that is based on carefully constructed experiments, especially those with some added field context. In fact, it is often effective to let the policy makers mingle with subjects or even participate in an experiment, as noted by Kenneth Binmore and Paul Klemperer: "We think that their experience in playing the roles of bidders within our experimental software had a significant effect on bolstering the confidence of non-economists on the auction team in the workability of the design. (By contrast, mathematical equations have very little persuasive power.)"[2]

Market-based resource allocations in the public sector typically involve auctions, which have the advantage of being fast, fair, and (usually) transparent. There is a bewildering array of alternative auction formats that can and have been used. The next two sections review cases in which laboratory experiments have been used successfully to select and refine the auction procedures: the 2001 Georgia Irrigation Reduction Auction and the 2004 Virginia NO_x Emission Allowance Auction. Then we consider two cases in which the experiments that were run seem to have had little lasting effect on public policy to date: experiments on the implementation of auction-based allocation of airport landing slots, and experiments designed to evaluate antitrust litigation of certain "facilitating practices" in the *Ethyl* case. The final section summarizes the advantages and disadvantages of using experimental evidence in the design of public policies.

II. The 2001 Georgia Irrigation Reduction Auction

In early 2000, just after the Y2K bug never materialized, much of the Southeastern United States was in the grips of a severe drought, which one Atlanta newspaper termed "Dry 2K." Concern over fresh water flow to the oyster fisheries in the Gulf raised the possibility that water would have to be released from sources that provide drinking water for Atlanta. In April 2000, the Georgia legislature passed a law that called for the use of an auction-like process to restrict agricultural irrigation in certain areas of South Georgia if the Director of the Georgia Environmental Protection Department (EPD) announced a drought emergency. The exact nature of the auction was unspecified, which made an ideal situation for laboratory testing of alternative auction methods that might be recommended.

The law focused on the Flint River watershed. This river, which essentially begins from a drainage pipe near the Atlanta Hartsfield Airport, grows to a size that supports barge traffic by the time it reaches the Florida state line and later empties into the Gulf. Most water usage in this river basin is for agricultural irrigation. A typical irrigation system is a circular array for 50-300 acres, covered by a state permit issued at no charge. Water is not metered, and as a result, it is pumped freely into the fields during dry periods, creating green circles visible from the air, which makes restrictions easy to monitor. The legislature set aside $10 million from the multi-state Tobacco industry settlement to pay farmers not to use one or more of their irrigation permits for the current growing season. This land would not be planted in the absence of irrigation permit.

The main idea behind the law was to pay farmers not to irrigate, which would reduce any negative political impact from farmers who typically resented intrusions from the state capital. A secondary effect was that a bidding process that selects low bids for compensation might retire relatively low-use-value land from irrigation, therefore boosting efficiency levels. The primary motivations, however, were environmental and political, not economic in nature. The goal of the EPD was to eliminate irrigation for as many acres as possible given their budget constraint and to minimize adverse reaction.

The auction format mentioned in the legislation was viewed by EPD officials as being non-coercive, fast, fair, and easy to implement. Speed was important, given the

short interval of several weeks between the March 1 deadline for the declaration of a drought emergency and the optimal planting time. In addition, bids would have to be collected from farmers in diverse locations, which suggested either the use of mail-in bids or web-based communications between officials at a number of bidding sites. In fact, the law had mentioned a "telecommunications process."

This would be a multi-unit auction, since the state could "purchase" many permits by agreeing to compensate permit holders for not irrigating the covered areas for the specified growing season. It became clear early on that bids should be tendered on a per-acre basis to make them comparable. One possible method, a discriminative auction, would have people submitting sealed bids, with the winning low bidders each receiving the amounts that they bid. For example, if the bids were $100, $200, $300, and $400 per acre for 4 permits, and if the two lowest were accepted, then the low bidder would receive $100 per acre and the second low bidder would receive $200 per acre. This auction is "discriminative" since different people receive different amounts for approximately the same amount of irrigation reduction per acre. In contrast, a uniform-price auction would establish a cutoff price and pay all bidders at or below this level an amount that equals the cutoff price. If the cutoff price were $200 in the above example, then the two low bidders would each receive $200, despite the fact that one bidder was willing to accept a compensation of only $100 per acre. Early discussions with state officials indicated a preference for discriminative auctions, in order to avoid the apparent "waste" of paying someone more than they bid, which would happen in a uniform-price auction in which all bidders would receive the same amount per acre. A discriminative auction could be run in several ways, as a one-shot sealed-bid auction or as some kind of a multi-round auction in which bids can be revised after each round. A multi-round auction would have the advantage of providing farmers with some feedback and a chance to learn in an unfamiliar environment. An additional consideration was the desire to find a procedure that would not be significantly affected by collusion among friends and relatives who might share a distrust of Atlanta-based officials.

The initial experiments were run in May of 2000, shortly after the law was passed and almost a year before the actual auction.[3] Although the experiments were not initially contracted by the EPD, officials were aware of them and interested in their outcomes.

After watching several of the experiments, the EPD officials selected two procedures to submit for public comment: a multi-round discriminative auction and a single-stage sealed bid discriminative auction. Subsequent experiments were focused on these two setups. Participants were recruited in groups ranging in size from 8 to over 42. Some of the participants were students from Atlanta, and others were farmers and locals from the South Georgia area. A final field test the following winter involved over 50 local participants bidding simultaneously at 3 different locations in South Georgia to test the software and communication procedures with officials who were watching on-line from Atlanta.

A typical laboratory economics experiment involves a series of interactions in a repetitive *Groundhog Day* setting that permits learning, even if the process is initially difficult to understand. In contrast, an actual irrigation auction would only be run once per year, with participants who would be very unfamiliar with a setting where the low bids are winners. As a result, most sessions began with an unusually long explanation of the auction format, followed by a single auction. It soon became clear from the experiments that a multi-round auction would remove a lot of the uncertainty that bidders would face in such a new situation. The idea was to collect bids in each round, rank them, and post provisional winners. However, these provisionally accepted bids would not be implemented if the EPD officials decided to request new bids in a subsequent round. To encourage serious bidding from the start, the farmers would not be told in advance how many rounds there would be, and this procedure also provided more flexibility for the auction officials. Bids that were unchanged between rounds would be carried over, but bidders would have the option of lowering or raising their bids, based on the provisional results. This process was intended to allow farmers to find out the approximate cutoff price and then compete at the margin to be included.

Laboratory experiments usually involve an intentionally neutral terminology, e.g. “units” of some unspecified commodity, in order to prevent participants from entering with “home-grown values” or preconceived bias. In contrast, a conscious decision was made in these experiments to provide an amount of context (“farmers,” “acres,” and “permits”) that would help subjects understand the situation and the fact that low bids would be accepted. Another unusual procedure was to let participants talk freely, and

even make speeches, in order to simulate the collusion possibilities. Refreshments were served, in order to create a relaxed atmosphere. The only restriction was that participants were not allowed to block access to the bid submission area during an experiment.

In most sessions, participants were given 3 permits, each with a specified number of acres and a "use-value" per acre that would determine earnings for that permit if the bid was not accepted. Subjects were told that the use value would be 0 for permits with accepted bids, in which case the earnings would be the per-acre bid times the number of acres covered by the permit. In both the single-round and multi-round experiments, bids were collected on bid sheets and ranked from low to high. Then starting with the low bids, the total expenditures would be calculated for adding permits with higher bids, until the total expenditure reached the amount allotted by the auctioneer. For example, if the amount to be spent had been announced to be 500, and if the lowest 30 bids yielded a total expenditure of 495, and a 31st bid would take the expenditure above this limit, then only 30 bids would be accepted. This cutoff would determine earnings in the single-round sessions. Earnings equaled the accepted bid or the use value for a permit with a bid that was not accepted. In the sessions with a multi-round setup, the cutoff would be calculated, and the permits with bids below the cutoff would be announced as being "provisionally accepted." Then new bids would be accepted, which would then be ranked, with a new announcement of which tracts had provisionally accepted bids. This process would continue until the experimenter decided to stop the auction, at which time the final acceptance decisions would be announced.

The setup with a fixed budget for permit purchases allows the calculation of a competitive price that can serve as a standard for judging bid levels. The supply function is determined by the use values (opportunity costs) arrayed from low to high, with horizontal segments of a length that is equal to the number of acres with a particular use value. The demand side is determined by the total dollar amount set aside for permit purchases. If B represents the budget available to purchase Q total acres at a price of P per acre, then all money is spent if $PQ = B$, or if $P = B/Q$. The resulting negative relationship between P and Q is graphed as a rectangular hyperbola, and the intersection with the supply function determines a price that would clear the market. The bids, of course, are not all equal, but competition does cause high bidders to come down if they

can in an effort to be included and the average of the accepted bids was generally quite close to this prediction.

One session that generated prices in excess of the competitive prediction provides a good example of how laboratory experiments can be used to pre-test procedures and discover unanticipated problems before they become costly and embarrassing mistakes. In this session (with student subjects), one person asked what would happen if there was a tie at the cutoff bid that exhausted the announced budget, and if there was not enough money in the budget to cover all tied bids. This possibility was not explicitly mentioned in the instructions, and the experimenter in charge generously announced that all tied bids would be included as provisional winners, or as final winners if this were the final round. A tie arose at a price about 5 percent above the competitive level, and all tied bids were provisionally included. In the next round, more bids came in at the previous cutoff level, and this accumulation of tied bids at the focal tie point continued in subsequent rounds. In the end, the resulting payments needed to include all tied bids were about twice the budgeted amount, which would be analogous to spending 20 million in the actual auction instead of the 10 million allocated by the legislature! The next session, run in the same day with a different group of students, used the same procedures and parameters, except that it was announced that a random process would be used to decide which bids to accept in the event of a tie at the cutoff. The average of the accepted bids in this session converged to the competitive level, and the budget constraint was satisfied.

Several other procedural changes were adopted as a result of the experiments. In particular, bids that were below the announced cutoff tended to increase in the next round. This tendency is not surprising, since there is less risk if low bidders could figure out about how high they could have gone in the previous round. This observation was the motivation for running some sessions in which only permit *numbers* (but not the actual bids) for the provisional winners were announced at the end of each round. This information change tended to reduce the upward creep of low bids in successive rounds, without reducing the tendency for high bidders to come down in a scramble to be included. These modifications in tie breaking and announcement procedures were adopted by the EPD in the subsequent auction.

The actual auction was conducted in April 2001, with assistance from a number of experimental economists, including one of the authors (Holt).[4] About 200 farmers showed up at 8 bidding sites, along with television cameramen, reporters and spectators. Bids were received on about 60 percent of the acres that were eligible to be retired from irrigation. The procedures matched those that had been implemented in the experiments, except that farmers' subjective evaluations of their own opportunity costs replaced the redemption use-values that had been induced in the experiment. The acreage amounts were those registered with the EPD. Bids were signed contracts and were entered via web-based forms by auction officials, as had been done in the experiments. The bids were collected at 8 locations and were displayed to top EPD officials, who met in the Experimental Economics Laboratory at Georgia State University where most of the experiments had been done. Then the officials discussed whether to stop the auction, and if not, how much money to release to determine provisional winners for that round. This budget flexibility differed from the procedures that had been used in the experiments, but it did not contradict any of the published auction rules.

Changes in the provisionally released budget caused the cutoff bid to stay roughly the same in the first four rounds, ending up at $125 per acre in round 4. Bids in the $130-$210 range fell round by round, and the number of provisionally included acres therefore increased. The director of the EPD then decided to release more money and raise the cutoff bid to $200 in round 5, after which termination announcement was made.

If the auction officials had used a fixed budget to determine provisional winners, as had been the case in the experiments, then the cutoff bid would have fallen from one round to the next. It is possible that the increasing budgets that prevented the fall in the cutoff price may have discouraged bid reductions at the margin. Moreover, the large budget increase in the final round might have had serious consequences for bidding in a future auction that used the same procedures. Despite these concerns, the 2001 Irrigation Auction was considered to be a success, with about 33,000 acres taken out of irrigation at an average price of about $135 per acre.

In the following year, the state officials decided to run the auction as a single-round discriminative auction by mail, with a maximum reserve price of $150 per acre. In this second auction, 41,000 acres were removed from irrigation, at an average cost of

$143 per acre.[5] This method was less expensive and easier to administer, and attendance at the auction site was less important, since farmers were already familiar with the low-bids-win feature of the previous year's auction. The use of mail-in procedures may have enhanced participation.

III. The Virginia NO_x Auction

The June 2004 auction of nitrogen oxide (NO_x) emission allowances in Virginia provides a fascinating case study of how experiments enabled a government agency to devise a successful procedure under a very tight time constraint.[6] This auction was the responsibility of co-author, Shobe, who was then at Virginia's budget department. The auction involved 1855 one-ton emission allowances for each of two years, 2004 and 2005. These allowances, which represented just over 5 percent of the state's total allotment of allowances, had previously been held and allocated to new users at no cost. The remaining 95 percent of the allowances were allocated to large emitters without cost in 5-year blocks based on the measured heat input at the firm's facilities in the previous 5-year period, a process known as "grandfathering." This use-based procedure has the potentially inefficient side effect of rewarding firms for increasing fuel input and for shifting pollution-generating activities to states with grandfathered allowance allocations. Auction-based allocations would not have these side effects, but the main factor in the state legislature's decision to authorize an auction was the desire to raise revenue in the midst of a budget crisis.

The NO_x allowances are tradable and can be used anywhere in a 19-state region of the eastern US. Although the spot market price provides a reliable measure of value, it would be difficult for a state official to forecast what this price would be at the time of sale. From the policy maker's perspective, there is significant risk in using a spot market for selling state assets. First, the variation in market prices may lead to questions about the timing of the sale, and second, questions may arise about whether particular buyers received special benefits from the sale. An auction provides a fast, transparent, and fair method of selling the permits, thus reducing the risk of claims that the sale was poorly or improperly managed.

The Virginia General Assembly included language in its 2002 budget bill authorizing the Department of Environmental Quality (DEQ) to "…auction the NOx emissions credits…" The budget bill language further required that DEQ report "no later than January 8, 2003," on the process by which the auction would be conducted. The DEQ report provided a brief review of some of the main issues in auction design from available texts and presented two key findings. First, the department tentatively recommended an auction design similar to the first price, sealed bid auction process used by the U.S. Environmental Protection Agency to auction sulfur dioxide allowances. Second, the department concluded that, due to the complexity of the issues involved, the department would need the assistance of "experts in the field" and input from "stakeholders" before finalizing auction design.

Final approval to go ahead with the planning and execution of an allowance auction did not take place until December 8, 2003, but legislative authorization for the auction was set to expire on June 30 of 2004. This meant that the process of auction design and implementation had to be collapsed into a six-month period. While Virginia government had substantial experience selling surplus property at auction, most of those involved felt that this did not provide a good model for auctioning emission allowances. In addition, previous emission allowance auctions had not been done for the purpose of raising revenue. This uncertainty led agency staff to bring in economists from the budget department and to contract for some very quick research help from economists at George Mason University. These researchers conducted a series of experiments to evaluate alternative procedures: a sealed-bid discriminative auction, two multi-round generalizations of this method, and a "clock" driven pricing mechanism. This last mechanism implemented an ascending-bid or "English clock" auction in which the price is raised in a sequence of steps. At each step, the bidders are asked to indicate the numbers of allowances that they are interested in purchasing at that price, and the process stops when the demand falls to a level that equals the number of allowances available.

The experiments were done at George Mason using groups of 8-12 financially motivated student subjects.[7] These group sizes turned out to be quite reasonable, given that there were 17-19 bidders in the actual auctions conducted in June 2004. Unlike the irrigation reduction experiments, the instructions for the allowance auction experiments

used neutral terminology, to avoid the negative connotations associated with pollution activities. Each group of people was involved in 12 auctions in a single session, three repetitions of each of the four alternative auction types mentioned above. Thus there was a considerable degree of repetition, which differed from the approach taken in the irrigation reduction experiments in which most groups only participated on one or two auctions in order to more closely match the fact that the actual auction would only be held once.

The demand-side structure of the market introduced another important experimental issue. One approach is to induce private values; each person is told their own values for one or more allowances, e.g. 200 for the each of the first 5 allowances, 190 for each of the next 5, etc. An alternative, common-value approach would be to have an unknown allowance value, with each person receiving a partially informative "signal" of the common value. Since the future spot market prices of allowances cannot be known in advance, the common value of the NO_x allowances would not be known with certainty, and there were some key uncertainties regarding the supply of allowances in future years. Allowances not used in a given year could be banked and used at a later time, but if the total number of banked allowances exceeded a certain threshold, then all *banked* allowances would lose half of their value (one-ton allowances would only be good for half a ton). There is certainly a common-value element in this case, but given the random-walk nature of the spot prices, it is possible that individuals have approximately the same information about future spot prices, in which case the main value variations across bidders would be due to production capacities and needs, which have largely private-value characteristics. In any case, time constraints precluded designing the experiment software to include common-value elements.

Just as the opportunity "use values" were arrayed into a supply curve in the irrigation reduction auction, the private values assigned to allowance "units" can be arrayed into a demand curve. The degree of individual variations in private values then determines the elasticity of this demand structure. In the absence of direct empirical evidence on the demands of the bidders, the George Mason team decided to run the experiments under several alternative demand conditions. The spot prices in June 2004 were slightly above $2,000 per ton for 2004 allowances and just above $3,000 per ton for

2005 permits. The experiments used value distributions that corresponded to four ranges of private values across bidders: Very Narrow: $2,000-$2,100, Narrow: $2,000-$2,500, Wide: $2,000-$3,000, and Very Wide: $2,000-$3,500. In all four treatment conditions, the multi-round auctions yielded the lowest sales revenues, and the one-round, sealed-bid discriminative auction yielded lower revenues than the English clock auction. The increases in revenue associated with the clock auction relative to the one-round discriminative auction baseline were estimated to be: 17% for the vary narrow value range, 13% for the narrow range, 7% for the wide range, and 2% for the very wide range. All of these difference estimates were statistically significant. Interestingly, these results came as something of a surprise. Initial theoretical analysis had led researchers to the conclusion that the multi-stage sealed bid and the English clock would both perform better than the "standard" sealed bid auction. This experimental result, however, led to the elimination of multi-round sealed bid designs as candidates for selection.

The experiments had been conducted in March of 2004 and reported to agency staff in April, but the results had not been publicized outside of Virginia government at the time that the request for proposal (RFP) for auction brokerage services was published in mid-May. To all agency staff involved, the tight deadline for holding the auction seemed to preclude all but a single-round, sealed bid auction, and all but one of the proposals from brokerage firms, which came back in late May, were limited to this method. The bid from one of the brokerage firms, *Amerex Energy*, surprised agency staff because it proposed using an English clock auction and presented credible evidence of an ability to implement the auction within the now extremely short time-frame. The Amerex proposal was based on their reading of a paper by Cramton and Kerr on the auction of emission allowances.[8] Amerex uncovered this paper through a routine internet search. The contract with Amerex was signed in early June, and the auction was held about 2 weeks later.

The state officials opted for the risky option of using a clock auction in an effort to generate maximum revenue. The final details of the auction were worked out in negotiations over an intense five-day period. The auction software had to be developed and tested in the following week. Some of the issues were whether to use a discriminative or uniform price auction, whether to use a single-round or multi-round

auction, and the order in which the 2004 and 2005 allowances would be sold. The state officials decided on a sequential approach, i.e. to sell the 2004 allowances before the 2005 allowances. The clock auction would begin with a low price, and each bidder could then submit a quantity demanded at that price. Based on advice from the experimenters, agency staff decided not to reveal the level of excess demand at the end of each round of bidding. The clock price would increase at pre-announced time intervals and new quantities would be submitted by bidders until the quantity demanded fell to the available supply (1855 permits for each year). If the quantity were to fall below the supply, then the state would use either the final price or the lower price, depending on which one maximized its revenue. If the bids at the lower price ended up being used, then the order in which bids were submitted would determine which bids would be accepted. Firms were given an additional incentive to bid aggressively early, since the quantity could not be increased from round to round.

Although the experimental design had involved a simultaneous auction of the two vintages of allowances, with the ability of bidders to shift bids between vintages between rounds, this design feature was not implemented in the actual auction. The brokerage firm argued against the multi-vintage design. They had great difficulty in drafting a document explaining the rules of the auction to potential bidders. In the end, the broker strongly advised that the more complicated bidding structure and the obscure mathematics used to select the revenue maximizing set of winning bids would significantly reduce participation in the auction. Thus, the theoretical niceties of combinatorial bidding were scrapped in order to ensure greater participation. This decision was a difficult one for the academic advisors to accept.

Even though the experiments implemented a pure private values structure, state officials gave some consideration to bidder fears of the winner's curse, i.e. the danger that the high bidders are those with value estimates that are biased upwards. In theory, bidders in an ascending clock auction could learn from the fact that others' demands were high enough to keep forcing prices upward. Due to the number of bidders, their geographic dispersion, and the generally competitive character of the emissions market, agency staff did not feel that collusion would be a significant risk. One unanticipated difficulty was the reluctance of some potential buyers to provide financial assurance prior

to the auction. To avoid having bidders strategically default, any bidders not having investment grade credit were required to place money in escrow to cover any bid they might place. Several potential buyers chose not to participate in the auction due to an unwillingness or inability to meet this requirement.

The auction was considered a success in several dimensions. The internet-based interface was inexpensive to run and worked well. The bidders seemed to understand the rules. Participation was broad; there were 19 bidders and 10 winners in the auction of 2004 permits, and 17 bidders and 5 winners in the auction for 2005 allowances. The final price for 2004 allowances was more than 3 percent above the spot market price for these permits. The auction for 2005 allowances, conducted later in the same day, yielded a price that was about 7 percent above the previous spot price. The costs associated with the auction were about $200,000 (including the initial research contract), and the net sales revenue was about $10.5 million, compared to the pre-auction expectations of around $9 million. This case serves as an example where experimental results concerning auction design played a crucial role in the choices made by the policy makers involved.

One unusual institutional factor played a key role in the implementation of the Virginia NO_x auction. Possibly because of the perceived political risks involved, no one in upper management of state government cared to be involved with this first-of-its-kind auction. After the initial contracting phase, this left one middle manager (Shobe), an economist in a civil service position, with full discretion on auction design and implementation. Such freedom of action in the implementation of a new policy is arguably rare but was, at least in this case, a critically important factor in the selection of a policy based largely on experimental results and not constrained by external political considerations.

IV. Airport Landing Slots

The dramatic increase in air travel has generated considerable congestion and delays, especially at airports with severe capacity constraints due to their urban locations. This overwhelming congestion and other efficiency problems might be due in part to the non-market mechanism used to allocate runway time slots that planes need in order to takeoff and land. In response, several proposals have recommended a switch in policy to

auctions or other market-based methods of allocating landing slots. In fact, experimental studies have been conducted to suggest alternative policies that might increase efficiency and reduce congestion. Although these experiments have had an impact on framing the discussion of policy alternatives, the resulting recommendations have not yet been implemented. Moreover, the effect of these experiments might have extended beyond the Federal Aviation Administration (FAA) and inspired experimental evaluations of policies in other contexts by other agencies. This section will focus on these landing-slot allocation experiments and their limited policy influence on the allocation method determined by the FAA, which is in contrast to the direct and immediate impact of the auction experiments discussed in the two previous sections.

Most airports allocate runway time according to a first-come, first-served rule where airlines (or other general aviation craft) pay low weight-based fees, even at peak times. The rationale behind the weight-based fees is that heavier planes inflict more damage on runways. These fees, however, do not represent the value of the runway time, especially for smaller planes during peak hours, which results in inefficient allocations of runway use. Due to congestion created from insatiable demand for runway time under this first-come, first-served rule, especially at large hub airports, a 1968 Federal Aviation Administration ruling set quotas on the number of "slots" (takeoffs and landings in a 30-60 minute timeframe) at four "high density" United States airports (La Guardia, Washington National, John F. Kennedy International, and O'Hare).[9] According to this ruling, slots at these high-density airports were allocated by the unanimous vote of scheduling committees, which were comprised of representatives from all certified airlines operating into or out of a given airport. In these meetings, no discussion was allowed concerning any other aspects of airline competition (city-pairs, profitability, fares, etc.) in order to prevent anti-competitive actions. If the committee could not reach a unanimous decision, the FAA determined the default allocation according to an unspecified rule.

In order to address the growing congestion and possible efficiency concerns associated with the unanimity committee procedure, Grether, Issac, and Plott prepared a report for the Civil Aeronautics Board and the Federal Aviation Administration in 1979.[10] In this report, they discuss results from experiments conducted to test the effectiveness of

the committee allocation method. They ran twenty-three committee experiments with student subjects that varied according to the default allocation, the dependence of payoffs on the outcomes of two consecutive meetings (meant to represent demand interdependencies across airports), the experience level of the subjects, and the number of committee members. The experimenters attended multiple meetings and designed the committee experiments to replicate the observed procedures followed by slot committees, thus adding significant realistic context. However, the instructions used neutral terminology, i.e. colored cards and flags were allocated instead of different time slots at various airports. This commonly applied experimental technique avoids bias due to preconceived values and opinions the subjects might have concerning airport landing slot allocation. Moreover, the size of the experimental committee was either nine or fourteen, a reasonable approximation of the number of airline representatives that attended any given meeting. Although it was difficult to determine exact values for different time slots across airlines due both to a lack of available data and the propensity for these conditions to change over time, the financial incentives in the experiment induced a demand structure that was motivated by key features of several of the high-density airports.

On the basis of these committee experiments, Grether, Isaac, and Plott concluded that outcomes were very sensitive to the threat allocation, which was incorporated into the experiments as either grandfathering (i.e. based on initial allocation), completely random, or mixed (where some slots were taken from those with large initial allocations and distributed randomly to individuals with smaller initial allocations). The results also indicated significant efficiency problems with the unanimous committee system of allocation.[11] First, the experimental data showed that, even when it was profitable and efficient to do so, large participants failed to expand their capacities, and smaller participants failed to grow large enough. Some very inefficient participants received a number of slots from the committee, thus redistributing slots away from high-value uses. Moreover, all entrants, regardless of their underlying costs, exerted a great deal of power because using their veto power would mostly likely result in a default under which they were allocated a positive number of slots. Lastly, the experimental evidence supported the notion that the committee method was unable to deal effectively with the interdependencies in demand for slots at different airports. That is, demand for landing

slots at a given airport depend on the takeoff slots an airline receives at a different airport when they supply service between these two cities.

Grether, Isaac, and Plott also investigated alternative potential allocation methods that might remedy some of the concerns with the committee mechanism. This second group of experiments focused instead on market-based methods but employed terminology and payoffs similar to those used in the committee experiments. Among some of the markets tested were one-price sealed-bid auctions and discriminative auctions (sometimes followed by secondary aftermarkets), "open book" markets, and negative auctions. The design allowed a direct comparison, with identical economic parameters and subjects, of the unanimity committee with an initial allocation default (representing "grandfathering") and a competitive auction with secondary markets. The results indicated a much higher efficiency level with the market process, which the experimenters attributed mostly to the inability of committees to deal with the interdependencies of demand. On the basis of these experiments, Grether, Isaac, and Plott proposed that the FAA consider one-price sealed bid auctions with "aftermarkets." The aftermarkets would allow free purchase and sale of slot permits acquired from the initial auction to increase efficiency. They also noted that the revenues collected should be used to increase runway capacity and that the market allocation should be introduced gradually.

A few years later, Rassenti, Smith, and Bulfin conducted other laboratory experiments that compared the auction method suggested by Grether, Isaac, and Plott to an alternative that incorporated combinatorial bidding.[12] A sealed-bid combinatorial auction allows airlines to submit contingent bids for slots or landing and takeoff slot packages across various airports. They argued that the one-price sealed bid auction of landing slots with aftermarkets did exhibit nicer properties than the committee process, but a primary combinatorial auction (with aftermarkets) would correct the inefficiency that exists because of interdependencies between demand for landing and takeoff slots at different airports. For example, an airline could specify desire to purchase a specific landing slot at airport A only if the bid for a takeoff slot at airport B is accepted. Once the bids were submitted, a complicated algorithm would allocate the slots to packages in order to maximize surplus (assuming bids were true revelations of underlying value).

The algorithm would also determine the prices for the packages, guaranteed to be equal to or lower than accepted bids, based on calculated slot resource shadow prices. Eight experiments were conducted and each experiment used six student subjects with backgrounds in economics and engineering. The treatments varied according to whether they were conducted according to the Grether, Isaac, and Plott recommendation noted above or the proposed combinatorial auction. The experimental data supported the hypothesis that the primary combinatorial auction would allocate more efficiently initially and save transaction costs in aftermarkets.

Although the auction formats suggested by these experimental studies have not yet been implemented, the FAA did shift the process of slot allocation in the direction of a market-based institution in 1986 with the introduction of a buy-sell program.[13] The need for change became apparent when the committees were no longer able to reach unanimous decisions. In response, the buy-sell program initially allocated slots by grandfathering them to existing users and then allowing airlines to purchase and sell the landing slots to other airlines in an attempt to increase efficiency. This program seemed similar to the aftermarkets described in the previous experimental studies, but with the initial grandfathering of existing slots, it inherits potential problems due to entry barriers, anticompetitive behavior, and high transaction costs.

More recently growing concerns have prompted the FAA to consider alternative long-term allocation options, including both administrative procedures and market-based institutions like auctions or congestion fees. Although some of these processes seem similar to those suggested by previous experiments, there is uncertainty as to why the earlier relevant experimental studies did not affect policy at the time. There might be several reasons for this inability of experimental evidence to convince policy makers of the positive attributes of alternative allocation methods that would help to ease congestion problems. One reason might relate to the consequences of transition to a new policy. For example, if the current policy favors large carriers or new entrants, these airlines might be reluctant to switch allocation methods to an auction where they are likely to become worse off (i.e. having to pay for landing slots that are currently allocated to them at no charge) and they can try to put significant pressure on the FAA to abort this structure. This scenario is hard to depict in a traditional experimental setting since

participants are typically unable to affect the design or rules of the game. However, it does seem possible to create an experiment design that incorporates the ability of participants to alter the allocation method with varying degrees of success. Another issue, especially with the combinatorial auction, might potentially arise out of the complexity of the "smart" computer-assisted allocation and pricing rules. There could be problems, or even lawsuits, if participants feel they were unjustly denied a slot due to the allocation based on a non-transparent algorithm. Moreover, there are costs associated with the introduction of a different means of allocation and the question of how to spend revenues from the auction. However, congestion still remains a problem and it will be interesting to see whether previous experimental research is cited or further experiments are performed to test potential new policies.

V. An Example from Antitrust Policy: the *Ethyl* Case

Many issues in antitrust analysis involve conjectures about alternative business practices or industry structures that cannot be directly observed, since mergers, divestitures, or other intrusive actions generally have irreversible consequences. In such situations, laboratory experiments may provide useful insights about policy alternatives being considered. The interpretation of laboratory results for antitrust purposes, however, is plagued by concerns about external validity due to the "distance" between the laboratory and natural market environments. At a minimum, this distance makes it difficult to take the data to a courtroom or other setting where individual culpability is not the same as theoretical likelihood, and where highly paid expert witnesses are eager to point out such problems. This section reviews the use of laboratory experiments in the study of several business practices that were the target of a particularly innovative antitrust case: *Ethyl Corporation* v. *Federal Trade Commission.*

Many economists have a strong intuition that the process of secret price cutting and aggressive discounting will enhance competition, even in relatively concentrated markets. One Federal Trade Commission investigation began with a quick review of transactions data collected from a number of producer-goods markets. This "sweep" revealed a nearly uniform practice of buyers' negotiations of discounts from list prices in most markets. In contrast, discounts were not reported in the market for lead-based

gasoline additives that were in widespread use at that time. This market attracted attention because of a somewhat unusual combination of sales contract provisions that was standard for such lead-based anti-knock compounds.

In particular, contracts for these gasoline additives had a most-favored-customer provision in which the seller promised the buyer the best available price, and that the buyer would automatically receive any discount subsequently offered to another buyer. This provision would obviously be attractive when the buyers are producers who compete with each other in a downstream market, as was the case. The second provision was a "meet-or-release" provision in which the seller promised to meet any lower price that the buyer could obtain, or to release the buyer from the contract. FTC investigators at the time felt that these provisions would deter the discounting, since the seller, who might be willing to grant discount to one buyer under pressure, would resist more forcefully if that discount had to go to all buyers. In addition, discounts designed to attract new business might only have the effect of getting the target buyers to obtain matching low prices from the sellers with which they are currently doing business. If discounts are restrained by these contracts and practices, the remaining problem for sellers is to coordinate on high prices. The FTC officials suspected this might be done through advance notice of price increases, a common practice in this market whereby one seller would announce an intended price increase in the trade press and then wait to see if others followed.

One of the authors was a consultant at the FTC at that time, and the discussion in the hallways centered on the conjecture that this combination of business practices might even allow the four sellers in this market to reach and maintain a *perfectly collusive*, joint-monopoly price, without much worry about new entry since the lead-based additives were to be phased out in the future. At the time the litigation began, these discussions were based largely on intuition, verbal arguments, and some very simple theoretical examples. In an effort to get a firmer view of the effects of the practices that were thought to "facilitate" collusion, the FTC arranged for some laboratory experiments to be conducted by David Grether and Charles Plott.[14]

Considerable effort was devoted to making the experiment design parallel the salient features of the market, in terms of numbers of buyers, sellers, the demand elasticity and cost structures, etc. But in the end, only one set of value and cost

parameters was used, due to the need to evaluate various combinations of the business practices. In thinking about the effects of these facilitating practices, it is useful to consider three benchmarks: the perfectly competitive price level determined by the intersection of demand and supply, the monopoly (perfectly collusive, joint-profit maximizing) price, and an intermediate price that would result if all sellers choose production quantities non-cooperatively in a Cournot equilibrium. The Grether and Plott experiments showed that prices were near competitive levels in a treatment in which none of the practices were used and prices were negotiated freely between buyers and sellers by phone. In contrast, the effect of all practices together tended to raise prices significantly, to levels that were approximately equal to the Cournot benchmark but well below the perfectly-collusive monopoly price.

In a subsequent paper, Charles Holt and David Scheffman construct a game-theoretic model in which it is shown that the meet-or-release and most-favored-customer contracts permit firms to maintain prices above competitive levels, up to the Cournot level, but no higher.[15] The intuition behind this result is that even though sellers choose price, so a Cournot equilibrium for a model with quantity-setting firms may not be relevant, the effect of the practices is to allow sellers to protect their own sales quantity by meeting competitors' price cuts. The most-favored-customer clause prevents sellers from cutting price secretly to attract new sales into the market without offering comparable price cuts to existing buyers. These features effectively transform the price competition into a process that has more in common with quantity competition.

In the end, the FTC won the case, but the Chairman of the FTC, James Miller, wrote a persuasive dissent in which he cautioned against antitrust policies that might depend on particular market conditions and combinations of business practices that could have other pro-competitive effects.[16] The FTC decision was later overruled on appeal in the federal courts. It is worth noting that the experimental evidence was not introduced in the judicial proceedings, due to a concern about how the defendants could counter with external validity arguments.

VI. General Observations

The cases considered in this paper indicate how experimental methods can be used effectively to help design and adjust public policy. The most obvious benefit is that well-crafted experiments give researchers the confidence to make strong recommendations where appropriate. Perhaps equally important is the use of experiments to help policy makers envision the effects of the options being considered, a process that sometimes involves letting the policy makers participate in experiments themselves. In addition, experiments are relatively quick and inexpensive to run, and treatments can be structured to match novel (even risky) policy proposals that have never been used previously. There are, of course, worries that the nature of the subject pool (typically students), the setting (repeated interactions in a laboratory), and incentives (induced values with low financial incentives) may limit the external validity of these experiments. These restrictions are often dictated by the tight time frame under which policy options are considered and implemented. With sufficient time, it is possible to use relevant professionals (e.g. farmers), high incentives, and more natural settings. One point to be made here is that the laboratory setting for auctions and other markets is often quite similar to a field situation, since both may involve the use of web-based interfaces for bid submission and the announcement of results. This distance is much greater for experiments designed to analyze antitrust issues and complicated institutional arrangements, which may explain the fact that experimental work seems to have had a larger effect on auction design than on the implementation of specific antitrust or industrial organization policies.

One common theme that has emerged is that transparency and simplicity are important in designing new policies. The reversal of the original decision on the *Ethyl* was based in part on concerns about policies that depended on an analysis of combinations of practices under specific industry conditions. The bidding instructions for the Georgia Irrigation Reduction auction had to be simple enough so that a one-page version could be read aloud at the eight bid sites (in front of television crews and some skeptical farmers). The tentative decision to run simultaneous auctions for 2004 and 2005 NO_x credits was abandoned when it became apparent that the process of deciding which bids to accept was so complicated that the instructions would not fit on a single page.[17]

These transparency considerations are particularly relevant for the upcoming Federal Communications Commission decisions about whether and how to allow package bidding for bandwidth licenses in the future. With a few exceptions, the FCC currently auctions off bandwidth licenses in a simultaneous ascending price auction, in which bids are collected in each round. Then bidders can view the current provisional winners for each license and decide whether to increase their bids, subject to a set of rather complicated activity rules. Each license is essentially being sold independently, even though the auction is simultaneous, which introduces some risks for firms that are trying to construct a network of licenses in adjacent areas. Some of these risks might be reduced by a combinatorial auction. For example, package bidding would allow a firm to bid 10 on license A, 10 on B, and 30 on AB. This would protect the firm from the "exposure problem" of bidding above its willingness to pay for A alone in the expectation of winning B, which would result in a loss if B were not acquired. Proposed package bidding schemes generally stipulate that the auctioneer will accept the set of package bids that maximizes that sales revenue, subject to the feasibility constraints. It is well known that this revenue maximization problem may not be solvable in finite time. In multi-stage package bidding auctions, this revenue-maximization problem would have to be solved at each stage before results at that stage could be released. Of course, the search for maximizing combinations can be truncated, but it is easy to imagine negative reactions that this might generate among bidders, and perhaps even in the courts if the results were to be challenged.[18] Combinatorial auction proposals have been tested with laboratory experiments, using integer-programming techniques to select the winning combination at each stage.[19] While these experiments indicate that the approach is feasible, the NO_x experience indicates that the tolerance for non-transparency may be lower in the interface between bureaucrats and firms than it is in a laboratory with paid participants.

This review of the use of economic experiments to guide policy decisions reminds us of the importance of the institutional context of economic decisions. Both private economic agents and policy makers must take their institutional environment into account in every decision. For the laboratory subjects and their counterparts in the non-experimental setting, there will always be important unobserved, and hence uncontrolled, constraints and margins of substitution. As a result, experiments may best be seen as

design tools for engineering incremental improvements in how policies are implemented. At each stage of this process, the economist as policy engineer will experiment at the margin expected to bring about the greatest incremental improvement in policy performance or the greatest expansion in the range of solvable policy problems.

Even in the relatively straightforward area of auction design, big differences will exist between the prototype in the lab and the product in the field. Much room is left for judgment in applying the new tools. In the Georgia irrigation buy-down and the NO_x auction, the policy makers used the plans provided by the economic engineers and modified them as needed to match the evolving information about how potential bidders in the respective auctions were likely to respond in actual practice. In Georgia, design changes were made on the fly to prevent evolving information from giving bidders perverse bidding incentives. In the NO_x auction, the theoretically attractive combinatorial auction, which would have permitted efficient arbitrage across auctions, was scrapped in the face of probable bidder aversion to the difficulty of knowing in advance what bid combinations would be selected by the computer-assisted revenue-maximization program.

From the policy maker's point of view, newly engineered designs broaden the range of available tools or help eliminate the obviously bad alternatives. Managers of the irrigation auction averted a political disaster by pre-auction design testing. Or, experiments may provide proof that certain behavioral margins are or are not important to policy choices. Even in the face of this evidence, policy choices may be constrained away from the optimum by cost, political considerations, equity, or lack of management support.

This view of experimental methods and policy suggests that experimental methods may be easier to apply where policy choices involve repeated incremental trials, i.e. the selling of emission allowances, spectrum rights, or irrigation reduction subsidies. In cases involving large discrete changes or irreversible changes, one may expect the uncertainties over the applicability of experimental methods to result in their being given less weight as a factor in the policy decision. Experiments are more likely to have the role of providing information about the importance of various behavioral margins then they are in determining the outcome of the ultimate policy choice. Thus, in the *Ethyl*

case, experiments provided information about the potential collusive value of Ethyl Corporation's pricing policies but could not provide specific evidence needed in the judicial proceedings to justify anti-trust action against *Ethyl*.

Economic policy may usefully be viewed as the output of a research and development process where product innovation advances through a series of stages: basic research, engineering, design, and implementation. As with any product innovation, once the product is actually put to use, the infinite variety of constraints and incentives of an actual rather than imagined economy will expose new margins for product improvement. Some of these margins for improvement will be exploited by the policy makers themselves, while others will likely result in another round of policy R&D. At each stage of the policy R&D process, experimental methods provide an increasingly effective tool for testing and refining ideas and designs. As these methods have become more sophisticated and more generally accepted in the economics profession as a test bed for economic ideas, their use as a guide for the improved design and implementation of public policy has increased dramatically. Laboratory experimentation in economics has clearly come of age as one of the essential tools of policy design.

Endnotes

[*] Holt and Smith: Department of Economics, University of Virginia; Shobe: Weldon Cooper Center for Public Policy, University of Virginia. This research was funded by grants from the National Science Foundation and the University of Virginia Bankard Fund.

[1] Alvin Roth, "Laboratory Experimentation in Economics," in T. Bewley, ed., Advances in Economic Theory, Fifth World Congress, Cambridge: Cambridge University Press, 269-299.

[2] Kenneth Binmore and Paul Klemperer, "The Biggest Auction Ever. The Sale of the British 3G Telecom Licences," *Economic Journal*, vol. 112 (2002), pp. 74-96, C85.

[3] Ronald Cummings, R., Charles A. Holt, and Susan K. Laury, "Using Laboratory Experiments for Policy Making: An Example from the Georgia Irrigation Reduction Auction," *Journal of Policy Analysis and Management*, 23(2), (Spring, 2004) 341.

[4] The auction was administered by Professors Susan Laury and Ronald Cummings from Georgia State. Other experimental economists assisting included Laura Taylor (Georgia State), Lisa Anderson (William and Mary), and Mark Van Boening (Mississippi).

[5] R. McDowell, "Going Once, Going Twice...," *GMDA News*, 2(2), 1.

[6] William Shobe, "Allowances for Sale: Virginia's NOx Allowance Auction." Paper presented at the fall meetings of the Emission Marketing Association, Toronto, Ontario, September 24, 2004.

[7] David Porter, Steven Rassenti, Vernon Smith, and Abel Winn, "Final Report for Virginia DEQ NOx Auction Design," George Mason University, Interdisciplinary Center for Economic Science, 2004.

[8] Peter Cramton, P., and S. Kerr (1999). "Tradeable Carbon Permit Auctions - How and why to auction not grandfather," Working Paper. University of Maryland, March 1999.

[9] More recently, in April 2000, the Wendall H. Ford Aviation Investment and Reform Act of the 21st Century (AIR-21) mandated exemptions from high density rule operation limits for certain flights either operated by new entrants or by smaller aircraft providing service to small airport communities. The AIR-21 also stated that the high-density rule would be phased out by 2007. As the number of exemptions grew and congestion was compounded, the FAA ruled in November 2000 to limit the number of exemption slots to 159 within a given period of the day and to allocate these slots via a lottery among all eligible flights.

[10] David M. Grether, R. Mark Isaac, and Charles R. Plott, "Alternative Methods of Allocating Airport Slots: Performance and Evaluation," prepared for Civil Aeronautics Board and Federal Aviation Administration, Polinomics Research Laboratories, Inc., Pasadena (1979).

[11] David M. Grether, R. Mark Isaac, and Charles R. Plott, "The Allocation of Landing Rights by Unanimity Among Competitors," *American Economic Review*, vol. 71(2) (May, 1981) pp. 166-171.

[12] S. J. Rassenti, Vernon L. Smith, and R. L. Bulfin, "A Combinatorial Auction Mechanism for Airport Time Slot Allocation," *The Bell Journal of Economics*, vol.13(2), (Autumn 1982), 402-417.

[13] See http://clarke.pair.com/ANCW.html for "Airport Congestion and Noise: Interplay of Allocation and Distribution," 1995.

[14] David M. Grether and Charles R. Plott, "The Effects of Market Practices in Oligopolistic Markets: An Experimental Examination of the *Ethyl* Case," *Economic Inquiry*, *24*, (1984) pp. 479-507.

[15] Charles A. Holt and David Scheffman, "Facilitating Practices: The Effects of Advance Notice and Best-Price Policies," *RAND Journal of Economics*, *18*(2) (Summer, 1987) pp. 187-197.

[16] William Breit and Kenneth Elzinga, *The Antitrust Casebook* (New York: Dryden Press, 1996), pp. 85-93.

[17] In particular, the simultaneous auction process that was used in the NO_x experiments involved a linear-programming procedure to calculate the revenue-maximizing combination of bids across the two markets. This allowed subjects to submit quantity orders for each market. The total quantity could not increase from round to round, but some units could be switched from one market to another, subject to some constraints, which could permit participants to take advantage of arbitrage opportunities that might arise from the fact that 2004 permits could also be banked and used in 2005 under some conditions. The simultaneous auction setup, however, meant that bidders would not know in advance what bids would be accepted at each stage, and any *ex post* challenges to the procedures due to misunderstandings might remove one of the main advantages of auctions: speed and finality.

[18] A Georgia farmer did sue without success, claiming that the auction officials had discriminated against him and others by excluded them on the basis of incorrect data about water flows and acreage. "Farmer's Irrigation Lawsuit Dismissed," *Southeast Farm Press*, January 16, 2002: http://southeastfarmpress.com/mag/farming_farmers_irrigation_lawsuit/

[19] "An Experimental Comparison of the Simultaneous Multi-Round Auction and the CRA Combinatorial Auction," Cybernomics Consulting Report to the FCC (March 15, 2000).