Mechanism Design Experiments

Yan Chen John O. Ledyard

October 18, 2006

Abstract
Mechanism design experiments bridge the gap between a theoretical mechanism and an actual economic process. In the domain of public goods, matching and combinatorial auctions, laboratory experiments identify features of mechanisms which lead to good performance when implemented among boundedly rational agents. These features include dynamic stability and security in public goods mechanisms, transparency in matching mechanisms, package bidding, simultaneity and iteration in combinatorial auctions.

Mechanism Design Experiments
Mechanism design is the art of designing institutions that align individual incentives with overall social goals. Mechanism design theory was initiated by Hurwicz (1972) and is surveyed in Groves and Ledyard (1986). To bridge the gap between a theoretical mechanism and an actual economic process that solves fundamental social problems, it is important to observe and evaluate the performance of the mechanism in the context of actual decision problems faced by real people with real incentives. These situations can be created and carefully controlled in a laboratory. A mechanism design experiment takes a theoretical mechanism, re-creates it in a simple environment in a laboratory with human subjects as economic agents, observes the behavior of human subjects under the mechanism, and assesses its performance relative to what it was created to do and relative to the theory upon which its creation rests. The laboratory serves as a wind-tunnel for new mechanisms, providing evidence which one can use to eliminate fragile ones, and to identify the characteristics of successful ones.

When a mechanism is put to test in a laboratory, behavioral assumptions made in theory are most seriously challenged. Theory assumes perfectly rational agents who can compute the equilibrium strategies via introspection. When a mechanism is implemented among boundedly rational agents, however, characteristics peripheral to theoretical implementations, such as transparency, complexity and dynamic stability, become important, or even central, to the success of a mechanism in a laboratory, and we suspect, ultimately in the real world. Mechanism design experiments cover several major domains, including public goods and externalities, matching, contract theory, auctions, market design and information markets. In what follows, we will review experimental results of some of these topics.

Public Goods and Externalities

With the presence of public goods and externalities, competitive equilibria are not Pareto optimal. This is often referred to as market failure, since competitive markets on their own either result in underprovision of public goods (i.e. the free-rider problem) or overprovision of negative externalities, such as pollution. To solve the free-rider problem in public goods economies, incentive-compatible mechanisms use innovative tax-subsidy schemes which utilize agents' own messages to achieve the Pareto optimal levels of public goods provision. A series of experiments test these mechanisms in the laboratory (see Chen (forthcoming) for a comprehensive survey).

When preferences are quasi-linear, the Vickrey-Clarke-Groves (VCG) mechanism (Vickrey 1961, Clarke 1971, Groves 1973, and Groves and Loeb 1975) is strategy-proof, in the sense that reporting one's preferences truthfully is always a dominant strategy. It has also been shown that any strategy-proof mechanism selecting an efficient public decision at every profile must be of this type (Green and Laffont (1977)). Two forms of the VCG mechanism have been tested in the field and laboratory by various groups of researchers. The Pivot mechanism refers to the VCG mechanism when the public project choice is binary, while the cVCG mechanism refers to the VCG mechanism when the level of the public good is selected from a continuum. Under the Pivot mechanism, misrevelation can be prevalent. Attiyeh, Franciosi and Isaac (2000) show that about ten percent of the bids were truthfully revealing their values. Furthermore, there was no convergence tendency towards value revelation. In a follow-up study, Kawagoe and Mori (2001) show that more information about the payoff structure helps reduce the degree of misrevelation. More recently, Cason, Saijo, Sjorstrom and Yamato (2006) provide a novel explanation for the problem of misrevelation in strategy-proof mechanisms. As Saijo *et al* (2003) point out, the standard strategy-proofness concept in implementation theory has serious drawbacks, i.e., almost all strategy-proof mechanisms have a continuum of Nash equilibria. They propose a new implementation concept, secure implementation, which requires the set of dominant strategy equilibria and the set of Nash equilibria to coincide. Cason *et al* (2006) compare the performance of two strategy-proof mechanisms in the laboratory: the Pivot mechanism where implementation is not secure and truthful preference revelation is a weakly dominant strategy, and the cVCG mechanism with single-peaked preferences where implementation is secure. Results indicate that subjects play dominant strategies significantly more often in the secure cVCG mechanism (81%) than in the non-secure Pivot mechanism (50%). The importance of secure implementation in dominant strategy implementation is replicated in Healy (2006), where he compares five public goods mechanisms, Voluntary Contributions, Proportional Taxation, Groves-Ledyard, Walker, and cVCG. The cVCG is found to be the most efficient of all mechanisms.

Although the VCG mechanism admits dominant strategies, the allocation is not fully Pareto-efficient. In fact, it is impossible to design a mechanism for making collective allocation decisions, which is informationally decentralized, non-manipulable and Pareto optimal.

This impossibility has been demonstrated in the work of Hurwicz (1975), Green and Laffont (1977), Roberts (1979), Walker (1980) and Mailath and Postlewaite (1990) in the context of resource allocation with public goods.

Many “next-best” mechanisms preserve Pareto optimality at the cost of non-manipulability, some of which preserve “some degree” of non-manipulability. Some mechanisms have been discovered which have the property that Nash equilibria are Pareto optimal. These can be found in the work of Groves and Ledyard (1977), Hurwicz (1979), Walker (1981), Tian (1989), Kim (1993), Peleg (1996), Falkinger (1996) and Chen (2002). Other implementation concepts include perfect Nash equilibrium (Bagnoli and Lipman (1989)), undominated Nash equilibrium (Jackson and Moulin (1991)), subgame perfect equilibrium (Varian (1994)), strong equilibrium (Corchon and Wilkie (1996)), and the core (Kaneko (1977)), etc. Apart from the above non-Bayesian mechanisms, Ledyard and Palfrey (1994) propose a class of Bayesian Nash mechanisms for public goods provision.

Experiments on Nash-efficient public goods mechanisms underscore the importance of dynamic stability, i.e., whether a mechanism converges under various learning dynamics. Most of the experimental studies of Nash-efficient mechanisms focus on the Groves-Ledyard mechanism (Smith (1979), Harstad and Marrese (1981, 1982), Mori (1989), Chen and Plott (1996), Arifovic and Ledyard (2006). Chen and Tang (1998) also compare the Walker mechanism with the Groves-Ledyard mechanism. Falkinger, Fehr, Gächter and Winter-Ebmer (2000) study the Falkinger mechanism. Healy (2006) compares Nash-efficient mechanisms to cVCG and other benchmarks.

Among the series of experiments exploring dynamic stability, Chen and Plott (1996) first assessed the performance of the Groves-Ledyard mechanism under different punishment parameters. They found that by varying the punishment parameter the dynamics and stability changed dramatically. For a large enough parameter, the system converged to its stage game Nash equilibrium very quickly and remained stable; while under a small parameter, the system did not converge to its stage game Nash equilibrium. This finding was replicated by Chen and Tang (1998) with more independent sessions and a longer time series in an experiment designed to study the learning dynamics.

[Figure 1 about here]

Figure 1 presents the time series data from Chen and Tang (1998) for two out of five types of players. Each graph presents the mean (the black dots) and standard deviation (the error bars) for each of the two different types averaged over seven independent sessions for each mechanism - the Walker mechanism, the Groves-Ledyard mechanism under a low punishment parameter (GL1), and the Groves-Ledyard mechanism under a high punishment parameter (GL100). From these graphs, it is apparent that GL100 converged very quickly to its stage game Nash equilibrium and remained stable, while the same mechanism did not converge under a low punishment parameter; the Walker mechanism did not converge to its stage game Nash equilibrium either.

Because of its good dynamic properties, GL100 had significantly better performance than GL1 and Walker, evaluated in terms of system efficiency, close to Pareto optimal level of public goods provision, less violations of individual rationality constraints and convergence to its stage game equilibrium.

These past experiments serendipitously studied supermodular mechanisms. Two recent studies systematically vary the parameters from below, close to, at and above the supermodualarity threshold to assess the effects of supermodularity on learning dynamics.

Arifovic and Ledyard (2006) conduct computer simulations of an individual learning model in the context of a class of the Groves-Ledyard mechanisms. They vary the punishment parameter systematically, from extremely small to extremely high. They find that their model converges to Nash equilibrium for all values of the punishment parameter. However, the speed of convergence does depend on the value of the parameter. As shown in Figure 2, the speed of convergence is U-shaped: very low and very high values of the punishment parameter require long periods for convergence, while a range of intermediate values requires the minimum time. In fact, the optimal punishment parameter identified in the simulation is much lower than the supermudularity threshold. Predictions of the computation model is validated by experimental data with human subjects.

[Figure 2 about here]

In a parallel research project on the role of supermodularity on convergence, Chen and Gazzale (2004) experimentally study the generalized version of the compensation mechanism (Varian 1994), which implements efficient allocations as subgame-perfect equilibria for economic environments involving externalities and public goods. The basic idea is that each player offers to compensate the other for the "costs" incurred by making the efficient choice. They systematically vary the free parameter from below, close to, at and beyond the threshold of supermodularity to assess the effects of supermodularity on the performance of the mechanism. They have three main findings. First, in terms of proportion of equilibrium play and efficiency, they find that supermodular and "near supermodular" mechanisms perform significantly better than those far below the threshold. This finding is consistent with previous experimental findings. Second, they find that from a little below the threshold to the threshold, the improvement in performance is statistically insignificant. This implies that the performance of "near supermodular" mechanisms, such as the Falkinger mechanism, ought to be comparable to supermodular mechanisms. Therefore, the mechanism designer need not be overly concerned with setting parameters that are firmly above the supermodular threshold - close is just as good. This enlarges the set of robustly stable mechanisms. The third finding concerns the selection of mechanisms within the class of supermodular mechanisms. Again, theory is silent on this issue. Chen and Gazzale find that within the class of supermodular mechanisms, increasing the parameter far beyond the threshold does not significantly improve the performance of the mechanism. Furthermore, increasing another free parameter, which is not related to whether or not the mechanism is supermodular, does improve convergence.

In contrast to the previous stream of work which identifies supermodularity as a robust sufficient condition for convergence, Healy (2006) develops a k-period average best response learning model and calibrates this new learning model on the data set to study the learning dynamics. He shows that subject behavior is well approximated by a model in which agents best respond to the average strategy choices over the last five periods under all mechanisms. Healy's work bridges the behavioral hypotheses that have existed separately in dominant strategy and Nash-efficient mechanism experiments.

In sum, experiments testing public goods mechanisms show that dominant strategy mechanisms should also be secure, while Nash implementation mechanisms should satisfy dynamic stability, if any mechanism is to be considered for application in the real world in a repeated interaction setting.

While experimental research demonstrates that incentive-compatible public goods mechanisms can be effective in inducing efficient levels of public goods provision, almost all of the mechanisms rely on monetary transfers, which limits the scope of implementation of these mechanisms in the real world. In many interesting real world settings, such as open source software development and online communities, sizable contributions to public goods are made without the use of monetary incentives. We next review a related social psychology literature, which studies contribution to public goods without the use of monetary incentives.

Social Loafing

Analogous to free-riding, social loafing refers to the phenomenon that individuals exert less effort on a collective task than they do on a comparable individual task. To determine conditions under which individuals do or do not engage in social loafing, social psychologists have developed and tested various theoretical accounts for social loafing. Karau and Williams (1993) present a review of this literature and develop a collective effort model, which integrates elements of expectancy-value, social identity and self-validation theories, to explain social loafing. A meta-analysis of 78 studies show that social loafing is robust across studies. Consistent with the prediction of the model, several variables are found to moderate social loafing. The following factors are of particular interests to a mechanism designer.

- Evaluation potential: Harkins (1987) and others show that social loafing can be reduced or sometimes eliminated when a participant's contribution is identifiable and evaluable. In a related public goods experiment, Andreoni and Petrie (2003) find a substantial increase (59%) in contribution to public goods compared to the baseline of a typical VCM experiment, when both the amount of individual contribution and the (photo) identification of donors are revealed.

- Task valence: The Collective Effort Model predicts that the individual tendency to engage in social loafing decreases as task valence (or perceived meaningfulness) increases.

- Group valence and group-level comparison standards: Social identity theory (Tajfel and Turner 1986) suggests that "individuals gain positive self-identity through the accomplishments of the groups to which they belong." (Karau and Williams 1993). Therefore, enhancing group cohesiveness or group identity might reduce or eliminate social loafing. In a closely related economics experiment, Eckel and Grossman (2005) use induced group identity to study the effects of varying strength of identity on cooperative behavior in a repeated public goods game. They find that while cooperation is unaffected by simple and artificial group identity, actions designed to enhance group identity contribute to higher levels of cooperation. This stream of research suggests that high degrees of group identification may limit individual shirking and free-riding in environments with a public good.

- Expectation of co-worker performance influences individual effort. This set of theories might be sensitive to individual valuations for the public good as well as the public goods production functions. The meta-analysis indicates that individuals loafed when they expected their co-workers to perform well, but did not loaf otherwise.

- Uniqueness of individual inputs: individuals loafed when they believed that their inputs were redundant, but did not loaf when they believe that their individual inputs to the collective product were unique. In an interesting application, Beenen et al (2004) conducted a field experiment in an online community called MovieLens. They find that users who were reminded of the uniqueness of their contributions rated significantly more movies than the control group.

- Task complexity: individuals were more likely to loaf on simple tasks, but less likely on complex tasks. This finding might be related to increased interests when solving complex tasks.

Exploring non-monetary incentives to increase contribution to public goods is an important and promising direction for future research. Mathematical models of social psychology theories are likely to shed insights on the necessary and sufficient conditions for a reduction or even elimination of social loafing.

Matching

Matching theory has been credited as "one of the outstanding successful stories of the theory of games" (Auman 1990). It has been used to understand existing markets and to guide the design of new markets or allocation mechanisms in a variety of real world contexts. Matching experiments serve two purposes: to test new matching algorithms in the laboratory before implementing them in the real world, and to understanding how existing institutions evolved. We focus on one-sided matching experiments, and refer the reader to Niederle, Roth and Sonmez [in this volume] for a summary of the two-sided matching experiments.

One-sided matching is the assignment of indivisible items to agents without a medium of exchange, such as money. Examples include the assignment of college students to dormitory rooms and public housing units, the assignment of offices and tasks to individuals, the assignment of students to public schools, the allocation of course seats to students (mostly in business schools and law schools), and timeshare exchange. The key mechanisms in this class of problems are the top trading cycles (TTC) mechanism (Shapley and Scarf 1974), the Gale-Shapley deferred acceptance mechanism (Gale and Shapley 1964), and variants of the serial dictatorship mechanism (Abdulkadiroglu and Sonmez 1998). Matching experiments explore several issues. For strategy-proof mechanisms, they explore the extent to which subjects recognize and use their dominant strategies without prompting. For mechanisms which are not strategy-proof, they explore the extent of preference manipulation and the resulting efficiency loss. As a result, they examine the robustness of theoretical efficiency comparisons when the mechanisms are implemented among boundedly rational subjects and across different environments.

For the class of house allocation problems, two mechanisms have been compared and tested in the lab. The random serial dictatorship with squatting rights (RSD) is used by many U.S. universities for on-campus housing allocation, while the top trading cycles mechanism is theoretically superior. Chen and Sonmez (2002) report the first experimental study of these two mechanisms. They find that TTC is significantly more efficient than RSD because it induces significantly higher participation rate of existing tenants.

Another application of one-sided matching is the time share problem. Wang and Krishna (forthcoming) study the top trading cycles chains and spacebank mechanism (TTCCS), and two status quo mechanisms in the timeshare industry, i.e., the deposit first mechanism and the request first mechanism, neither of which is efficient. In the experiment, the observed efficiency of TTCCS is significantly higher than that of the deposit first mechanism, which in turn, is more efficient than the request first mechanism. In fact, efficiency under TTCCS converged to 100% quickly, while the other two mechanisms do not show any increase in efficiency over time.

More recently, the school choice problem has received much attention. We review two experimental studies. Chen and Sonmez (2006) present an experimental study of three school choice mechanisms. The Boston mechanism is influential in practice, while the Gale-Shapley and Top Trading Cycles mechanisms have superior theoretical properties. Consistent with theory, this study indicates a high preference manipulation rate under Boston. As a result, efficiency under Boston is significantly lower than that of the two competing mechanisms in the designed environment. However, contrary to theory, Gale-Shapley outperforms Top Trading Cycles and generates the highest efficiency. The main reason is that a much higher proportion of subjects did not realize that truth-telling was a dominant strategy under TTC, and thus manipulated their preferences and ended up worse off. While Chen and Sonmez (2006) examine these mechanisms under partial information, where an agent only knows his own preference ranking, and not those of other agents, a follow-up study by Pais and Pinter (2006) investigates the same three mechanisms under different information conditions, ranging from complete ignorance about the other participants' preferences and school priorities to complete information on all elements of the game. They show that information condition has a significant effect on the rate of truthful preference revelation. In particular, having no information results in a significantly higher proportion of truthtelling than under any treatment with additional information. Interestingly, there is no significant difference in the efficiency between partial and full information treatments. Unlike Chen and Sonmez (2006), in this experiment, TTC outperforms in terms of efficiency. Furthermore, TTC is also less sensitive to the amount of information in the environment.

Due to their important applications in the real world, one-sided matching experiments provide insights on the actual manipulability of the matching mechanisms which are valuable in their real world implementations. Some issues, such as the role of information on the performance of the mechanisms, remains open questions.

Combinatorial Auctions

In many applications of mechanism design, theory is not yet up to the task of identifying the optimal design or even comparing alternative designs. One case in which this has been true is in the design of auctions to sell collections of heterogeneous items with value complementarities which occur when the value of a combination of items can be higher than the sum of the values for separate items. Value complementarities arise naturally in many contexts, such as broadcast spectrum rights auctioned by the Federal Communications Commission, pollution emissions allowances for consecutive years bought and sold under the RECLAIM program of the South Coast Air Quality Management District in Los Angeles, aircraft takeoff and landing slots, logistics services, and advertising time slots. Because individuals may want to express bids for combinations of the items for sale, requiring up to 2N bids per person when there are N items, these auctions have come to be known as Combinatorial Auctions.

As was discussed earlier under public goods mechanisms, theory has identified the Vickrey-Clarke-Groves mechanism as the unique auction design that would implement an efficient allocation assuming bidders use dominant strategies. Theory has not yet identified the revenue maximizing combinatorial auction, although we do know that it is not the VCG mechanism. Theory has also been of little use in comparing the expected revenue collection between different auction designs. This has opened the way for many significantly different auction designs to be proposed, and sometimes even deployed, with little evidence to back up various claims of superiority.

To give some idea of the complexity of the problem we describe just some of the various design choices one can make. Should the auctions be run as a sealed bid or should some kind of iterative procedure be used? And, if the latter, should iteration be synchronous or asynchronous? What kinds of bids should be allowed? Proposals in for allowable bids include only bids for a single item, bids for any package, and some which allow only a limited list of packages to be bid on. What stopping rule should be used? Proposals have included fixed stopping times, stop after an iteration in which revenue does not increase by more than x percent, stop if demand is less than or equal to supply, and an imaginative but complex system of eligibility and activity rules created for the FCC auctions. Should winners pay what they bid or something else? Alternatives to pay what you bid include VCG prices and second best prices based on the dual variables to the program which picks the provisional winners. What should bidders be told during the auction? Some designs provide information on all bids and provisional winners and the full identity of the bidders involved in them. Some designs provide minimal information such as only the winning bids without even the information on who made them. The permutations and combinations are many. Because theory has not developed enough to sort out what is best, experiments have been used to provide some evidence.

The very first experimental analysis of a combinatoric auction can be found in Rassenti et al. (1982) where they compared a sealed bid auction (RSB) allowing package bids to a uniform price sealed bid auction (GIP), proposed by Grether, et al. (1981), that did not allow package bids. Both designs included a double auction market for re-trading after the auction results were known. The RSB design yielded higher efficiencies than the GIP design. Banks et. al. (1989) compared a continuous, asynchronous design (AUSM) a generalization of the English auction - with package bidding to a synchronous iterative design with myopic VCG pricing and found AUSM to yield higher efficiencies and revenues on average. Ledyard et. al. (1997) compare the continuous AUSM to a synchronous iterative design (SMR) developed by Milgrom (2000) for the FCC auctions which only allowed simultaneous single item bids. The testing found that ASUM yielded significantly higher efficiencies and revenues. Kwasnica et. al. (2005) compare an iterative design (RAD) with package bidding and price feedback to both AUSM and SMR. RAD and SMR use the same stopping rule. Efficiencies observed with RAD and AUSM are similar and higher than those for SMR, but revenue is higher in SMR since many bidders lose money due to a phenomenon known as the exposure problem identified in Bykowsky et. al. (2000). If it is assumed that bidders default on bids on which they make losses and thus set the prices of such bids to zero, revenues are in fact higher under AUSM and RAD than under SMR. At the behest of the FCC Banks et. al.(2003) ran an experiment to compare an iterative, package bidding design (CRA) from Charles River Associates (1998) with the FCC SMR auction format. They also found that the package bidding design provides more efficient allocations but less revenue due to bidder losses in the SMR.

Parkes and Unger (2002) proposed an ascending price, generalized VCG auction (iBEA) that maintains non-linear and non-anonymous prices on packages and charges VCG prices to the winners. The design would theoretically produce efficient allocations as long as bidders bid in a straightforward manner. Straightforward bidding is myopic and non-strategic and involves bidding on packages that yield the locally highest payoff in utility. There is no evidence that actual bidders will actually behave this way. Chen and Takeuchi (2005) have experimentally tested iBEA against the VCG sealed bid auction and found that VCG was superior in both revenue generation and efficiency attained. Lin et. al. (2006) tested RAD against VCG and found that RAD generated higher efficiencies, especially in the earlier auctions. They were using experiments to test combinatoric auctions as a potential alternative to scheduling processes in situations with valuation complementarities. In many cases current procedures request orderings from users and then employ a knapsack algorithm of some kind to choose good allocations without any concern for incentive compatibility. Lin et. al. (2006) find that both RAD and VCG yield higher efficiencies than the knapsack approach. Ledyard et. al. (1996) found similar results when comparing a more vanilla combinatoric auction to an administrative approach. These findings suggest there are significant improvements in organization performance being overlooked by management.

Porter et. al. (2003) proposed and tested a Combinatorial Clock (CC) auction. After bids are submitted, a simple algorithm determines the demand for each item by each bidder and for those items that have more than one bidder demanding more units than are available the clock price is raised. They test their design against the SMR and CRA. They do not report revenue but in their tests the CC design attained an almost perfect average efficiency of 99.9%. CRA attained an average of 93%, while SMR attained only 88%. Brunner et. al. (2006) have carried out a systematic comparison of SMR and three alternatives, CC, RAD, and a new FCC design called SMRPB, which takes the basic RAD design and changes two things. SMRPB allows bidders to win at most one package and the pricing feedback rule includes some inertia that RAD does not. They find that in terms of efficiency RAD is better than CC which is equivalent to SMRPB which is better than SMR. In terms of revenue, they find CC is better than RAD which is better than SMRPB which is better than SMR.

Most of these papers compare only two or three auction designs at a time and the environments used as the basis for comparison is often different in different papers. Further, environments can often be chosen which works in favor of one auction over another. To deal with this, many research teams do stress test their results by looking at boundary environments collections of payoff parameters that give each auction under examination its best or worst chance of yielding high revenue or efficiency. But it is still unusual for a research team to report on a comparative test of several auctions in which their design ends up being out-performed by another. Nevertheless, there are some tentative conclusions one can draw from this research.

The easiest and most obvious conclusion is that allowing package bidding improves both efficiency and revenue. In all the studies listed, anything that limits bidders ability to express the full extent of their willingness to pay for all packages will interfere with efficiency and revenue. Less obvious but also easy to see is that simultaneity and iteration are also good design features. Bidding in situations in which value complementarities exist can be difficult since bidders need to discover where their willingness to pay is more than others but also where they fit with others interests. Getting this right improves efficiency and revenue. Iteration and relevant price-feedback both help here. Stopping rules also matter. Although this is an area that could benefit from more research, it is pretty clear that in many cases, complicated stopping rules that allow auctions to proceed for very long periods of time provide little gain in revenue or efficiency.

Summary

Mechanism design experiments identify features of mechanisms which lead to good performance when they are implemented among real people. Experiments testing public goods mechanisms show that dominant strategy mechanisms should also be secure, while Nash-efficient mechanisms should satisfy dynamic stability if it is to be considered for application in the real world in a repeated interaction setting. For matching mechanisms, transparency of the dominant strategy leads to better performance in the laboratory. Lastly, in combinatorial auctions, package bidding, simultaneity and iteration are shown to be good design features. In addition to the three domains covered in this article, there have been a growing experimental literature on market design, information markets and contract theory. We do not cover them in this article, due to lack of robust empirical regularities. However, they are excellent areas to make a new contribution.

Yan Chen and John O. Ledyard

References

Andreoni, J. and R. Petrie (2004). Public goods experiments without confidentiality: a glimpse into fund-raising. Journal of Public Economics 88: 1605 1623.

Arifovic, J. and J. Ledyard (2006). Computer Testbeds and Mechanism Design: Application to the Class of Groves-Ledyard Mechanisms for Provision of Public Goods, Manuscript, Caltech.

Attiyeh, G., R. Franciosi and M. Isaac (2000), Experiments with the Pivot Process for Providing Public Goods, Public Choice.

Bagnoli, M. and B. Lipman (1989), Provision of Public Goods: Fully Implementing the Core Through Private Contributions, Review of Economic Studies 56: 583-602.

Banks, J. S., Ledyard, J. O. and Porter, D. P. (1989) Allocating Uncertain and Unresponsive Resources: An Experimental Approach, Rand Journal of Economics, 20(1), 1-25.

Banks, J., M. Olson, D. Porter, S. Rassenti and V. Smith (2003) Theory, Experiment and the Federal Communications Commission Spectrum Auctions, Journal of Economic Behavior and Organization, 51, 303-350.

Beenen, G., K. Ling, X. Wang, K. Chang, D. Frankowski, P. Resnick, and R. Kraut (2004). In Proceedings of ACM CSCW 2004 Conference on Computer Supported Cooperative Work, Chicago, IL. pp. 212-221.

Brunner, C., J. Goeree, C. Holt, and J. Ledyard (2006) Combinatorial Auctioneering, Caltech working paper.

Bykowsky, M., R. Cull, and J. Ledyard (2000) Mutually Destructive Bidding: The FCC Auction Design Problem, Journal of Regulatory Economics, 17(3): pp. 205-228, May 2000.

Charles River and Associates Inc. and Market Design Inc. (1998) Report No. 1351-00.

Chen, Y. (2002), A Family of Supermodular Nash Mechanisms Implementing Lindahl Allocations, Economic Theory 19:773-790.

Chen, Y. (forthcoming), Incentive-Compatible Mechanisms for Pure Public Goods: A Survey of Experimental Literature. The Handbook of Experimental Economics Results edited by Charles Plott and Vernon Smith.

Chen, Y. and R. Gazzale (2004), Supermodularity and Convergence: An Experimental Study of the Compensation Mechanism, American Economic Review 94(5): 1505-1535.

Chen, Y. and C. R. Plott (1996), The Groves-Ledyard Mechanism: An Experimental Study of Institutional Design, Journal of Public Economics 59: 335-364.

Chen, Y. and T. Sonmez (2002), Improving Efficiency of On-Campus Housing: An Experimental Study. American Economic Review 92 (5): 1669-86.

Chen, Y. and T. Sonmez (2006), School Choice: An Experimental Study. Journal of Economic Theory 127: 202-231.

Chen, Y. and F.-F. Tang (1998), Learning and Incentive Compatible Mechanisms for Public Goods Provision: An Experimental Study, Journal of Political Economy 106-3: 633-662.

Chen, Y. and K. Takeuchi (2005) Multi-object Auctions with Package Bidding: An Experimental Comparison of Vickrey and iBEA, working paper.

Eckel, C. and P. Grossman (2005). Managing diversity by creating team identity, Journal of Economic Behavior & Organization, 58(3): 371-392.

Falkinger, J. (1996), Efficient Private Provision of Public Goods by Rewarding Deviations from Average, Journal of Public Economics 62:413-422.

Falkinger, J., E. Fehr, S. Gchter and R. Winter-Ebmer (2000), A Simple Mechanism for the Efficient Provision of Public Goods - Experimental Evidence, American Economic Review 90:247-264.

Green, J. and J.-J. Laffont (1977), Characterization of Satisfactory Mechanisms for the Revelation of the Preferences for Public Goods, Econometrica 45:427-438.

Grether, D., M. Isaac, and C. Plott (1981) The allocation of landing rights by unanimith among competitiors, American Economic Review 71, 166-171

Groves, T. and J. Ledyard (1977), Optimal Allocation of Public Goods: A Solution to the 'Free Rider' Problem, Econometrica 45(4):783-809.

Groves, T. and J. Ledyard (1987), Incentive Compatibility since 1972, in: T. Groves and R. Radner, eds., Essays in Honor of Leonid Hurwicz (University of Minnesota Press, Minneapolis).

Harstad, R. M., and M. Marrese (1981), Implementation of Mechanism by Processes: Public Good Allocation Experiments, Journal of Economic Behavior and Organization 2:129 -151.

Harstad, R. M., and M. Marrese (1982), Behavioral Explanations of Efficient Public Good Allocations, Journal of Public Economics 19:367 - 383.

Healy, P. J. (2006), Learning Dynamics for Mechanism Design: An Experimental Comparison of Public Goods Mechanisms. Journal of Economic Theory 129: 114-149

Hurwicz, L. (1972), On Informationally Decentralized Systems, in: C. McGuire and R. Radner, eds., Decision and Organization. (North Holland, Amsterdam) 297-336.

Hurwicz, L. (1979), Outcome Functions Yielding Walrasian and Lindahl Allocations at Nash Equilibrium Points, Review of Economic Studies 46:217-225.

Isaac, R. and D. James (2000) Robustness of the incentive compatible combinatorial auction, Experimental Economics, 3(1), 31-53

Kawagoe, T. and T. Mori (1998a), Can Pivotal Mechanism Induce Truth-telling? An Experimental Study, Discussion papers in Economics No. 233, Nagoya City University.

Kawagoe, T. and T. Mori (1998b), A Short Report on Pivotal Mechanism Experiment, mimeo, Nagoya City University.

Kim, T. (1986), On the Nonexistence of a Stable Nash Mechanism implementing Lindahl Allocations, University of Minnesota mimeo.

Kim, T. (1993), "A Stable Nash Mechanism Implementing Lindahl Allocations for Quasi-linear Environments", Journal of Mathematical Economics 22:359-371.

Kwasnica, A. M., J. O. Ledyard, D. Porter, and C. DeMartini (2005) A New and Improved Design for Multi-Object Iterative Auctions, Management Science, 51(3), March, 419-434.

Ledyard, J., M. Olson, D. Porter, J, Swanson, and D. Torma (2002) The First Use of a Combined Value Auction for Transportation Services, Interfaces, 32(5):4-12,

Ledyard, J. and T. Palfrey (1994). Voting and Lottery Drafts as Efficient Public Goods Mechanisms. The Review of Economic Studies, 61 (2): 327-355.

Ledyard, J., D. Porter, and C. Noussair (1996) The Allocation of a Shared Resource within an Organization, Economic Design, Vol. 2, No. 2, November 1996, pp. 163-192.

Ledyard, J., D. Porter, and A. Rangel (1997) Experiments Testing Multiobject Allocation Mechanisms, Journal of Economics and Management Strategy, Vol. 6(3), 639-675.

Lin, J., K. Takeuchi, Y. Chen, T. Finholt (2006) Facilitation Large-scale Collaboration Through Effective Mechanism Design: An Economic Examination of Resource Akllocation in Collaboratories, working paper

Mailath, G. and A. Postlewaite (1990), Asymmetric Information Bargaining Problems with Many Agents, Review of Economic Studies 57:351-367.

McAfee, P. R. and McMillan, J. (1996) Analyzing the Airwaves Auction, Journal of Economic Perspectives, 10(1), 159-175.

McCabe, K., Rassenti, S. and Smith, V. (1989) Designing 'Smart' Computer Assisted Markets, European Journal of Political Economy, 5, 259-283.

Millgrom, P. (2000) Putting Auction Theory to work: the simultaneous ascending auction. Journal of Political Economy, 108, 245-272a

Milgrom, P., and J. Roberts (1990), Rationalizability, Learning and Equilibrium in Games with Strategic Complementarities, Econometrica 58(6):1255-1277.

Milgrom, P., and C. Shannon (1994), Monotone Comparative Statics, Econometrica 62(1):157-180.

Mori, T. (1989), Effectiveness of Mechanisms for Public Goods Provision: An Experimental Study, Economic Studies 40(3):234-246.

Peleg, B. (1996), Double Implementation of the Lindahl Equilibrium by a Continuous Mechanism, Economic Design 2:311-324.

Porter, D. P. (1999) The Effect of Bid Withdrawal in a Multi-Object Auction, Review of Economic Design, 4(1), 73-97.

Porter, D., S. Rassenti, A. Roopnarine and V. Smith (2003) Combinatorial Auction Design, Proceedings of the National Academy of Sciences, 100(19), September, 11153-11157.

Rassenti, S., Smith, V. and Bulfin, R. (1982) A Combinatorial Auction Mechanism for Airport Time Slot Allocation, Bell Journal of Economics, 13, 402-417.

Roberts, J. (1979), Incentives and Planning Procedures for the Provision of Public Goods, Review of Economic Studies 46:283-292.

Smith, V. (1979a), Incentive Compatible Experimental Processes For the Provision of Public Goods, in: Smith Research, ed. Experimental Economics vol. 1 (JAI Press Inc, Greenwich, Connecticut).

Smith, V. (1979b), An experimental comparison of three public goods decision mechanisms, Scandinavian Journal of Economics 81:198-251.

Tian, G. (1989), Implementation of the Lindahl Correspondence by a Single-Valued, Feasible, and Continuous Mechanism, Review of Economic Studies 56:613-621.

Walker, M. (1980), On the Impossibility of a Dominant Strategy Mechanism to Optimally Decide Public Questions, Econometrica 48:1521-1540.

Walker, M. (1981), A Simple Incentive Compatible Scheme for Attaining Lindahl Allocations, Econometrica 49:65-71.

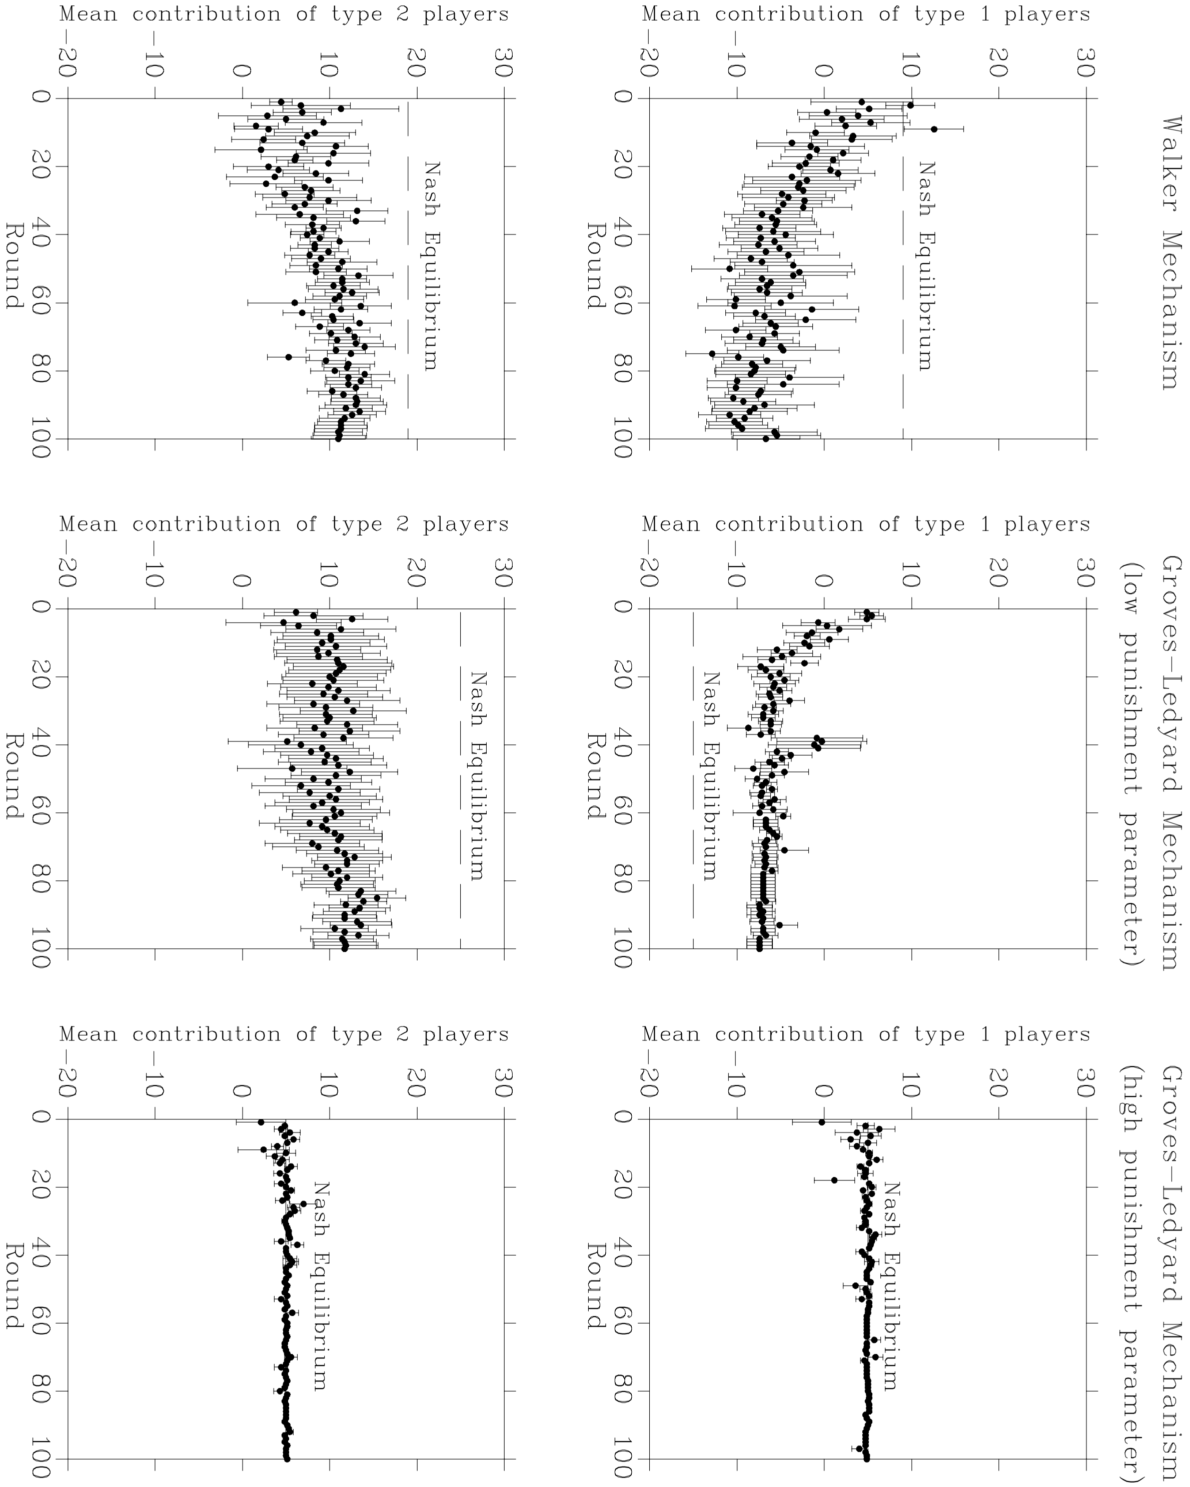

Figure 1: Mean Contribution and Standard Deviation in Chen and Tang (1998)

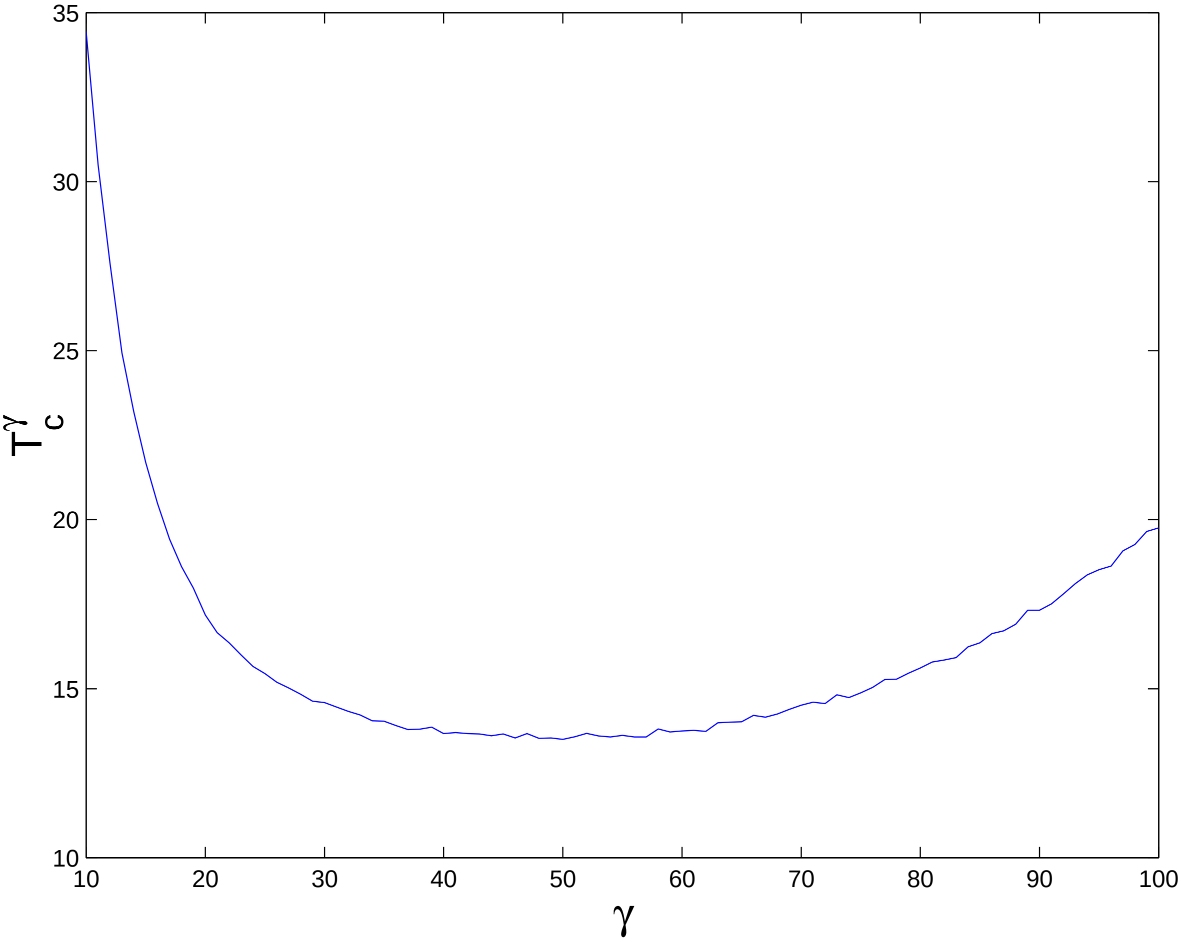

Figure 2: Convergence Speed in Groves-Ledyard in Arifovic and Ledyard (2006)