

Ideal Bootstrapping and Exact Recombination: Applications to Auction Experiments

Carl Bergstrom

University of Washington, Seattle, WA
and Santa Fe Institute, Santa Fe, NM,

Ted Bergstrom

University of California, Santa Barbara,

Rod Garratt*

University of California, Santa Barbara

Current Version: July 16, 2009

Previous Version: January 7, 2008

Abstract

We provide simple formulas that can be used to calculate ideal bootstrap or exact recombination estimates of group statistics from experimental data. When resampling is done over small numbers of groups, the ideal bootstrap is more accurate than the exact recombinant estimate, however the former is biased. For large sample sizes there is no discernable difference between the two approaches and both produce estimates that are more accurate than those obtained from observed group outcomes.

1 Introduction

A researcher conducts auction experiments in s separate sessions. Each session includes n subjects who bid against each other for an object. The experi-

*Corresponding author. Department of Economics, University of California, Santa Barbara, CA 93106, (805) 893-2849, garratt@econ.ucsb.edu

menter randomly induces valuations for each subject, by making independent draws from a specified distribution.

Subjects are informed of their own valuations, the number of bidders in their group and the rules of the auction. For each group, a single auction is conducted, the bid of each subject is recorded and net payoffs are awarded according to the rules of the auction.

How can the results of this experiment be used to predict the probability distribution of the outcome if the experiment were to be repeated? Suppose, for example, that the researcher would like to predict the seller's revenue from this type of auction. A simple approach would be to observe the revenue collected in each of the s groups, calculate the mean and standard deviation of revenue across these groups, and make statistical predictions based on the assumption that the subjects in future experiments would be drawn from a normal distribution with mean and standard deviation equal to those found for the s groups that were sampled.

This approach has two serious drawbacks. One is that there is no reason to believe that the probability distribution of revenue from randomly constructed auctions will be normally distributed. Even if the joint distribution of values and bids is normal, auction revenue is based on order statistics from this distribution and the order statistics of a normal distribution are not normally distributed. More importantly, to record only the results of the groups that actually formed to determine individual payoffs is to discard a great deal of useful information. An experimental auction is a one-shot game in which no player communicates with others in his group before making his own bid. Although an individual's payoff depends on the bids of others in his group, his own bid would be no different if he were assigned to any other group. We can improve our estimates by examining the distribution of outcomes that would result from random reassignments of subjects to groups.

We would like to use the experimental results to answer statistical questions such as "Suppose that we were to repeat the experiment with another auction in which n bidders are randomly drawn from the same population as were the original participants: What is the probability distribution of revenue from the group chosen in this draw? What is the probability distribution of the "efficiency" of the auction in this new draw?"

Mullin and Reiley (2005) suggest a procedure that they call "recombinant estimation." This procedure is to estimate the distribution of outcomes that would be generated by constructing a new group of n subjects drawn randomly *without replacement* from the set of $ns = T$ bid-value combinations

observed in the experiment. The recombinant approach was also employed earlier by Mitzkewitz and Nagel (1993) to estimate the probability distribution of profits in an ultimatum game and by Mehta *et al* (1994) in a matching game.

A related approach that is more familiar to statisticians is “bootstrap estimation ” (Efron and Tibshirani (1993)). The bootstrap approach estimates the distribution of outcomes that would be generated by constructing a new group of n subjects drawn randomly *withreplacement* from the set of T bid-value combinations observed in the original experiment.

For the two-player games studied by Mitzkewitz and Nagel and by Mehta *et al*, exact calculation of the probability distribution of outcomes from random recombination of players is straightforward and the authors present such calculations. Reiley and Mullin (2005) consider a problem where larger groups are selected and where brute force calculation of probability distributions did not appear practical. They propose estimating the recombinant distribution by Monte Carlo simulations. Similarly, most applications of the bootstrap methods involve numerical simulations. Efron and Tibshirani explain that it is possible in principle to state exact probability distributions for bootstrap problems. Such probabilities are known as “ideal bootstrap” probabilities.

This paper shows that for bidding experiments, simple tricks make it possible to calculate exact probability distributions of the results of resampling, either by the bootstrap or the recombination method. This method works even where the number of distinct partitions that can be obtained by resampling is extremely large.¹ We evaluate the alternative approaches by using them to compute expected revenue in simulated second-price auction experiments with known bidding distributions. The mean squared deviation between the true expected selling price and the estimated selling price is lower under the ideal bootstrap than the recombinant method when the sample size (i.e., the number of auctions that are combined to compute the estimates) is small. However, the mean estimated selling price obtained by the bootstrap method is slightly lower than the true mean for small sample sizes (the bootstrap estimate is biased) whereas the recombinant mean is correct. Hence, there is a tradeoff between the two approaches for small sample sizes.

¹For example, $T = 100$ subjects can be partitioned into groups of 5 bidders in more than seventy-five million possible ways.

For large sample sizes the bias in the bootstrap estimate disappears and the mean squared deviation between the true expected selling price and the estimated selling price obtained by the two approaches converges. Both the ideal bootstrap and the recombinant method predict the true selling price better than observed group outcomes.

2 Applications to Auctions

2.1 Preliminaries

Suppose that the experimenter has data on bids and values from s groups, each of which has n subjects. Array these T bids in ascending order to construct a list b . In general, the items in the list b will not all be distinct, since more than one bidder may submit the same bid. To handle ties, we construct a second list b' , which consists of all *distinct* bids, arrayed in ascending order. Let T' be the number of elements of the list b' . For each $i = 1, \dots, T'$, define $L_i(b)$ to be the number of elements of b that are no larger than the i th element of b' .²

2.2 Revenue with First-Price Sealed Bid Auctions

In a first-price sealed bid auction, each subject submits a single bid, without observing the bids of others. An object is sold to the highest bidder in each group at a price equal to the high bidder's bid.

2.2.1 The ideal bootstrap estimate

Calculating the ideal bootstrap distribution is strikingly simple if all bids are distinct (no ties). Suppose that each group has n bidders, then where b_i is i th element of the vector of bids arrayed from smallest to largest, the bootstrap estimate of the probability that b_i is the largest bid in a group of size n is just

$$\left(\frac{i}{T}\right)^n - \left(\frac{i-1}{T}\right)^n. \quad (1)$$

²If the elements of b are all distinct, $L_i(b) = i$ for each bid i , and the results presented below are familiar properties of order statistics; see Evans et al, 2006.

When not all bids are distinct, the computation is only slightly more complicated. For each b'_i in the list b' of distinct bids, the probability that b'_i is the highest bid in a randomly selected group of n individuals drawn with replacement is equal to the probability that all n draws are no larger than b'_i , minus the probability that all draws are no larger than b'_{i-1} . This probability is

$$p^F(b'_i) = \left(\frac{L_i(b)}{T}\right)^n - \left(\frac{L_{i-1}(b)}{T}\right)^n. \quad (2)$$

Expected revenue is simply

$$\sum_{i=1}^{T'} b'_i p^F(b'_i). \quad (3)$$

Direct calculation of the other moments of the distribution of revenue is also straightforward.

2.2.2 The exact recombinant estimate

Exact recombinant estimates differ from ideal bootstrap methods only in that random groups are constructed by resampling *without replacement* from the set of T bids submitted by subjects. If groups of size n are chosen without replacement, the lowest possible top bid in a group is b'_n . For $i \geq n$, the number of groups of size n in which b'_i is the highest bid is $\binom{L_i(b)-1}{n-1}$. Therefore the probability that b'_i is the winning bid in a randomly selected group of size n drawn without replacement is

$$\tilde{p}^F(b'_i) = \frac{\binom{L_i(b)-1}{n-1}}{\binom{T}{n}} \text{ if } L_i(b) \geq n \text{ and } \tilde{p}^F(b'_i) = 0 \text{ if } L_i(b) < n. \quad (4)$$

Expected revenue is simply $\sum_{i=n}^{T'} b'_i \tilde{p}^F(b'_i)$.

2.3 Revenue with Second-Price Auctions

In a second price auction, the sale item goes to the high bidder at a price equal to the second highest bid. The probability distribution of revenue is simply the probability distribution of the second highest bid in a randomly selected group of n subjects.

2.3.1 The ideal bootstrap estimate

For each b'_i in the list b' of distinct bids, define $P^S(b'_i)$ to be the probability that the second highest bid is no larger than b'_i . Thus $P^S(b'_i)$ is the probability that no more than one bid larger than b'_i is selected in a sample of size n drawn with replacement from the list b' . The probability that a single draw will be less than or equal to b'_i is $L_i(b)/T$. Therefore

$$P^S(b'_i) = \left(\frac{L_i(b)}{T}\right)^n + n \left(\frac{T - L_i(b)}{T}\right) \left(\frac{L_i(b)}{T}\right)^{n-1}. \quad (5)$$

For each b'_i in the list b' , the probability that the second highest bid is exactly b'_i is the difference between the probability that the second highest bid is less than b_i and the probability that the second highest bid is less than b_{i-1} . Therefore the probability that the second highest bid is exactly b'_i is

$$p^S(b'_i) = P^S(b'_i) - P^S(b'_{i-1}). \quad (6)$$

We have thus produced an estimate of the full probability distribution of revenue and we can readily calculate the mean or any other moment of this distribution. In this case, expected revenue is $\sum_{i=1}^{T'} b'_i p^S(b'_i)$.

2.3.2 The exact recombinant estimate

The recombinant approach is to find the probability distribution of second-highest bids if each of the $\binom{T}{n}$ groups of n bids that could be selected from the populations of bids were equally likely.

A bid b'_i in the list b' will be the second-highest bid in a group of n bidders if there are $n - 2$ bids less than or equal to b'_i and one other bid at least as large as b'_i . Thus the total number of groups of size n in which b'_i is the second-highest bid is

$$\binom{L_i(b) - 1}{n - 2} \times (T - L_i(b))$$

provided $L_i(b) \geq n - 1$ and 0 otherwise. Since the number of distinct groups of size n that can be formed is $\binom{T}{n}$, then where $L_i(b) \geq n$, the probability that b'_i is the second highest bid is

$$\tilde{p}^S(b'_i) = \frac{\binom{L_i(b)-1}{n-2} \times (T - L_i(b))}{\binom{T}{n}} \quad (7)$$

and for $L_i(b) < n$, $\tilde{p}^S(b'_i) = 0$.

3 Efficiency

Auction theorists are interested in the “efficiency” of auctions, where efficiency is measured by the ratio of realized total profits for buyers and sellers to the maximum potential total profits. To estimate the efficiency of auctions, we need to look at bidders’ valuations (which have been induced by the experimenter) as well as their bids. We will show how to calculate this measure by means of the ideal bootstrap procedure. Similar calculations can be made for exact recombinant estimates.

Suppose that an experiment has generated T observations of bid-value pairs, (b_j, v_j) . We first wish to find the expected surplus yielded by a perfectly efficient outcome. This value will be equal to the expected revenue in a first price auction where every bidder j bids his true value v_j . This expected revenue is given by Equation 3, where we replace the list b of bids by a list v of bidders’ values and the list b' of distinct bids by the list v' of distinct values arrayed in ascending order. The expected highest valuation is

$$\sum_{i=1}^{T''} v'_i p^F(v'_i), \quad (8)$$

where T'' is the number of distinct values in the list v' .

In second price auctions, as well as first price auctions, the object is sold to the highest bidder. Thus in either case, we need to compute the expected value of the object to the highest bidder. For either type of auction, the probability $p^F(b'_i)$ that b'_i is the winning bid in a first price auction is given by Equation 3. Let us define \bar{v}_i to be the mean of the valuations, v_j , of those subjects j whose bids are b'_i in a first price auction. Thus if b'_i is the winning bid, the expected value of the object to the buyer, is \bar{v}_i . Thus for either a first price or a second price auction, the expected value of the object to the winning bidder is therefore

$$\sum_{i=1}^{T'} \bar{v}_i p^F(b'_i). \quad (9)$$

Let us measure the efficiency of an auction as the ratio of expected value of the object to the winning bidder to the expected value of the object to the

bidder with highest value. Thus we have

$$E^F = \frac{\sum_{i=1}^{T''} \bar{v}_i p^F(\bar{v}'_i)}{\sum_{i=1}^{T'} v'_i p^F(b'_i)}. \quad (10)$$

4 Hypothesis Testing

Suppose a seller is planning to sell an object to a population that he believes is very similar to the sample population used to generate our bootstrap estimates of expected revenue. He wishes to know which auction format is most likely to produce the most revenue. The bootstrap procedure can also be used to put confidence intervals around the estimates of expected revenue. Let R^A denote the expected revenue of auction format A , $A \in \{F, S\}$. The bootstrap estimate of the standard deviation of this estimate is $\mathcal{S}^A = \sqrt{\sum_{i=1}^{T'} p^A(b'_i)(b'_i - R^A)^2}$. This can be used to compute confidence intervals around the revenue estimates.³

5 Examples

5.1 A lazy experimenter's auction

To clarify the difference between naive estimation, recombinant estimation, and bootstrap estimation of auction results it is instructive to consider their workings in a very simple class of examples.

A researcher is interested in estimating the expected revenue that a supplier would raise by selling an object in a two-person, second bidder auction. The researcher does not know the distribution of willingness to pay in the population at large and decides to estimate these returns by experimental methods. To reduce our own computational burden, let us assume that this is very lazy researcher who chooses a sample of just two persons and sells to one of them by means of a sealed-bid second price auction.

The lazy experimenter selects her two subjects at random from a large population, half of which values the object at \$1 and half of which values the object at \$0. Let us assume that both subjects play their weakly dominant

³These confidence intervals will not be accurate if, as may well be the case, the distribution of the winning bids is far from normal.

strategy of bidding their true valuation. A seller will receive revenue of \$1 if both bidders have \$1 values and will receive \$0 otherwise. Therefore the true expected value of revenue in a two-person auction with randomly chosen buyers is $1/4$. How well the alternative estimation procedures perform?

Since the population contains equal proportions of each type, the probability is $1/4$ that both subjects value the object at \$1, $1/2$ that one subject values it at \$1 and the other values it at \$0, and $1/4$ that both value the object at \$0. If the investigator follows the naive procedure of using observed revenue to estimate expected revenue, then with probability $3/4$, the estimate will be 0 and with probability $1/4$ the estimate will be 1. Since the entire sample has only two bidders, recombination by sampling without replacement does not produce any new combinations of bidders and so the recombinant estimate must be the same as the naive estimate. The naive procedure and the recombinant procedure are both “unbiased” in the sense that the expected value of the estimate $3/4 \times 0 + 1/4 \times 1 = 1/4$ is equal to the true expected value of revenue. On the other hand, the investigator’s estimate is never exactly right. The mean squared error of this estimate is $3/4(1/4)^2 + 1/4(3/4)^2 = 3/16$.

The bootstrap procedure draws new samples *with replacement* from the selected subject pool. If both persons in the original sample are of the same type, bootstrap resampling produces no new combinations and yields the same estimate as the naive and recombinant methods. But if one person sampled has value \$1 and the other value \$0, the bootstrap estimates the probability of drawing two \$0 types to be $1/4$, the probability of drawing one person of each type to be $1/2$ and the probability of drawing two \$1 types to be $1/4$. In this case she estimates expected revenue to be $1/4$. Since the bootstrap estimate is \$1 with probability $1/4$ and $1/4$ with probability $1/2$, the expected value of the bootstrap estimate of expected revenue is $3/8$. The bootstrap estimate is therefore “biased” since the true expected revenue from two randomly selected bidders is $1/4$. Although the bootstrap estimator is “biased” and the naive and recombinant estimators are “unbiased”, the bootstrap estimate is strictly “better than” the naive and recombinant estimates in the following sense. Whenever the two bidders drawn are of the same type, the bootstrap estimate is the same as the naive and recombinant estimates, but whenever the two bidders drawn are of two different types, the bootstrap estimate is exactly correct and the other two procedures underestimate expected revenue. The mean squared error of the bootstrap procedure is $1/4(3/4)^2 + 1/4(1/4)^2 = 1/16$.

Consider a slightly more diligent experimenter, who recruits two pairs of two subjects and runs a sealed-bid second-price auction with each pair. With a bit of calculation, we find that if she uses naive estimation, her estimate of expected revenue will be \$0 with probability $9/16$, $\$1/2$ with probability $3/8$, and $\$1$ with probability $1/16$. This estimate for revenue of $\$1/4$ is unbiased and has mean squared error of $9/16(1/4)^2 + 3/8(1/4)^2 + 1/16(3/4)^2 = 9/32$. With recombinant estimation, she will estimate expected revenue to be 0 with probability $5/16$, $1/6$ with probability $3/8$, $1/2$ with probability $1/4$ and 1 with probability $1/16$. This estimate, which is also equal to $\$1/4$, is also unbiased and has mean squared error of $5/16(1/4)^2 + 3/8(1/12)^2 + 1/4(1/4)^2 + 1/16(3/4)^2 = 7/96$. With bootstrap estimation, she will estimate expected revenue to be 0 with probability $1/16$, $1/16$ with probability $1/4$, $1/4$ with probability $3/16$, $9/16$ with probability $1/4$ and 1 with probability $1/16$. The expected value is $5/16$, so the bootstrap estimate is biased upward. However the mean squared error is $1/16(1/4)^2 + 1/4(3/16)^2 + 1/4(3/16)^2 + 1/4(5/16)^2 + 1/16(3/4)^2 = 37/1024$, which is lower than both the naive and recombinant estimate. Hence the bootstrap estimate is more accurate.

Does it make an important difference to compute the ideal bootstrap or exact recombinant instead of using the naive approach of simply forming each auction once and looking at the results, and if it does matter, which should approach is better, the ideal bootstrap or the exact recombinant? To answer these questions we consider five-player second price auctions for the case where we know the underlying distribution of bids to be normal with mean 10 and standard deviation 2. For this case, we can work out the actual "true" expected selling price to be 10.99 with standard deviation of 1.12. We want to see how well each of these estimation procedures: naive formation of auctions using each point only once, ideal bootstrap, and ideal recombinant, perform as a function of the sample size (total number of five player auctions that we run in our experiment). We do that by simulating a very large number of experiments (100,000) for each sample size.

Figure 1 shows the mean squared deviation between the true expected selling price and the estimated expected selling price as a function of sample size. The dotted line shows us what we get for the naive procedure. The dashed line shows us what we get for the exact recombinant. The solid line shows us what we get for the ideal bootstrap (The figure is plotted on a log scale). As expected, mean squared deviation drops with sample size. We also see that the ideal bootstrap is strictly better than the other procedures, but that as the sample size gets large the ideal recombinant does approximately

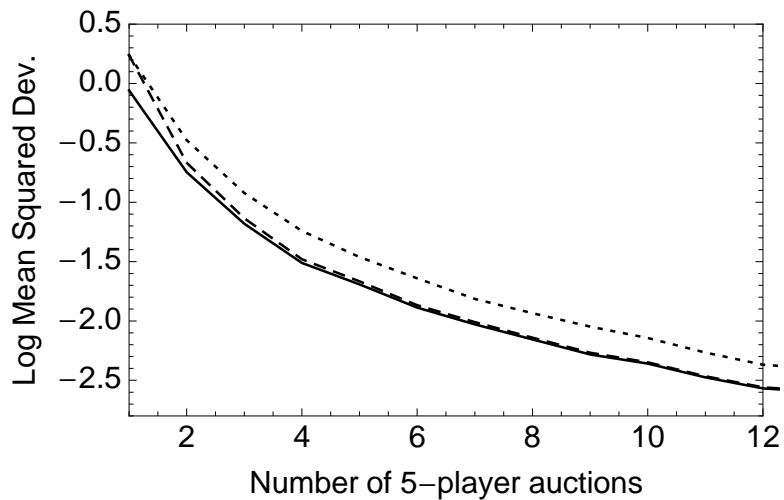


Figure 1: Comparison.

as well. So this figure shows us that it does make a difference whether we use some sort of recombinant sampling or bootstrap, and hints that for small sample sizes the bootstrap might be best. For large sample sizes, it doesn't matter which of the two one uses.

Why wouldn't you use the bootstrap? A comparison of figures 2-4 suggests one possible answer. The solid lines in these figures show the mean value across all 1000 experiments of the estimated mean selling price. In the naive and recombinant procedures, the mean estimated mean selling price is right on target for all sample sizes — that is, these estimators are unbiased. In the bootstrap procedure, the mean estimated selling price is somewhat lower than the true mean for small sample sizes — that is, the bootstrap estimator is biased. For practical purposes, where the experimenter is running some specified number of sessions once (rather than replicating the whole experiment thousands of times) the issue of bias seems less important than accuracy. Hence, the advantage of the lower mean squared deviation for the ideal bootstrap probably outweighs the unbiased nature of the exact recombinant. The dashed lines in figures 2-4 show the ranges in which 97.5% of the estimated means lie. The ranges are tighter for the ideal bootstrap and broader for the naive approach.

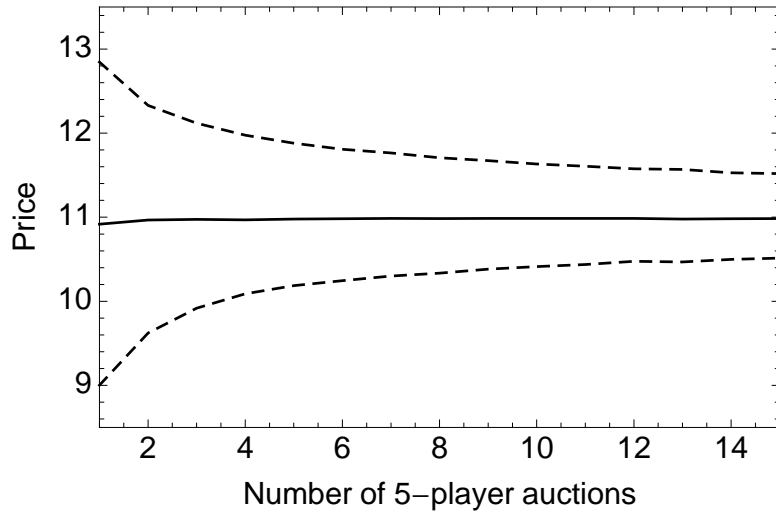


Figure 2: Ideal bootstrap.

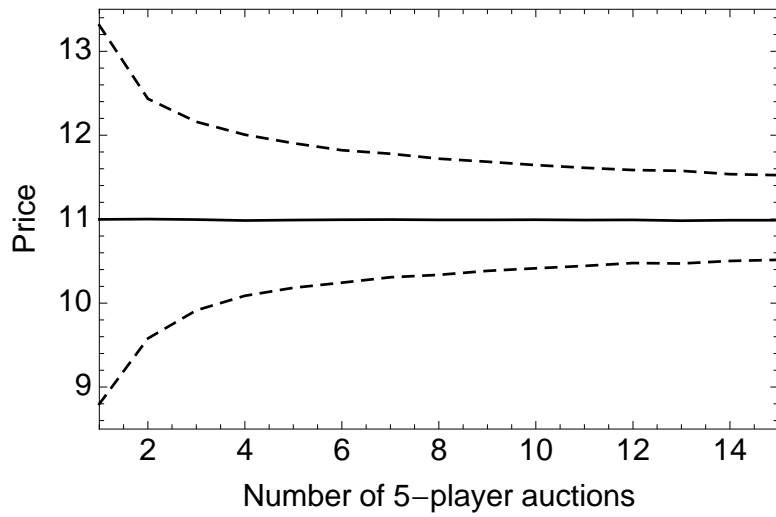


Figure 3: Recombinant sampling.

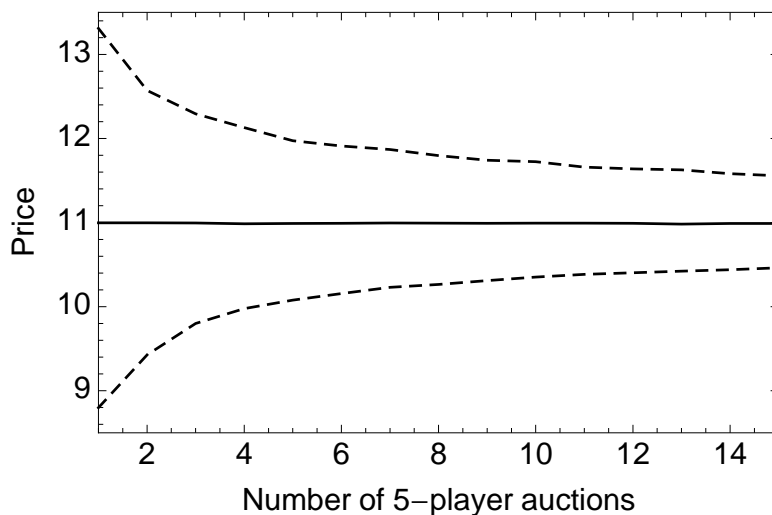


Figure 4: Naive pairing.

6 Other applications

Other types of behavior can also be analyzed using the above techniques. For instance, in a war of attrition a group of players compete to be the unique survivor and receive a prize. The expenditures by each player are equivalent to bids in an auction and the award goes to the highest bidder. Other interactions such as lobbying, political campaigns, lawsuits, standing in line for tickets, and some forms of price-setting oligopoly can be modelled as auctions (See Klemperer, 2000), and hence experimental treatments of these interactions produce bid-value vectors which can be analyzed using recombinant or bootstrap methods. For example, Dufwenberg et al (2006) conduct an experiment in which subjects choose prices in a Bertrand oligopoly game. In this case the “winner” is the player who chooses the lowest price. The ideal bootstrap and exact recombinant estimates of expected price in these markets can be obtained simply by replacing the $L(\cdot)$ function in Equations 2 and 4 with the function $G(\cdot) : B' \rightarrow Z_+$, defined so that for any $b'_i \in B'$, $G(b'_i) = |\{b_j \in B : b_j \geq b'_i\}|$.

References

- [1] M. Dufwenberg, U. Gneezy, J. K. Goeree, and R. Nagel, "Price floors and competition." *Economic Theory*, forthcoming.
- [2] B. Efron and R. Tibshirani, 1993, "An Introduction to the Bootstrap." Chapman and Hall, New York.
- [3] D. Evans, L. Leemis, and J. Drew, 2006, "The Distribution of Order Statistics for Discrete Random Variables with Applications to Bootstrapping." *INFORMS Journal on Computing*, 18, 19-30.
- [4] P. Klemperer, "Applying Auction Theory to Economics." Department of Economics Discussion Paper, Number 1, April 2000, Oxford.
- [5] J. Mehta, C. Starmer, and R. Sugden, 1994, "The nature of salience: An experimental investigation of pure coordination games." *American Economic Review*, 84, 658-73.
- [6] M. Mitzkewitz and R. Nagel, 1993, "Experimental results on ultimatum games with incomplete information." *International Journal of Game Theory*, 22, 171-98.
- [7] C. H. Mullin and D. H. Reiley, 2006, "Recombinant Estimation for Normal-Form Games, with Applications to Auctions and Bargaining." *Games and Economic Behavior*, 54, 159-182.