Ideal Bootstrapping and Exact Recombination

Ted Bergstrom and Rod Garratt UCSB

February 7, 2007

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 experimenter 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 revenue from this experiment. 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 outcome 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 formed.

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 use only the revenue results of the groups actually formed is to discard a great deal of useful information. The experimental auction is a one-shot game in which no player communicates with others in his group before making his own bid. Thus each individual's behavior depends only on his own characteristics and not on those of the other members of his group.

An experiment can be viewed as the result of two distinct random processes. In the first of these processes, T different bid-value combinations are drawn from an unknown joint distribution of bids and values. In the second process, the T bid-value combinations are partitioned into s groups of size n.

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 there are n bidders are randomly drawn from the same population as that of 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?

To answer such questions, we need to specify the way in which a new sample would be drawn. Two reasonable alternatives present themselves. In the "bootstrap approach," (Efron and Tibshirani (1993)) one draws n bid-value combinations *with replacement* from a population with the empirical distribution of bid-value combinations observed in the original experiment.

A variant of the bootstrap approach, "recombinant estimation," is advocated by Mullin and Reiley (2005). Recombinant estimation assumes that if the experiment were done again, the new participants would be drawn randomly *without replacement* from the set of 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.

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 Tibsharani 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 resampling is extremely large.[1]

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 $i = 1, \ldots, T'$, define $L_i(b)$ to be the number of elements of b that are no larger than the ith element of b'_i. (In the special case where the elements of b are all distinct, then $L_i(b') = i$ for each bid i.)

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

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

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

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

Expected revenue is simply

$$\sum_{i=1}^{T'} b'_i p^F(b'_i). \tag{2}$$

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

$$\tilde{p}^F(b'_i) = \frac{\binom{L_i(b)-1}{n-1}}{\binom{T}{n}} \text{ if } L_i(b) \geq n \text{ and } \tilde{p}^F(b'_i) = 0 \text{ if } L(b'_i) < n. \tag{3}$$

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

$$P^S(b'_i) = \left(\frac{L_i(b)}{T}\right)^n + n\left(\frac{T - L_i(b)}{T}\right)\left(\frac{L_i(b)}{T}\right)^{n-1}. \quad (4)$$

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

$$p^S(b'_i) = P^S(b'_i) - P^S(b'_{i-1}). \quad (5)$$

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

$$\tilde{p}^S(b'_i) = \frac{\binom{L_i(b)-1}{n-2} \times (T - L(b'_i))}{\binom{T}{n}} \quad (6)$$

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.

3.1 Expected highest valuation

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 2, 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

$$\sum_{i=1}^{T''} v'_i p^F(v'_i), \tag{7}$$

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

3.2 Calculating expected surplus

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 2. 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

$$\sum_{i=1}^{T'} \bar{v}_i p^F(b'_i). \tag{8}$$

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

$$E^F = \frac{\sum_{i=1}^{T''} \bar{v}_i p^F(\bar{v}_i')}{\sum_{i=1}^{T'} v_i' p^F(b_i')}. \tag{9}$$

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.[2]

5 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 (1) and (**??**) with the function $G(\cdot) : B' \to Z_+$, defined so that for any $b_i' \in B'$, $G(b_i') = |\{b_j \in B b_j \geq b_i'\}|$.

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

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] R. Garratt, M. Walker, J. Wooders, "Experienced Bidders in On-line Second-Price Auctions." (2005), unpublished paper.

[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] D. H. Reiley and C. Mullin, "Recombinant Estimation for Normal-Form Games, with Applications to Auctions and Bargaining." (January 2002), forthcoming in Games and Economic Behavior.