ECONOMIC RESEARCH REPORTS

CAN AFFIRMATIVE ACTION BE COST EFFECTIVE? AN EXPERIMENTAL EXAMINATION OF PRICE-PREFERENCE AUCTIONS

by **Allan Corns**
and
Andrew Schotter

RR # 96-02 **January 1996**

C.V. STARR CENTER FOR APPLIED ECONOMICS

NEW YORK UNIVERSITY
FACULTY OF ARTS AND SCIENCE
DEPARTMENT OF ECONOMICS
WASHINGTON SQUARE
NEW YORK, NY 10003-6687

Can Affirmative Action Be Cost Effective? An Experimental Examination of Price-Preference Auctions

Allan Corns and Andrew Schotter

July 23, 1996 (Revised since RR#96-02 published January 1996)

ABSTRACT

One of the most controversial questions facing our economy and society in recent years is the question of affirmative action. The debate centers around an issue of minority representation in which the rights of the majority are depicted as being compromised by affirmative action programs set up to increase the welfare of minorities. Implicit in this discussion is the assumption that all affirmative action programs must be cost increasing. This result, it is claimed, follows trivially from economic theory since any interference into the competitive process which prevents the most capable from being chosen must be wasteful and costly.

In this paper we challenge this implicit result by demonstrating that price-preference auctions, in which high cost minority firms are given preferential treatment in the awarding of contracts, can be programs that both enhance minority representation and are cost effective in that they decrease the cost of government procurement. In this paper we ask: If a government agency has instituted a price-preference policy in an effort to give preferential treatment to a particular subset of contractors, what effect would this price-preference policy have on the cost of government purchasing? In particular, is it possible that such a policy might not only benefit that subset but actually lower the purchasing cost of the government as well? We answer this question in the affirmative both on theoretical and empirical (experimental) grounds[1].

.

Keywords: Affirmative Action, Auctions, Experiments
JEL Classification: C92, D44

[1]The authors would like to express their thanks to Huagang Li for his computational assistance. We would also like to thank J.P. Benoit, and John Riley for sharing his BIDCOMP[2] program with us. In addition, the research assistance of Jia Lu Yin, Gautam Barua, and Jeff Davis is gratefully acknowledged as is the financial assistance of the C.V. Starr Center for Applied Economics at New York University.

1. Introduction

One of the most controversial questions facing our economy and society in recent years is the question of affirmative action. The debate centers around an issue of minority representation in which the rights of the majority are depicted as being compromised by affirmative action programs set up to increase the welfare of minorities. Implicit in this discussion is the assumption that all affirmative action programs must be cost increasing. This result, it is claimed, follows trivially from economic theory since any interference into the competitive process which prevents the most capable from being chosen must be wasteful and costly.

In this paper we challenge this implicit result by demonstrating that price-preference auctions, in which high cost minority firms are given preferential treatment in the awarding of contracts, can be programs that both enhance minority representation and are cost effective in that they decrease the cost of government procurement. In this paper we ask: If a government agency has instituted a price-preference policy in an effort to give preferential treatment to a particular subset of contractors, what effect would this price-preference policy have on the cost of government purchasing? In particular, is it possible that such a policy might not only benefit that subset but actually lower the purchasing cost of the government as well? We answer this question in the affirmative both on theoretical and empirical (experimental) grounds.

Such a result is not without precedent, however. In the theoretical literature on asymmetric auctions (see, Maskin and Riley (1995)) and the optimal auction literature (see, Myerson (1981)) it has been known for quite some time that in asymmetric auctions in which firms draw their costs from different probability distributions, it is optimal to not always award the good to the lowest cost firm. In fact, in a paper on international trade, McAfee and McMillian (1989) ask a question identical to the one investigated here and demonstrate that a price-preference can be cost effective. Our analysis follows theirs closely.

In the experimental literature, a similar question was asked previously in the context of economic tournaments by Schotter and Weigelt (1992) where it was found that in situations where historical discrimination was severe, affirmative action programs which favor high cost workers or workers who, because of previous discrimination, have a higher cost of effort than their privileged counterparts, actually increase the effort levels of both minority and majority workers and hence the output of the enterprise in which they work.

What we find in our experiments is that the imposition of a price-preference rule can lead to both an increase in minority representation and cost effectiveness if the degree of price-preference is chosen correctly. By this we mean that in our experimental auction in which a 5% price-preference was used, the cost of purchasing for the laboratory auctioneer decreased while at the same time the frequency with which high cost (minority) firms won a contract increased. For the auctions run using a 10% or 15% price preference this was not the case. In those auctions, while there was an increase in minority representation, it came at a cost of higher purchasing prices for the laboratory auctioneer. These results are in accordance with theory however, indicating that if a government purchasing agent has some accurate prior information about the distributions from which the cost of buyers are drawn, then it might be possible to choose a price-preference which will be cost effective. A note of caution must be raised, however, since choosing the wrong price-preference rule may reverse this result.

In this paper we will proceed as follows. In Section 2 we will review the need for experiments in answering the question posed above. In addition we will review the incidence of price-preference auctions among government agencies as well as the empirical work in the field. Section 3 presents the theory that underlies these experiments while in Section 4 we present our experimental design and presents the results of our experiments. Finally, in Section 5 we offer some conclusions.

2. Price-Preference Auctions – The Need for Experiments and Their Use in Practice.

2.1. Using Experiments

The main advantage of using laboratory techniques is that the laboratory offers an opportunity to analyze controlled changes in circumstances and parameters which is not available in the analysis of field data. For example, take the question asked above. In order to answer this question correctly an investigator would have to be able to discover how a bidder changed his or her behavior in response to an introduction of price-preference rules. To understand this process it would be necessary to know the true costs of bidders and then be able to also observe their bid on any contract. In addition, the contract being offered for sale would have to be held constant so as to allow the investigator to compare bids on identical projects under different sets of auction rules. In the real world little of this is possible. While bids are observable costs are not. If costs could be observed, there would be no need for an auction since the purchasing agency could simply add a reasonable profit percentage on to the cost of the lowest cost firm and award the contract at that price. In addition, the contracts let for bid are many times not identical and not bid for by the same set of contractors.

Hence, field data does not present us with the opportunity to make the type of controlled comparisons we would like and more indirect methods would have to be employed. In addition, simulations of the type performed by McAfee and McMillian(1989), while informative about the cost effectiveness of auctions with perfectly rational bidders using equilibrium bid strategies, tell us little about the way real bidders, who use rules of thumb that are less than perfectly rational, might behave. The laboratory allows us to get information about real bidders in a controlled setting.

2.2. The Incidence of Price-Preference Auctions

Price-preference auctions, of one form or another, have been a common feature of the American economy for many years. A number of surveys have been conducted by the National Institute of Governmental Purchasing (NIGP 1993) and the National Association of State Purchasing Officials (1994) whose purpose is to summarize the purchasing practices of various governmental agencies. The results indicate that 2% of the 402 Federal, State, City, and administrative bodies responding indicated that they used minority price-preferences in their procurement programs. In fact, of all minority preference programs existing about 8% are price-preference programs with the remainder split between set aside and goal based programs. Probably the most famous use of price-preferences, at least for economists, has been their use in the recent FCC auction of radio spectrum licenses. In setting up this auction the government's concern for minority representation led it to devise a preference program where minority and women owned firms ("designated entities") were given a 40% bidding preference[2]. Further, of the 50 states, 15 had some form of in-state price preference program with percentages ranging from 2% to 5%. On the Federal level there is a long history of the Buy-American purchasing preference program (starting with the Buy-America Act of 1933) with price-preferences ranging from 6% to 12%. On defense contracts, however, the price-preference can be as high as 50%. Because these price-preference programs are seen as creating substantial barriers to trade, either among the states or across nations, they have been challenged and studied a fair amount.

2.3. A Review of the Empirical Literature

Despite their common use, we are not aware of any study which investigates the procurement cost consequences of minority price-preference auctions except for the recent investigation by Ayres and Cramton (1996, forthcoming) of the FCC auction data. The reason for this lack of attention is that in all of the studies we have seen which mention price-preference programs, it has been assumed that these programs are put

[2] In authorizing these auctions, Congress required the FCC to "ensure that ... businesses owned by members of minority groups and women are given the opportunity to participate in the provision of spectrum-based services, and for such purposes, consider the use of tax certificates, bidding preferences, and other procedures." U.S.C. 309(j)(4)(D).

in place strictly for minority representation purposes with the tacit assumption that higher purchasing costs for the government will inevitably result. These programs, like minority set-aside programs, are assumed to be anti-competitive because, by offering one set of economic agents an artificial advantage, they are seen as diminishing competition. When minority preference programs are discussed, the attention tends to focus on set-aside programs and the surrounding discussion of disparity studies done to justify them. (See LaNoue (1994) for a survey of the issues relating to disparity studies and Clegg (1993) for a discussion not only of disparity studies but also of the theory of discrimination and how it relates to price-preference programs.).

While little attention has been focused on minority price-preference programs, there has been some research on the related (and theoretically identical) issue of State-Preference and "Buy-American" programs of State and Federal purchasing authorities. When studying State and Federal price-preference programs, there seems to be an implicit ad hoc methodology hinted at of simply taking all contracts awarded by an in-state preference and estimating the cost of the preference as the difference between the lowest bid and the actual price that the contract was awarded. This view of what industry seems to think of as the cost of a price-preference program is summarized by Jordan (1978):

> "The major effect of such laws is to increase the cost of government. Indeed, an obvious cost is any amount which must be paid in excess of the lowest bid from non-residents." (Jordan (1978), page 215).

This naive approach does not , however, consider the effect of the price-preference rule on bidding behavior. For example, consider a procurement auction in which one minority firm and one non-minority firm participate as partially represented by Table 2.1. The minority firm can draw either costs x or y with equal probability. Likewise, the non-minority firm can draw either costs w or z with equal probability. Assume $w < y < z < x$.

Take the case when the minority firm draws y and suppose its equilibrium bid when drawing this cost in a no-preference auction is \$94, while in a 10% preference auction, with the minority bidder receiving the preference, its equilibrium bid is \$96[3].Suppose that the non-minority firm draws a cost of z and that its equilibrium bid when drawing this cost in a no-preference auction is \$93 and in a 10% preference auction is \$90. Therefore if the purchasing authority was using the 10% preference rule, the minority firm would win the auction at a price of \$96 since $\$96 < 1.1 \times \$90$. Thus, the above mentioned methodology would claim that the cost of using the 10% preference would be \$96-\$90=\$6. This approach, however, ignores the fact that if instead the purchasing agency had implemented a no-preference auction the non-minority firm would have won the contract, but at a price of \$93, not the \$90 it bid in the 10% preference auction. So actually in this case the cost of the preference is only \$3

If we then assume that when the non-minority firm draws cost w it bids \$90 in the no-preference auction and \$86 in the 10% preference auction, when the minority firm draws cost y the non-minority firm will win in the no-preference auction with a price of \$90 and will also win in the 10% preference auction with a price of \$86. Thus, in this case, the preference rule actually saves the purchasing agency \$4. Finally, if we assume that the equilibrium bids of the minority firm when it draws cost x are such that it does not win in either auction form no matter the cost draw of the non-minority firm we can see that the imposition of the 10% preference rule would save the purchasing agency \$3 if the non-minority firm drew cost z, as it would lower its winning bid from \$93 to \$90, and again save \$4 if it drew cost w. Therefore, the actual expected cost to the government of imposing the 10% preference rule is not the $\left(\frac{1}{4}\right)\$6 = \$1.5$ as predicted by the ad hoc methodology but is rather $\left(\frac{1}{4}\right)$(\$3 - \$4 - \$3 - \$4) = -\$2, or an expected savings of \$2

[3]This example is constructed to show in the simplest possible way how price-preferences affect bidder behavior and thereby procurement costs. While space considerations prevent us from providing a full model from which such equilibrium bids might arise, the behavior of both bidders in response to the introduction of a price-preference is consistent with the type of response theory dictates.

Table 2.1

Example Auctions' Procurement Cost Comparisons

Given Minority Bidder Draws Cost y

	0% Preference	**10% Preference**
Non-Minority Draws z $\left(\Pr = \frac{1}{2}\right)$		
Minority Bid	\$94	\$96
Non-Minority Bid	\$93	\$90
Naive Cost of Preference = \$96 - \$90 = \$6		
Real Cost of Preferenece = \$96 - \$93 = \$3		
Non-Minority Draws w $\left(\Pr = \frac{1}{2}\right)$		
Minority Bid	\$94	\$96
Non-Minority Bid	\$90	\$86
Naive Cost of Preference = \$86 - \$86 = \$0		
Real Cost of Preference = \$86 - \$90 = -\$4		
Expected Cost of Preference Given Minority Draws $y = \frac{1}{2}(\$3 - \$4) = -\$.50$		

Though this example is a bit simplistic, it points out two flaws in the naive, ad hoc approach: First, when a favored high cost firm wins due to a preference, the cost of the preference is overstated. Second, this naive approach does not account at all for the fact that when the low cost firm wins a positive preference auction, it is winning with a lower bid than it would have made in a no-preference auction, thus actually saving the purchasing agency money. Most of the literature on this topic seems to ignore these effects of a price-preference rule on bid behavior. While the ad hoc approach might be an appealing assumption to make in large world markets, it is clearly inappropriate for the thin auction markets typically found for government contracts[4].

[4]Using a fixed, world-market price assumption, Lowinger(1976) does an analysis to measure the implicit tariff defined by the "Buy-America" program. He calculates an implicit tariff of between

3. Some Theory

3.1. Optimal Auction Design

From a theoretical point of view it should not be surprising that price-preference programs, in which the lowest cost producer is not awarded the contract, can reduce the cost of purchasing for the government. This is the standard result derived for asymmetric auctions. For example, the theory of optimal auction design as outlined by Myerson (1981) and applied to a setting almost identical to the one studied here by McAfee and McMillian [5] clearly indicates that under the optimal auction mechanism the lowest-cost bidder is not always awarded the good. More precisely, consider a model where there are n_1 firms of type 1 and n_2 firms of type 2. Types differ by the distributions from which they draw their costs with type 1 firms drawing their costs from a continuously differentiable distribution function G_1 and firms of type 2 drawing their costs from a similar distribution function G_2. Letting c_i^l and c_i^h, $i = 1$, 2 be the lowest and highest per unit costs respectively of firms of type i, define c_{ij}, as the cost draw of the j^{th} firm of type i, such that $c_i^l \leq c_{ij} \leq c_i^h$. Finally, define

$$J_i(c_{ij}) = c_{ij} + \frac{G_i(c_{ij})}{g_{i(c_{ij})}} \; j = 1, ..., n_i; i = 1, 2 \tag{3.1}$$

which is a scalar that can be derived by an auctioneer knowing only the distribution function $G_i(c_{ij})$ and the cost of firm i of type j. c_{ij}, however, will be revealed at the equilibrium of the optimal direct mechanism. Myerson (1981) has demonstrated the optimal direct mechanism involves bidders submitting their costs and the contract being awarded to that bidder with the lowest $J_i(c_{ij})$ at a price determined by the bid of his competitors. Note, however, that because of the differing distribution functions G_i, it is possible for the lowest cost producer, i.e. that firm with the lowest $c_{ij,}$ not to receive the contract and hence it can be optimal to discriminate against the low

26% and 43%. Of course, as stated above, in price-preference auctions such a fixed price assumption is not appropriate.

[5] For an intuitive explanation of the theory see Bulow and Roberts (1989).

cost firm[6].

McAfee and McMillian (1989) investigate this identical problem in an international trade context. For them, the two types of agents in the model are domestic and foreign firms which have different cost distributions. Optimal discrimination involves the existence of a discrimination function $z(c_1)$ which the government uses to compare the costs of domestic and foreign firms. More precisely, a domestic firm with cost c_1 will be designated a winner against a foreign firm with cost c_2 if $z(c_1) < c_2$. The price paid will be that of the low cost firm. If for some c, $z(c) < c$, then the foreign firm will be said to be discriminated against and the contract could go to the firm with the highest cost.

To construct the optimal discriminatory policy or $z(c)$ function, we need to choose $z(c)$ such that $J_1(c) = J_2(z(c))$ since this $z(c)$ gives that discrimination level consistent with the optimal auction allocation rule. In the case where both distributions are uniform, and related by $c_2^l = \psi(c_1^l)$ and $c_2^h = \psi(c_1^h)$, $0 < \psi < 1$, the optimal discriminatory rule adds a (negative) constant term to the costs of the favored class of firms (See McAfee and McMillian (1989) Corollary 5). In fact, in these circumstances the optimal policy yields the following simple rule:

$$z(c) = c + \frac{c_2^l - c_1^l}{2} \tag{3.2}$$

Note that the constant added to the cost reported by any firm in the optimal discriminatory auction is merely a function of the lower bounds of their two distributions. As we will see later using our experimental data, the region between these two values is critical for the cost of government purchasing in the face of price-preference rules. More generally $z(c) < c$ if and only if $\frac{G_2(c)}{G_1(c)}$ is strictly decreasing in c. (McAfee and McMillian (1989) Theorem 3).

[6]To implement this auction, the following procedure is used: 1) Have each bidder announce his true cost. 2) Translate each bidder's cost via his J function. 3) Award the contract to the bidder with the lowest J, call it J_1, at a price determined by taking the second lowest J, call it J_2, and paying the winner the amount $J_1^{-1}(J_2)$, which is equivalent to the cost the winner would have had to have to have $J_1 = J_2$. Note that truth telling is optimal in this auction form.

3.2. Price-Preference Auctions

While price-preference auctions in which a percentage preference is given to one type of bidder is not an optimal auction form, we study such price-preference rules in this paper because they are the empirically relevant auction form used by governments. While not optimal, it is possible, however, that by choosing the right preference percentage a government might be able to reduce its cost of purchasing[7]. The reason why this occurs is simple. Price-preferences, while accused of being anti-competitive, actually increase competition among the set of firms of which the government wants to spur competition – the low cost firms. To understand why price-preferences work, consider an auction with 2 high cost firms and 4 low cost firms. By low cost and high cost we mean that the low cost firms draw their costs from a uniform distribution with support $[c^l, c^h]$ while the high cost firms draw their costs from a uniform distribution with support $[\lambda c^l, \lambda c^h]$ where $\lambda > 1$. Note the asymmetry in the situation. While high cost firms face one other high cost firm and four low cost firms, low cost firms face three other low cost firms and two high cost firms. In other words, low cost firms face less competition than high cost firms. When price-preferences are instituted, it makes high cost firms look more like low cost firms and as a result of this increase in "effective" competition, the low cost firms bid more aggressively, i.e. they bid closer to their cost. While, by analogy, high cost firms now face less competition and hence bid less aggressively. If the preference is chosen correctly, the reduction in bids by low cost firms (the firms who are more likely to have lower costs of production) more than compensates for the increased bids of high cost firms.

To see this more formally, assume that there is a single contract to be awarded in a price-preference first price auction. Let $N = \{1, ..., n\}$ be the set of bidders for this contract. A subset containing $k < n$ of these bidders are assumed to be of type A. The rest of the bidders are of type B. Each bidder privately draws an individual-specific

[7]Especially in the uniform distribution case, price-preference rules might be quite successful in approximating both the allocation and purchasing costs of the optimal auction.

cost, c, of fulfilling the contract[8] from their own type-specific probability distribution over costs, $G_i(.), i = A, B$. $G_i(.)$ is be assumed to be a continuous, differentiable distribution, defined over the interval $[\underline{c}_i, \bar{c}_i]$, $i = A, B$, with density function $g_i(.)$. It will be assumed that $\bar{c}_A \geq \bar{c}_B$ and $\underline{c}_A \geq \underline{c}_B$, so type A bidders will be the high cost bidders. Each individual bidder knows his own type, how many bidders of each type are present, his own private cost draw, and the distributions of costs faced by all other bidders.

Given the assumptions on the supports of the cost draws for each type, if it is true that for any given cost, c, in either support, $H_B(c) > H_A(c)$ where

$$H_i(c) = \frac{g_i(c)}{1 - G_i(c)} \tag{3.3}$$

then bidders of type B have a *cost advantage* over bidders of type A. In other words, given any cost level, c, contained in either interval, there is a higher probability that a bidder of type B will draw a cost lower than c than of a bidder of type A drawing a cost lower than c.

What separates the price-preference auction from the simple asymmetric case is that bidders of type A are given a *bidding advantage* in the auction for purposes of awarding the contract. After all bids are submitted, the auctioneer adjusts all type B bids by multiplying them by one plus the amount of the preference (we will denote this sum as θ) to form the comparison bids for bidders of type B. The bids of type A remain unchanged for purposes of comparison, so for all type A bidders, the submitted bid and comparison bid will be identical. The bidder who submits the lowest *comparison* bid will then be awarded the contract at a price equal to his *submitted* bid. Therefore, the adjustment procedure only determines of the winner of the auction and does not affect the ex-post payoff of the winner.

[8]We drop subscripts for simplicity.

3.3. Equilibrium Bid Functions

To be able to find the equilibrium bid functions of the two types, one must prove that the bid functions are monotonically increasing and that they are defined over a closed and bounded interval. Once these two facts are established, it follows that the bid functions are differentiable almost everywhere on that interval. Differentiability allows us to take the first order conditions for optimization of expected profits for both types. Solving these two equations simultaneously yields the equilibrium bid functions of each type. Since the proofs of these two pre-requisites are easily available elsewhere[9] we will simply proceed by assuming monotonicity for the bid functions as well as differentiability.

3.3.1. First Order Conditions for Optimization

The problem facing a bidder i of type A is to maximize expected profits given his cost draw, c_i. This can be written as

$$\max_b (b - c_i) Pr(b \leq b_j \ \forall j \in A) Pr(b \leq \theta b_k \ \forall k \in B) \tag{3.4}$$

We do not know the distributions of bids for the types, but we do know the distributions of costs and that the bid functions are strictly monotonic. Therefore, we can define the inverse bid functions for both types as

$$y_j(b_i) = c_i, \ j = A, B, \ i \in j$$

We can rewrite the optimization problem facing a member of type A as

$$\max_b (b - c)(1 - G_A(y_A(b)))^{k-1}(1 - G_B(y_B(\frac{b}{\theta})))^{n-k} \tag{3.5a}$$

given that there are n total bidders, k of which are of type A. Analogously, we can define the problem faced by members of type B as

$$\max_b (b - c)(1 - G_A(y_A(\theta b)))^{k}(1 - G_B(y_B(b)))^{n-k-1} \tag{3.5b}$$

[9] For proof of differentiability, see, for example, Kolmogorov and Fomin (1970)

Differentiating with respect to b and after some manipulation, the first order condition for members of type A is

$$\frac{1}{b - y_A(b)} = \frac{g_A(y_A(b))y'_A(b)(k-1)}{1 - G_A(y_A(b))} + \frac{g_B(y_B(\frac{b}{\theta}))y'_B(\frac{b}{\theta})(n-k)}{\theta(1 - G_B(y_B(\frac{b}{\theta})))} \tag{3.6a}$$

Similarly, the first order condition for members of type B can be written as

$$\frac{1}{b - y_B(b)} = \frac{g_A(y_A(\theta b))y'_A(\theta b)\theta k}{1 - G_A(y_A(\theta b))} + \frac{g_B(y_B(b))y'_B(b)(n-k-1)}{1 - G_B(y_B(b)))} \tag{3.6b}$$

We can further simplify this set of first order conditions by using the $H_i(\cdot)$ functions from equation 3.3 and write

$$\frac{1}{b - y_A(b)} = H_A(y_A(b))y'_A(b)(k-1) + H_B(y_B(\frac{b}{\theta}))y'_B(\frac{b}{\theta})\frac{(n-k)}{\theta} \tag{3.7a}$$

$$\frac{1}{b - y_B(b)} = H_A(y_A(\theta b))y'_A(\theta b)(\theta k) + H_B(y_B(b))y'_B(b)(n-k-1) \tag{3.7b}$$

If we assume that both cost distributions are uniform, then

$$H_i(c) = \frac{1}{\bar{c}_i - c}, \; i = A, \; B$$

Replacing in the above system of equations gives us

$$\frac{1}{b - y_A(b)} = \frac{1}{\bar{c}_A - y_A(b)}y'_A(b)(k-1) + \frac{1}{\bar{c}_B - y_B(\frac{b}{\theta})}y'_B(\frac{b}{\theta})\frac{(n-k)}{\theta} \tag{3.8a}$$

$$\frac{1}{b - y_B(b)} = \frac{1}{\bar{c}_A - y_A(\theta b)}y'_A(\theta b)(\theta k) + \frac{1}{\bar{c}_B - y_B(b)}y'_B(b)(n-k-1) \tag{3.8b}$$

It is this system of differential equations that needs to be solved, given appropriate boundary conditions, to yield the symmetric, within-type equilibrium bid functions.

The upper boundary conditions of the system for both type A and type B firms can be determined analytically and can serve as terminal conditions to pin down the equilibrium bid functions. Let $\bar{b}_A = b(\bar{c}_A)$ and $\bar{b}_B = b(\bar{c}_B)$ be the equilibrium bids for type A and type B , respectively, when $\bar{c}_A$ and $\bar{c}_B$ are drawn. Consider two cases: Case 1, $\theta\bar{c}_B > \bar{c}_A$, and Case 2, $\theta\bar{c}_B < \bar{c}_A$. We will demonstrate that the upper bounds for the system in Case 1 will be $\hat{b}_A = \bar{b}_A = \bar{c}_A$ and $\hat{b}_B = \frac{\bar{c}_A}{\theta}$. An analogous argument, left to the reader, demonstrates that in Case 2 $\hat{b}_A = \theta\bar{c}_B$ and $\hat{b}_B = \bar{b}_B = \bar{c}_B$.

Assume that Case 1 is true, a firm of type A draws cost $\bar{c}_A$, and that in equilibrium any firm, regardless of type, drawing a cost whose associated bid has a zero probability of winning will bid its cost. Submitting a bid below $\bar{c}_A$ cannot be part of an equilibrium bid function since such a bid results in a negative expected profit. Bidding strictly above $\bar{c}_A$ also cannot be part of an equilibrium due to within-type competition between members of type A since it would result in each firm of type A having incentive to lower its maximum bid in the direction of $\bar{c}_A$, thus gaining positive expected profits. Hence, within-type competition establishes $\bar{b}_A = \bar{c}_A$ in Case 1 for type A firms. In addition for all $c_A < \bar{c}_A$, $b(c_A) > c_A$ since setting $b(c_A) = c_A$ implies zero expected profit while bidding strictly above cost can result in a positive probability of winning and hence positive expected profits. Firms of type B, recognizing the value of the maximum bid of type A, will know that any bid greater than $\frac{\bar{c}_A}{\theta}$ will lose with probability one. Hence by assumption, for all cost realizations greater than $\frac{\bar{c}_A}{\theta}$ bidders of type B will bid their cost. Thus, in Case 1 $\hat{b}_B = \frac{\bar{c}_A}{\theta}$ and $\hat{b}_A = \bar{c}_A$. Case 2 can be solved for in an analogous manner.

As with most systems of nonlinear differential equations, it is not possible to find a closed-form solution for the inverse bid functions, and thereby for the actual bid functions themselves. We used the algorithm of Riley and Li(1993) to numerically solve for the equilibrium bid functions in our auction[10].

4. Experimental Design

The experiments performed were a straight forward implementation of a price-preference auction. Students were recruited by announcement in undergraduate economics classes during the Summer and Fall of 1995. Volunteers were told to come to a classroom that was reserved for the experiment. Upon arrival six students were then randomly selected for each experimental session. Then the subjects were randomly assigned

[10]Here again let us take the opportunity to thank Huagang Li for all of his time and effort in helping us use this algorithm and for the actual time he took in working on this problem.

to be either a type A or a type B bidder. Type B bidders drew their costs in any round of the experiment from a uniform distribution with support $[100, 200]$. Type A subjects drew their costs from a uniform distribution with support $[110, 220]$. Hence, type A subjects were high cost bidders while type B subjects were low cost bidders. There were 4 type B bidders and 2 type A bidders.

After subjects read the instructions and had them read out loud by the experimental administrator, any questions the subjects had were answered and the experiment began[11]. In the beginning of each round of the experiment (there were 20 rounds in all) an experimental administrator walked around the room with two bags of chips marked A and B. In each bag was a number of chips representing uniform distributions for the integers in the supports of the two types of distributions. The bag identities were hidden from the subjects so that no one in the room knew who among them was a low or high cost type. If a subject was an A type, the administrator would give him or her the A bag and he or she would pull out a chip with a number written on it. This number would be the cost for the subject in that round.

After a cost was drawn, each subject would record that cost on his or her worksheet and then take out one of their 20 bid slips upon which they would write their bid. These bids were then collected by the experimental administrator and then depending on the rules of the auction run, a winner would be determined. The experimental administrator would then write on the blackboard the number of the subject who had won, the price he or she won at, and whether that subject was an A or B type. (Actually subjects numbered 1 or 2 were A types while those numbered 3, 4, 5,or 6 were B types). Subjects would then record their payoffs and the next round would start in an identical manner. In each experiment there were 20 rounds and the final payoff to subjects was the sum of their payoffs over the entire 20 round history of the experiment. Subjects were paid at the end of the experiment and dismissed. A post-experimental questionnaire was also administered in 13 out of the 20 experimental

[11]Instructions for the 10% preference experiment are contained in the Appendix to this paper.

sessions run. The results of this questionnaire are available upon request from the authors.

Four different experiments which differed only with respect to the rules used were run. In other words, in all experiments the number of A and B types remained the same, as did the supports of the cost distributions, etc. The preference rules used were a 0%, 5%, 10%, and 15% preference for type A. For any x% preference rule greater than 0%, after the bids were submitted the bids of the B types were increased by x% before any comparison of bids was made. The lowest post-preference bid was then awarded the contract at the price of the submitted bid and the experiment proceeded to the next round. Subjects were paid $5.00 for showing up and average payoff for the one hour experimental session was $10.03. Motivation was quite high. Table 4.1 describes our design.

Table 4.1

Experimental Design

Preference Experiment	**Number of Groups**	**Types**	**Costs**	**Number of Subjects**
0%	5	2 A-Types 4 B-Types	$c_A \in [110, 220]$ $c_B \in [100, 200]$	30
5%	5	2 A-Types 4 B-Types	$c_A \in [110, 220]$ $c_B \in [100, 200]$	30
10%	5	2 A-Types 4 B-Types	$c_A \in [110, 220]$ $c_B \in [100, 200]$	30
15%	5	2 A-Types 4 B-Types	$c_A \in [110, 220]$ $c_B \in [100, 200]$	30

Our discussion of the experiment's results will be broken into two sections, one dealing with the original policy question asked and another dealing with more theoretical issues. In our first section we will investigate the representation question of whether the imposition of a price-preference rule increased the likelihood that a high cost firm would win an auction. We will also be concerned with whether the imposition of a price-preference rule increased cost effectiveness and which price-preference rule resulted in the lowest cost of purchasing for the auctioneer (experimental administrator).. In the second section we will investigate how well the theory of auctions predicted the behavior of the subjects and how behavior changed over time as the auctions were repeated with the same set of subjects.

4.1. Auction Performance: Does Price Preference Decrease Cost and Increase Representation of High Cost Firms

In general, the institution of an appropriate price-preference rule can not only decrease the cost of government purchasing but also raise the probability that a high cost firm will win an auction. In short, if the price-preference rule is chosen correctly one can obtain increases in both cost effectiveness and representation of high cost firms. The descriptive results of the experiment are shown in Table 4.2 which we will use to discuss both the representation and procurement cost outcomes of the experiment.

Table 4.2 Here

In this table we list the number of winners of each type in the four auctions run (Winners) and the associated percentage of type A winners (Type A Win %), the theoretically predicted percentage of type A winners (Theoretical Type A Win %) for comparison, the number of type A that won because of the preference rule (Wins by Preference), the average cost drawn by winning subjects in each (Avg. Cost Winners), the average bid entered by these winners, (Avg. Bid Winners), as well as the average profit realized by these winning bidders (Avg. Profit Per Unit). We also list the average price paid by the auctioneer to purchase goods in each auction (Avg. Observed Price) and the theoretically predicted expected purchase price (Theoretical Price). Finally, as we describe later, we present two estimates of the cost of these price-preference rules. One estimate, what we call the naive estimate (Naive Cost of Preference), uses the standard methodology and simply takes the difference between the price paid by the auctioneer when the price preference is invoked and the lowest bid entered for that good and averages these costs over all good sold. The second, (Actual Cost) incorporates the fact that bidding behavior changes when we change the biding rule as described above.

4.2. Representation of High Cost Firms

Discussing representation first, we see that as the price-preference given to high cost firms increases, the fraction of contracts awarded to them increases from 12% in the 0% or no-price-preference auction to 43% in the 15% preference auction. More importantly, however, it appears that as the price-preference is increased, it becomes responsible for more and more of these wins, as it turns what would have been losing high cost bidders under lower preference regimes into winners. For example, while 10/22 or 45.4% of the type A winners in the 5% experiment won because of the price preference 63.6% of the wins for type A subjects in the 10% experiment were the result of the preference rule. The fact that the percentage falls from 63.6% to 53.4% in the 15% case is likely the result of the cost realizations occurring in the 15% experiment where, at least in one experiment, the cost realizations for the low-cost B types were particularly high. Hence there is little dispute that at least in our laboratory setting, price-preference rules increase the probability that high-cost (minority) firms win contracts.

4.3. Procurement Cost

While on the surface it might appear easy to present evidence supporting the idea that price-preference auctions increase cost effectiveness, it is actually slightly complicated. The problem is that cost realizations across experiments differ, and hence the cost of purchasing for the laboratory government might be more reflective of these cost differences than differences in the bidding behavior of subjects. For example, if subjects in the 5% price-preference auction happened to have drawn a skewed set of cost realizations that were "too high", then even if the price-preference rule led them to bid aggressively, the actual costs determined by the 5% auction might be higher than those in the 0% auction where costs were more consistent with the uniform distribution. Hence we look for a method to present our data which will purge this cost effect and standardize our comparisons. We offer several such methods.

4.3.1. Procurement Cost Comparisons

Looking at Table 4.2 we see that the average price paid per auction in our four experiments was 121.24, 119.29, 122.78, and 124.41 for the 0%, 5%, 10%, and 15% auctions respectively. Looking simply at the average price paid per auction, the 5% price-preference rule seems to be the best, followed by the 0% auction and then the 10% and 15% auctions. It is hard to rely on these findings as a solid demonstration that the cost of purchasing is lowest when the 5% preference is used since as stated above, these differences might reflect the fact that the costs drawn in this experimental auction were lower than those drawn in others and not behavioral differences in bidding. If we could have had our subjects receive identical realizations across the 5% and other auctions, it might have turned out that, given their behavior, purchasing costs might have been lower with, say, a 10% price-preference auction or, more importantly, with a 0% preference. The distribution of costs drawn are shown in Figures 4.1a-b where we present the histograms of cost realizations for A and B types in our four auctions as well as the cumulative distributions of these costs.

Figures 4.1a-b Here

As can be seen in Figures 4.1a - 4.1b, while little difference appears in the random costs drawn by type B subjects, the type A subjects in our 0% experiment seem to have drawn more costs in the region below 140 than did type A subjects in the other three experiments. When this difference is tested using a Kolmogorov-Smirnov test we see that no difference exists between the distribution of drawn costs between any two experiments for any type[12]. In other words, statistically the costs drawn came

[12] Kolmogorov-Smirnov Tests for Equality of Cost Distributions A-Types

from the same population.

Despite these statistical tests, however, it is still possible that the cost of purchasing differences observed were functions of the particular costs drawn in particular experiments and not behavioral differences. As a result, in order to make the correct comparison we asked what would the cost of purchasing be if the bidders in an auction with an $x\%$ price-preference rule had received the cost realization of bidders in an auction with a $y\%$ rule and vice versa? This is exactly the comparison we make in Table 4.3.

Table 4.3

Purchasing Cost Comparisons

Cost Realizations	Cost in 0% Auction	Cost in 5% Auction	Cost in 10% Auction	Cost in 15% Auction
0%	121.24	119.74	121.63	123.46
5%	121.09	119.29	121.04	122.76
10%	122.9	121.60	123.60	125.78
15%	123.16	121.42	122.67	124.41

In Table 4.3 we use the bid functions for type A and type B subjects estimated from a pooled data set in each experiment (pooled over rounds and subjects). Using these bid functions we ask what the cost of purchasing would have been in the $x\%$

Experiment Pair	P-Value
0% vs 5%	.232
0% vs 10%	.091
0% vs 15%	.187
5% vs 10%	.953
5% vs 15%	.953
10% vs 15%	.669

Kolmogorov-Smirnov Tests forEquality of Cost Distributions B-Types

Experiment Pair	P-Value
0% vs 5%	.843
0% vs 10%	.607
0% vs 15%	.334
5% vs 10%	.252
5% vs 15%	.131
10% vs 15%	.980

experiment if round-by-round they would have received the cost realizations generated by the $y\%$ experiment. For example, down the first column we have listed the experiments whose data we are to use in our calculation. Across the rows we see what the cost of purchasing would have been had we received those cost realizations but had subjects who were behaving according to the bid functions in the various experiments. For example, looking at the entry in row 1, column 2 we see what the government's cost of purchasing would have been had the subjects in the 5% experiment received the cost realizations that occurred in the 0% experiment. The cost of purchasing for the government would have been 119.74 instead of the 119.29 it actually was in the 5% auction. Conversely, looking at row 2, column 1, we see that had the cost realizations in the 5% experiment been offered to the bidders in the 0% experiment, the cost of purchasing would have increased from 121.09 to 121.24. In other words, if we look across row 1 we find that using the cost realizations in the 0% auction and the behavior in other auctions, the government would have saved 1.23% if it had imposed a 5% preference on bidders while it would have cost it .3% to use a 10% preference rule and 1.79% if it had used a 15% preference rule. In addition, as we know from our previous discussion of representation, not only would the 5% rule have been cost effective but it also would have increased the percentage of contracts going to high cost (minority) firms from 13.7% in the 0% preference case to 28.2% when 5% was used. Interestingly, McAfee and McMillian (1989) estimate that with the parameters used in our experiment a 4.1% price-preference rule would have been optimal among the class of percentage price-preference rules[13].

If we had used the naive approach to procurement cost here and simply taken the difference in cost between the price paid for a contract in all those instances when the

[13] Because even equilibrium purchasing costs are not predicted to be dramatically different between the 0% and 5% price-preference auctions, it is hard for the differences we observe to pass a significance test. Still, when using the 5% data and comparing purchasing costs in the 5% and 0% auctions we find that the difference between 119.29 and 121.09 was significant at the 5% level (p-value 0.0503). The reverse, using the 0% data for comparison only yielded a difference significant at the 19.36% level. A Wilcoxon test performed on the actual prices formed in the 0% and 5% auctions, unconditional on cost realizations, shows a difference significant at the 11.68% level.

price-preference rule was invoked and the lowest bid for that contract made by any bidder of type B, we would have come to the conclusion that imposing the 5% rule would have cost us .3% while the 10% rule would have cost us 1.22% and the 15% rule 1.76%.

To demonstrate the relationship between the cost of purchasing and the price-preference rule used we present Figure 4.2 where we place the price-preference rule along the horizontal axis and the cost of purchasing along the vertical. Inside the figure we plot two graphs. One is the theoretical relationship between the cost of purchasing and the price-preference rule generated by the theory while the other is the actual cost of purchasing determined by the experiments. What we see is that while both curves have a common shape indicating that the 5% preference should be the best among these four price preference rules, the actual cost of purchasing is everywhere below the predicted cost. This fact is consistent with the fact that in all of these experiments subjects tend to bid more aggressively than the theory predicts (see Section 4.5 below).

Figure 4.2 Here

Another way to see if our price-preference rule is cost effective is to estimate a bid function (pooled) for the subjects in each price-preference auction and see how bidding behavior changes as the price-preference rule is increased. Such estimated bid functions offer us insight into the mean bid made by subjects conditional on their cost realization which is exactly the information we desire. These regressions were at first run in a non-linear fashions using a second degree polynomial to fit the data. However, since the quadratic term proved to be insignificant in all regressions run, we will report the results only of our linear regressions.

From a policy perspective, however, we are not interested in the behavior of the bid functions over domains where the probability of winning is close to zero (these functions will be estimated later when we turn attention to testing the theory of auctions). For example, as Table 4.2 indicates, in the 0% experiment no type A subject won an auction who drew a cost greater than 125 while no type B subject won with a cost realization greater than 156. In fact no bidder of any type in any laboratory auction run ever won with a cost realization greater than 161. (This is not strange since the distribution of winning bids has more in common with the distribution of the lowest order statistic than the distribution of the parent distribution). Since cost realizations beyond these bounds were soon known to imply no chance of winning, the behavior of bidders when these cost realizations arose is irrelevant to the policy question posed since we do not care about bidding behavior in those regions of the domain where winning is a low probability event. Some bidders with high cost realizations bid their cost while others submitted throw-away bids of arbitrarily high amounts. Hence, estimating a bid function using these cost-bid pairs would greatly distort the estimated bid functions. To rectify this situation we truncated the domain of the bid-function space to coincide with that region where the empirical probability of winning was positive. These "truncated" regressions then contain observations for costs drawn in the interval where at least one bidder had won an auction. Figures 4.3-4.4 present the estimated regression lines of these truncated bid function regressions while Tables 4.4a-4.4b and 4.5a-4.5b present the actual estimates of the coefficients for the truncated and untruncated bid functions, respectively.

Figure 4.3-4.4 Here

As we can see from Figures 4.3 and 4.4, while the changes in bidding behavior were subtle they do present a consistent pattern. For type B subjects the imposition of price-preference rules leads to more aggressive bids over a substantial portion of the range of winning cost realizations. This can be seen by the fact that the 0% auction regression line is above those of the other auctions over the range of cost realizations where all wining bids come from. For type A bidders the situation is more complex. While the bid functions for the 5% and 15% auctions are initially below that of the 0% auction, they quickly rise above it. Hence, for A types, the imposition of price-preference rules does lead to less aggressive bidding over large portions of the domain of costs[14]. Still, as theory indicates, the more aggressive bidding by B types more than compensates for this behavioral change on the part of A types. Still note, however, that over a significant portion of the domain where winning occurs, the 5% bid function for A types is below that of A type subjects in the 0% auction.

Tables 4.4-4.5 Here

4.4. Summary

Our representation and procurement cost results are summarized in Table 4.6. In this table, a "plus" is placed where, when compared to the 0% auction, either the representation of high cost firms or the cost effectiveness of an auction has increased and a "minus" where it was decreased.

Table 4.6

Summary of Representation of High Cost Firms and Cost Effectiveness Results

Experiment	Representation	Cost Effectiveness
5%	+	+
10%	+	-
15%	+	-

As we can see from this table, when the 5% price-preference is used, both representation of high cost firms and cost effectiveness are enhanced when compared to the no-preference (0%) rule. For the 10% and 15% rules, however, while representation is enhanced, cost effectiveness is not.

As we noted in Section 3.2 the percentage price-preference auctions we ran are not optimal auctions. We investigate price-preference auctions because they are the ones most frequently used. Still, one might ask how much better the government might be

[14]In post-experimental surveys given subjects, 16 out of 22 B-Type subjects who participated in experiments with positive preferences indicated that the existence of the price-preference rule led them to bid lower than they would have if no preference existed. For A-Types, 3 of 9 subjects indicated that the preference caused them to raise their bids.

if it used the optimal auction form explained in Section 3.1. These calculations are presented in Table 4.7.

As we see, in theory the optimal mechanism, by definition, gives lower procurement costs for the government than does any price-preference auction. However, the experimental price-preference auctions performed better than the optimal mechanism would have had it been implemented using the actual cost realizations of our experimental auctions. Compare the bottom two rows of Table 4.7. For example, while the optimal mechanism predicts an expected procurement cost of 125.53 for the realizations of the 5% price-preference groups, when the 5% price-preference rule was used in our experiment the expected procurement cost was 119.29. This is clearly attributable to the underbidding of both types of subjects in this group.

Table 4.7

Procurement Cost Comparisons:

Optimal Mechanism vs. Price Preference Rules

Procurement Cost	**0% Group**	**5% Group**	**10% Group**	**15% Group**
Theoretical				
% Preference Rule	134.13	133.24	134.50	136.06
Optimal Mechanism	132.19	132.19	132.19	132.19
Actual				
% Preference Rule	121.24	119.29	123.60	124.41
Optimal Mechanism*	124.97	125.53	126.49	127.70

* The costs in this row are the expected procurement costs that would be generated had we used the optimal mechanism with the actual cost realizations of each group.

4.5. Theory and Behavior

The theory of auctions makes precise predictions about the behavior of subjects. In this section we attempt to examine whether the observed behavior of our subjects was consistent with the predictions of the theory, and if not, in what direction did deviations occur.

Since the theory of auctions makes a prediction about the bids of subjects conditional on any cost drawn, our test of the theory will consist of the regression

$$b_a = \alpha + \beta b_p \tag{4.9}$$

where b_a is the actual bid made by a subject with a particular cost draw and b_p is the theoretical predicted bid for that cost. Clearly, if the theory predicts well, we expect $\alpha = 0$ and $\beta = 1$.

The results of the regression run for our four experiments are presented in Tables 4.8a-b.

Tables 4.8a-b Here

As we see in Tables 4.8a-b, the theory was not supported by the regressions we ran. More precisely, for all experiments the null hypothesis that $\alpha = 0$ and $\beta = 1$ was rejected in favor of the one-tail hypothesis that $\alpha < 0$ and $\beta > 1$ at the 5% level of significance. These results imply that in not following the theory, bidders tended to under bid relative to the prediction of the theory for most of the domain of costs.

To illustrate this point, consider Figures 4.5a-4.5d and 4.6a-4.6d which show the estimated actual bid functions used by our subjects juxtaposed to the bid functions predicted by the theory.

Figures 4.5-4.6 Here

As was demonstrated in the regression results of the actual versus predicted bids, in all experiments the actual bid functions of our subjects were below those predicted by the theory. In fact, this divergence was greatest in the lower regions of the cost domain where there was the highest probability of winning for any bidder. Hence, these figures also illustrate the fact that our bidders exhibited a "utility of winning" not contained in the theory since they tended to sacrifice some monetary surplus in order to increase their probability of winning[15].

4.6. Learning, Performance, and Theory

All of our analysis of the performance of price-preference auctions has been time-aggregated in the sense that we have pooled our observations over time to create one large data set. However, since our auctions were run for 20 periods, it is logical to ask whether behavior over the last 10 periods was different from behavior over the first 10. It can be argued that once subjects learn the strategic situation they are in they might change their behavior and this behavioral change might reverse our policy prescriptions.

To investigate this possibility, we disaggregated some of the calculations made earlier and ran some of the same regressions on disaggregated data (data taken from the first ten and last ten periods). These recalculations are presented in Tables 4.9 and 4.10.

In these tables we present the performance characteristics of our auctions (Table 4.9) over the first ten and last ten periods of the auction and look to see if the 5% price-preference rule still out-performs the other three rules. In addition, we look to see if there is any general tendency for behavior to change over time which would indicate that with learning our cost effectiveness results might disappear.

Table 4.9 Here

As we can see from Table 4.9, as time goes on in these auctions it appears that both A-Types and B-Types bid in a more aggressive manner. In three of the four experiments (all except the 5% price-preference rule auction) average price falls between

[15] Other experimentalists have noted this "utility of winning" phenomenon. For example, Cox, Smith, and Walker (1988) use this phenomenon to explain the divergence of bids in their auctions from the theoretical predictions for risk-neutral bidders, while Bull, Schotter and Weigelt (1992) found similar behavior in the tournaments they ran in which subjects exerted "too much" effort in an attempt to win the big prizes offered.

the first and last ten rounds as does the average profit per unit (i.e. the mark-up of bids over costs). For example, while B-Type subjects in the 15% price-preference auction experiment tended to bid, on average, 8.5 units above their cost in the first ten rounds of their experiment, this mark-up falls to 4.52 over the last ten rounds. The only subjects for which this fact was not true were the B-Types in the 0% experiment.

As was true of our aggregate analysis, if we are to properly compare the government cost of purchasing across price-preference rules, we must find a way to hold cost constant and attribute differences in the cost of purchasing to differences only in bidding behavior. This is what we did in Table 4.3 for our aggregate data and what we repeat here in Table 4.10a and 4.10b in a time-disaggregated fashion.

Table 4.10a-b Here

The interesting thing about these tables is the fact that when we move from the first to the last ten rounds we preserve our basic policy conclusion that the 5% price-preference auction is best among the class of four price-preference rules explored here. (Note that the entry in the 5% column is the lowest in each row of the matrix indicating that using the cost realizations from any experiment, the purchasing cost of the laboratory government would have been lowest if the 5% rule was used. This was true in both the first and last ten rounds). Furthermore, in comparing Table 4.10a to 4.10b (moving from the first to the last ten rounds) we see that the purchasing cost of our laboratory government decreases in 12 of the 16 cell comparisons indicating once more a shift toward more aggressive bidding in the last ten as opposed to the first two rounds of our experiments. These comparisons tends to indicate that experience would not be responsible for an alteration of our basic policy conclusion that affirmative action price-preference auctions are cost effective.

To get a better idea of how the cost of purchasing behaves over time consider Figure 4.7 which presents the actual cost of purchasing over the first and last ten rounds of our experiment. In fact, this diagram plots the costs portrayed on the diagonals of Tables 4.10a and 4.10b. As one can see, the realized cost of purchasing fell during the last 10 rounds of the experiment for all price-preference rules except the 5% rule where it rose minimally. However, note that for both the first and last ten rounds of the experiment, the 5% price-preference rule offered a lowest cost of purchasing than any other price-preference rule tested.

Figure 4.7 Here

While the focus of our analysis in this section has been to ask whether different price-preference auctions perform differently as subjects learn more about the auctions they are functioning in, we might also be interested in whether they more closely approximate the behavior predicted by the theory as time progresses. To investigate this question we re-estimated the untruncated bid functions of our subjects in a disaggregated manner running one regression on the first ten and one on the last ten rounds for each auction and each type of bidder. The results are summarized in Figures 4.8a and 4.8b, and the regression results are presented in Table 4.11a and 4.11b.

Figures 4.8a and 4.8b Here
Tables 4.11a and 4.11b Here

As we can see from Figure 4.8, there is no difference between the bid functions employed by subjects in the first ten versus the last ten rounds of our experiment.

5. Conclusions

In this paper we have demonstrated that the use of affirmative action programs need not force policy makers to have to choose between the benefits of minority representation and cost effectiveness. In certain circumstances, such as the price-preference auctions investigated here, both minority representation and cost effectiveness can be enhance simultaneously if the proper price-preference rule is used. This result seems robust to learning on the part of bidders in the sense that after subjects have experience with bidding in such auctions (i.e. after the first ten rounds of any experiment) the auctions run using a 5% price-preference rule continue to outperform those in which no preference is given and in general bidders bid more aggressively (closer to their cost) as the auction progresses.

Our results do raise a note of caution, however, since if the government fails to use the optimal price-preference rule, it could increase its cost of purchasing and hence fail to reap the benefits that such price-preference rules offer. McAfee and McMillian (1989) offer a rough rule of thumb to help decision makers choose the correct preference rule (when costs are drawn from uniform distributions) which is to take one-third the difference in the mean cost difference between the two types of firms (i.e. the percentage difference in their means). In our experiments their rule suggests a 3.3% preference which is close to the 5% rule which proved to be best among our limited set of four rules.

References

Ayres, Ian and Cramton, Peter, “Pursuing Deficit Reduction Through Diversity: How Affirmative Action at the FCC Increased Auction Competition”, Stanford Law Review, (forthcoming 1996).

Bulow, Jeremy and Roberts, John, “The Simple Economics of Optimal Auctions”, Journal of Political Economy, Vol. 97, 1993

Cox, James C., Vernon L. Smith and James M. Walker (1988) “Theory and Individual Behavior of First Price Auctions”, Journal of Risk and Uncertainty, March 1988

Clegg R. ed., Racial Preference in Government Contracting, National Legal Center for the Public Interest, Washington D.C., (1993)

Jordan, D. (1978) “In-State Preferences in Public Contracting: State’s Rights vs. Economic Sectionalism”, University of Colorado Law Review, Vol. 49 pp. 205-231

Kolmogorov, A. and Fomin, S., Introductory Real Analysis, Dover Press, New York (1975).

LaNoue, G. Minority Programs and Disparity Studies, National League of Cities, Washington D.C., (1994)

Lowinger, T. “Discrimination in Government Procurement of Foreign Goods in the U.S. and Western Europe”, Southern Economic Journal, Vol. 42, pp.451-460, (1976)

Maskin, Eric and John G. Riley (1995) “Asymmetric Auctions”, mimeo , U.C.L.A., June 1995

McAfee, R. Preston and John McMillan (1987) “Auctions and Bidding”, Journal of Economic Literature, Vol. 25, pp. 699-738

McAfee, R. Preston and John McMillan (1989) “Government Procurement and International Trade”, Journal of International Economics, Vol. 26, pp. 291-308

Myerson, Roger (1981) “Optimal Auction Design”, Mathematics of Operations Research, Vol. 6, pp., 58-73.

National Association of State Purchasing Officials, “State and Local Government Purchasing”, Fourth Edition, National Association of State Purchasing Officials, Lexington, Kentucky, 1994

Results: 1993 Procurement Practices, National Institute of Government Purchasing, Reston Virginia (1994)

Riley, John and Li, H. "BIDCOMP^2 −−A Freeware Program to COMPute Equilibrium Bids and COMPare Expected Revenue" Mimeo, U.C.L.A. (1993).

Schotter, Andrew (1995) "Price Preference Procurement Auctions: Background and Methodology", New York University, mimeo.

Schotter, A. and Weigelt, W., "Asymmetric Tournaments, Equal Opportunity Laws, and Affirmative Action: Some Experimental Results", Quarterly Journal of Economics, May 1992, pp. 511-539.

Table 4.2
Descriptive Results of Price Preference Auctions

Auction Type	**0%**	**5%**	**10%**	**15%**
Winners				
A-Types	12	22	28	43
B-Types	87	78	52	57
Type A Win %	12.1%	22.0%	35.0%	43.0%
Theoretical Type A Win%	19.1%	24.9%	29.8%	36.9%
Wins by Preference		10	19	23
Avg. Cost Winners				
A-Types	117.17	120.68	127.64	123.74
B-Types	113.82	111.44	112.62	112.77
Avg. Bid Winners				
A-Types	124.42	126.95	134.14	131.56
B-Types	120.80	117.13	117.92	119.02
Highest Cost Winner	A's-125 B's-151	A's-151 B's-155	A's-160 B's-144	A's-148 B's-146
Avg. Profit Per Unit				
A-Types	7.25	6.27	6.50	5.47
B-Types	6.99	5.69	5.31	5.68
Average Observed Price	121.24	119.29	123.60	124.41
Theoretical Price	134.13	133.24	134.50	136.06
Naive Cost of Preference	0	0.39	1.51	2.04
Actual Cost	0	-1.50	0.39	2.22

Table 4.4a
Truncated Bid Function Regressions: (A-Type Subjects)
Equation: Bid = Intercept +slope·(Cost-110)

Auction	Intercept (α)	Coefficient (β)	Observations	R^2
0%	121.73 (22.17)	.669 (3.72)	43	.256
5%	117.88 (19.97)	.889 (20.40)	82	.802
10%	121.73 (18.77)	.811 (17.56)	82	.793
15%	119.44 (18.60)	.93 (15.25)	82	.743

.
.

Table 4.4b
Truncated Bid Function Regressions: (*B*-Type Subjects)Equation
Bid = Intercept +slope·(Cost-100)

Auction	Intercept (α)	Coefficient (β)	Observations	R^2
0%	109.72 (19.69)	.876 (38.08)	242	.849
5%	107.08 (18.62)	.933 (39.87)	251	.864
10%	107.97 (13.96)	.927 (20.42)	166	.717
15%	108.33 (25.01)	.882 (40.09)	210	.877

(t-stats in parentheses.)

Table 4.5a
Untruncated Bid Function Regressions: (A-Type Subjects)
Equation: Bid = Intercept +slope·(Cost-110)

Auction	Intercept (α)	Coefficient (β)	Observations	R^2
0%	118.86 (169.78)	.944 (78.41)	188	.97
5%	116.52 (141.76)	.964 (63.11)	176	.96
10%	118.51 (129.94)	.936 (63.14)	159	.96
15%	118.69 (141.25)	.941 (72.45)	198	.96

Table 4.5b
Untruncated Bid Function Regressions: (B-Type Subjects)
Equation: Bid = Intercept +slope·(Cost-100)

Auction	Intercept (α)	Coefficient (β)	Observations	R^2
0%	108.46 (214.89)	.938 (96.91)	398	.96
5%	106.33 (186.22)	.971 (87.71)	398	.95
10%	107.28 (144.40)	.968 (69.70)	306	.94
15%	106.61 (193.43)	.980 (87.72)	363	.96

(t-stats in parentheses.)

Table 4.8a
Regression of Actual on Predicted Bids (B-Type Subjects)

Auction	Intercept (α)	Coefficient (β)	Observations	R^2
0%	−26.546 (−17.367)	1.140 (116.46)	396	.9717 (.9717)
5%	−33.376 (−15.879)	1.183 (87.378)	398	.9750 (.9705)
10%	−21.345 (−8.54)	1.105 (69.763)	306	.9701 (.9412)
15%	−34.636 (−16.457)	1.204 (87.333)	363	.9771 (.9548)

Table 4.8b
Regression of Actual on Predicted Bids (A-Type Subjects)

Auction	Intercept (α)	Coefficient (β)	Observations	R^2
0%	−22.481 (−9.237)	1.126 (77.731)	188	.9849 (.9701)
5%	−34.771 (−11.061)	1.182 (62.838)	176	.9786 (.9577)
10%	−24.493 (−7.94)	1.108 (63.212)	159	.9809 (.9621)
15%	−44.249 (−14.696)	1.224 (71.518)	198	.9813 (.9630)

(t-stats or Adjusted R^2 in Parentheses)

Table 4.9
Descriptive Results of Price Preference Auctions: First and Last Ten Rounds

First Ten Rounds				
Auction Type	**0%**	**5%**	**10%**	**15%**
Winners				
A-Types	**7**	10	16	26
B-Types	42	40	24	24
Wins by Preference		6	11	13
Avg. Profit per Unit				
A-Types	7.57	6.60	7.44	8.42
B-Types	6.67	6.40	6.50	8.50
Average Price	121.71	119.04	124.45	127.10
Naive Cost of Preference	0	.46	1.8	2.42
True Cost		-0.62	1.92	3.81
Last Ten Rounds				
Auction Type	**0%**	**5%**	**10%**	**15%**
Winners				
A-Types	5	12	12	17
B-Types	45	38	28	33
Wins by Preference		4	8	10
Avg. Profit per Unit				
A-Types	6.80	6.0	5.25	6.88
B-Types	7.29	4.95	4.29	4.52
Average Price	120.78	119.54	122.75	121.74
Naive Cost of Preference	0	.32	1.18	1.98
True Cost		-1.74	-0.78	-0.03

Table 4.10a
Purchasing Cost Comparisons: Rounds 1-10

Cost Realizations	Cost in 0% Auction	Cost in 5% Auction	Cost in 10% Auction	Cost in 15% Auction
0%	121.71	121.09	121.63	125.52
5%	120.81	119.04	121.24	123.49
10%	122.85	121.92	124.45	127.15
15%	125.05	124.18	127.01	127.10

Table 4.10b
Purchasing Cost Comparisons: Rounds 11-20

Cost Realizations	Cost in 0% Auction	Cost in 5% Auction	Cost in 10% Auction	Cost in 15% Auction
0%	120.78	119.04	120.00	120.75
5%	121.00	119.54	120.68	123.10
10%	123.05	121.23	122.75	124.88
15%	121.25	119.28	120.77	121.74

Table 4.11a
Bid-Function Regressions Rounds 1-10 and 11-20: Untruncated Data
Type *A* Subjects

Experiment	Rounds	Number	Intercept, α	Slope, β	R^2
0%	1-10	97	120.28 (151.97)	.923 (68.02)	.97
0%	11-20	91	117.09 (98.51)	.971 (47.39)	.96
5%	1-10	90	117.95 (78.35)	.960 (34.50)	.93
5%	11-20	86	115.18 (189.19)	.965 (84.82)	.99
10%	1-10	80	118.91 (99.62)	.934 (46.71)	.97
10%	11-20	79	117.98 (83.28)	.940 (56.17)	.96
15%	1-10	98	119.84 (96.88)	.933 (46.45)	.96
15%	11-20	100	117.44 (113.36)	.950 (58.14)	.97

Table 4.11b
Bid-Function Regressions Rounds 1-10 and 11-20: Untruncated Data
Type B Subjects

Experiment	Rounds	Number	Intercept, α	Slope, β	R^2
0%	1-10	197	108.73 (161.95)	.942 (71.93)	.96
0%	11-20	200	107.93 (220.40)	.930 (1010.9)	.96
5%	1-10	199	107.34 (135.08)	.957 (60.08)	.95
5%	11-20	199	105.25 (128.52)	.987 (64.03)	.95
10%	1-10	154	108.73 (93.57)	.963 (45.74)	.93
10%	11-20	152	105.96 (118.64)	.970 (56.17)	.95
15%	1-10	178	108.27 (113.00)	.966 (50.71)	.94
15%	11-20	186	105.48 (177.54)	.984 (79.90)	.97

Figure 4.1a: Cost Realizations (Type A)

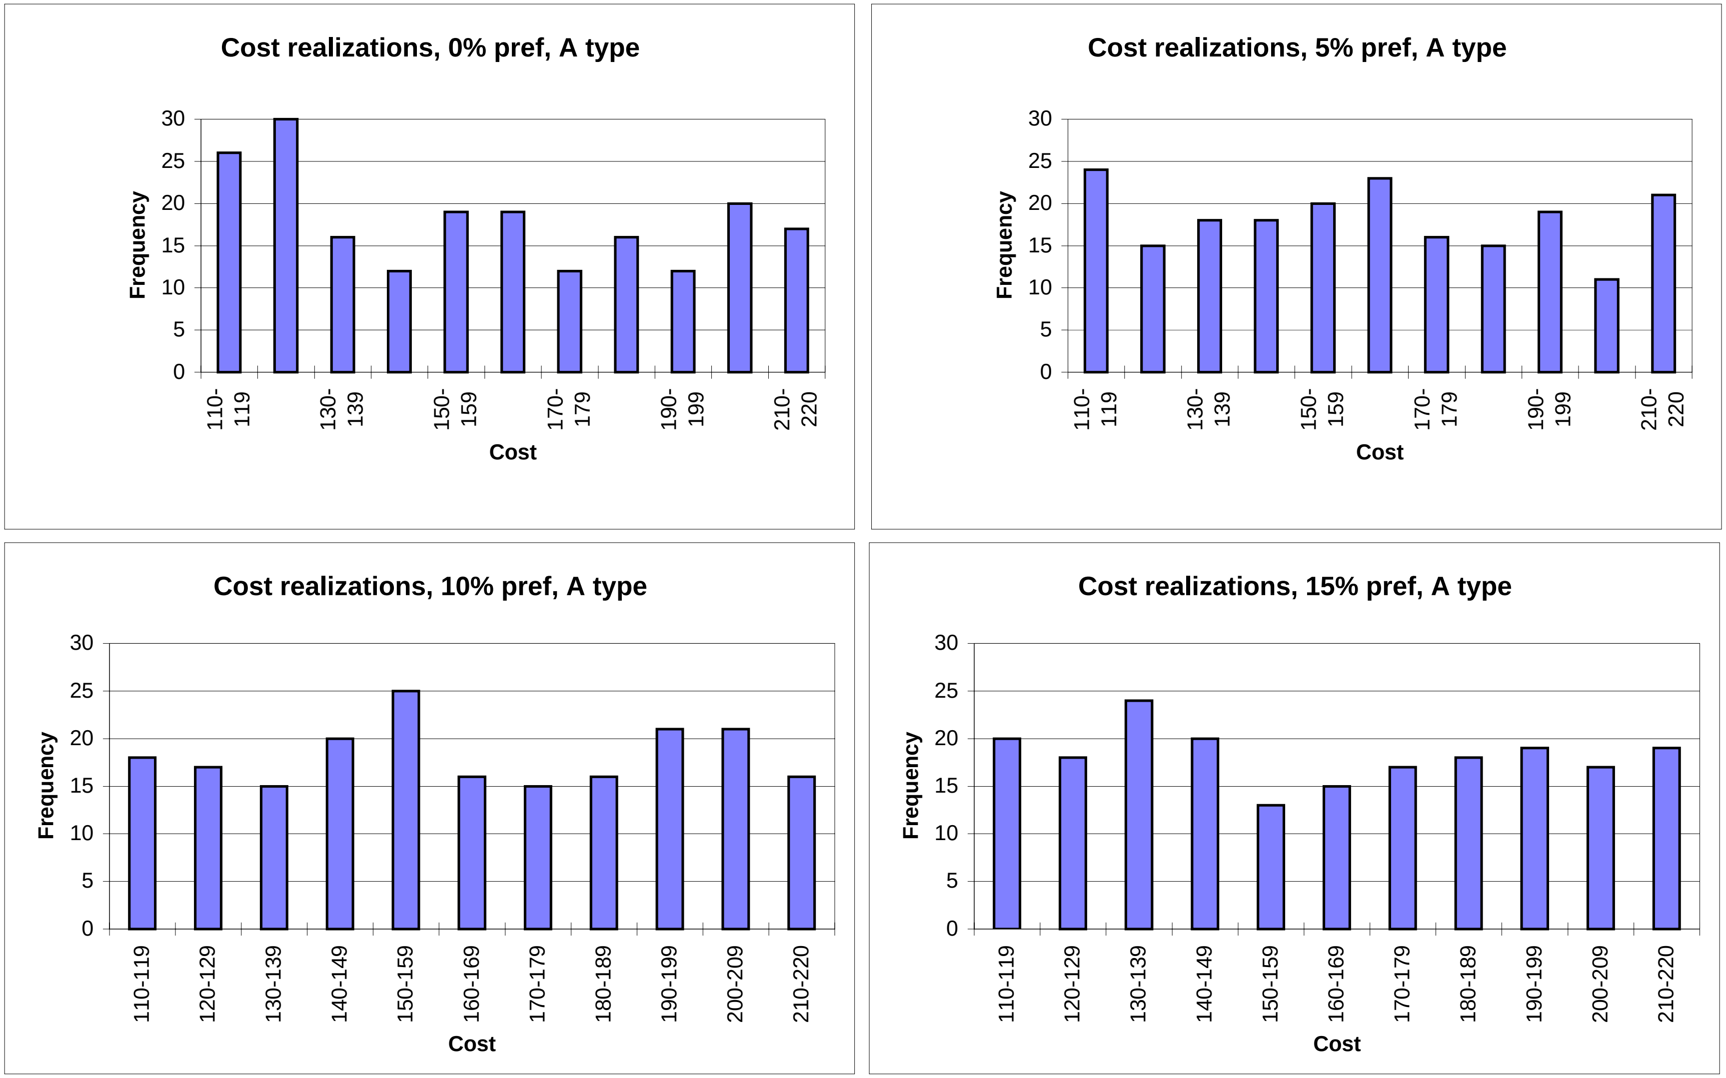

Figure 4.1b: Cost Realizations (Type B)

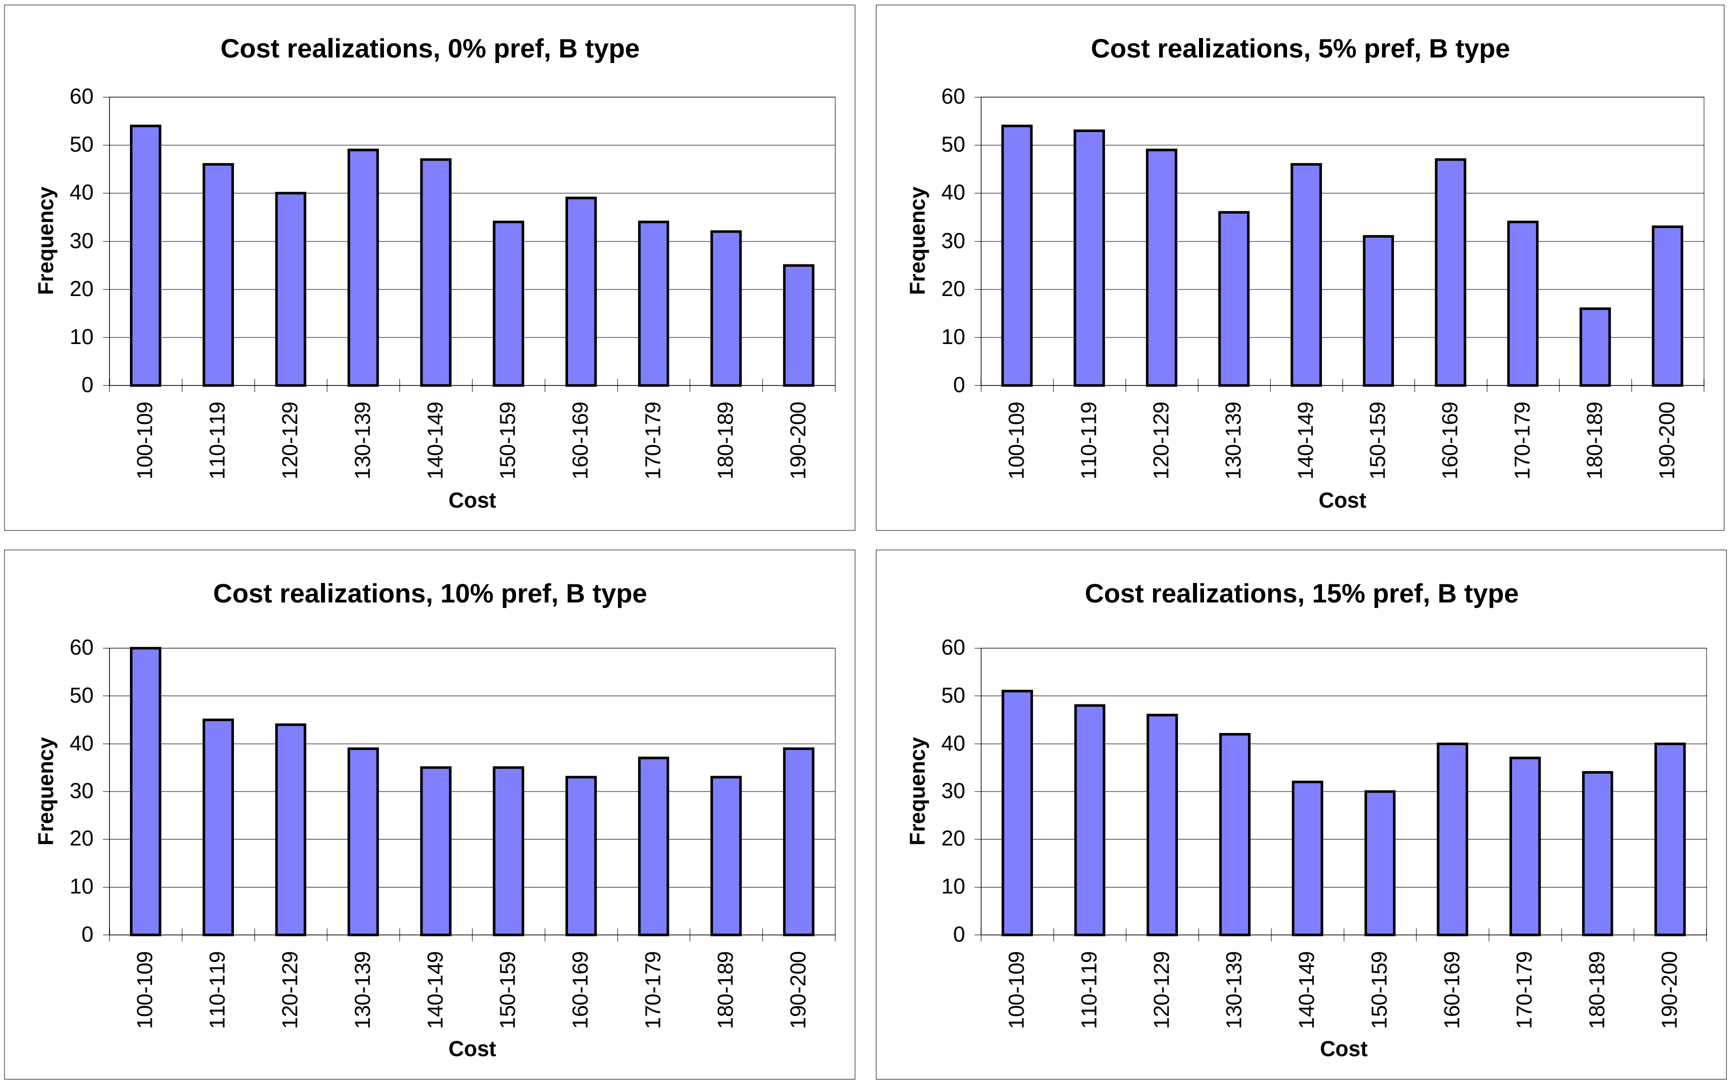

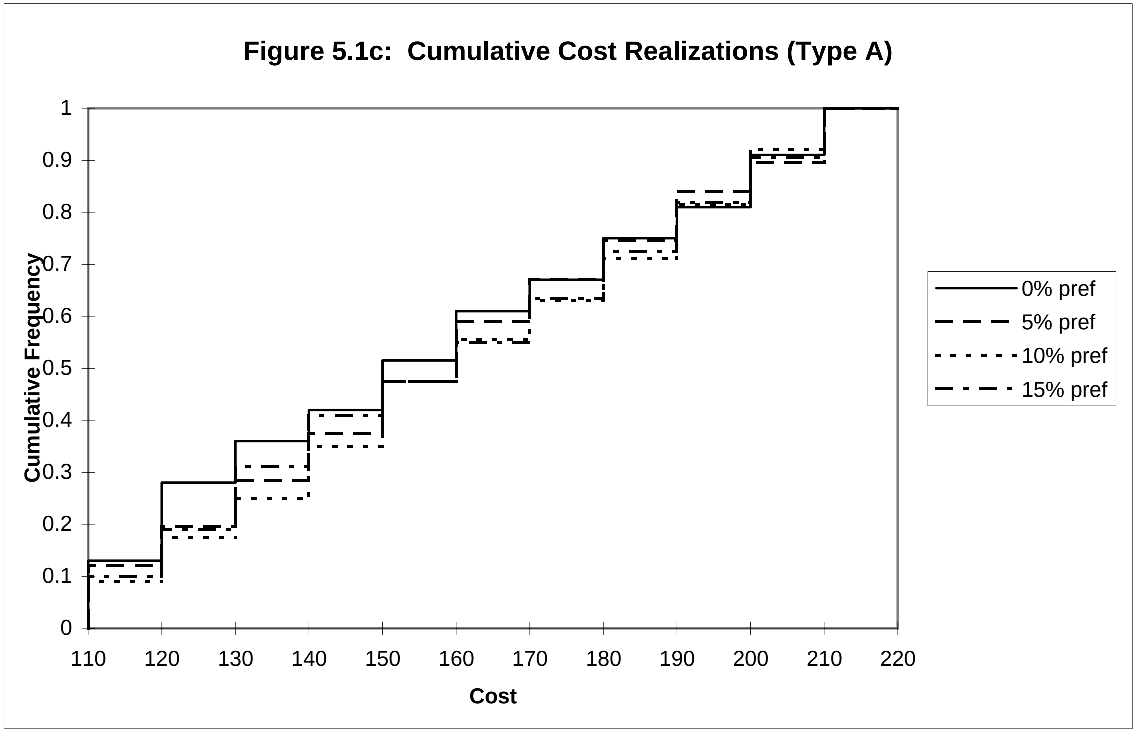

Figure 5.1c: Cumulative Cost Realizations (Type A)
Cumulative Frequency
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
110
120
130
140
150
160
170
180
190
200
210
220
Cost
0% pref
5% pref
10% pref
15% pref

Figure 5.1d: Cumulative Cost Realizations (Type B)