Procurement Auctions with Renegotiation and Wealth Constraints*

Wei-Shiun Chang†
National Cheng Kung University

Timothy C. Salmon‡
Southern Methodist University

Krista J. Saral§
Webster University Geneva
CNRS – GATE Lyon St Etienne

August 2015

Abstract

Renegotiation is a common practice in procurement auctions which allows for post-auction price adjustments and is nominally intended to deal with the problem that sellers might underestimate the eventual costs of a project during the auction. Using a combination of theory and experiments, we examine the effectiveness of renegotiation at solving this problem. Our findings demonstrate that renegotiation is rarely successful at solving the problem of sellers misestimating costs. The primary effect of allowing renegotiation is that it advantages sellers who possess a credible commitment of default should they have underbid the project. Renegotiation allows these weaker types of sellers to win more often and it also allows them to leverage their commitment of default into higher prices in renegotiation from a buyer.

JEL Codes: C91, D44, D82

Key Words: Procurement auctions, renegotiation, bankruptcy, default, experiments

---

*Financial support for this project was provided by IFREE. We thank seminar participants at ESA NYC, ESA Zurich, ASFEE Lyon and Oklahoma University. We in particular would like to thank Georgia Kosmopoulou, Jason Shachat, and Marco Pagnozzi for useful suggestions.

†Institute of International Management, National Cheng Kung University No.1, University Road, Tainan City 701, Taiwan (R.O.C.) wschang@mail.ncku.edu.tw Phone: +886-6-2757575 ext.53566

‡Southern Methodist University, Department of Economics, 3300 Dyer Street, Suite 301 Umphrey Lee Center, Dallas, TX 75275-0496. tsalmon@smu.edu, Phone: 214-768-3547, Fax: 214-768-1821.

§George Herbert Walker School of Business and Technology, Webster University Geneva, Route de Collex 15, CH-1293 Bellevue, Switzerland. CNRS – GATE Lyon St Etienne Email: kjsaral@webster.ch

1 Introduction

A common problem in procurement auctions is that a seller will not know the true cost of completing the project at issue when placing a bid. Projects that require a longer time frame to complete could experience substantial labor and materials cost shifts over the course of the project that may not be foreseeable prior to the auction. Also, unexpected difficulties may emerge requiring additional design work, or different production technologies that were unanticipated during the bidding phase. Despite these difficulties, costs must be estimated prior to an auction in order for a seller to be able to construct a bid they expect to be profitable. Due to the cost uncertainty it is likely that some sellers will overestimate the final cost while others will underestimate it. If the procurement auction takes the form of a standard low-price auction in which the lowest bid wins and the winning seller is paid their bid, then it is likely to be the case that the winning seller will be the one that underestimated the cost of the project by more than any of his competitors. This adverse selection is what leads to the well-known Winner's Curse phenomenon and if not taken into account during bid formulation it can lead to substantial losses for the sellers.[1]

One of the more commonly used mechanisms to deal with the possibility that a seller has underbid is renegotiating the terms of a contract after it is awarded and costs are realized. Two well-known examples of post-award renegotiation would be the North South metro line in Amsterdam and the Big Dig in Boston.[2] While these are highly visible examples of renegotiation following a competitive bidding process, the phenomenon occurs in a wide range of situations, including smaller-scale procurement projects which are particularly susceptible to bankruptcy due to the weaker financial standing of many of the firms (Calveras, Ganuza, and Hauk (2004)). Decarolis (2014) examines a data set of procurement auctions in Italy and finds that much of the price reduction achieved by using auctions for the procurement process are lost through renegotiation. Bajari, Houghton, and Tadelis (2014) provides estimates of what they call adaptation costs in highway paving contracts which amount to an 8-14% increase in the final price over the initial price set by auction. Wang (2000) provides a theoretical argument for how allowing for post-auction renegotiation could help a buyer obtain lower prices but this is in a private value environment

[1] The existence and prevalence of the Winner's Curse has been well documented in both the lab and the field (see for example Kagel and Levin (1986); Charness and Levin (2009); Bajari and Hortacsu (2003)).

[2] Contractors working for the City of Amsterdam began construction of a new North-South line of their metro to go through the central part of the city in 2002. The initial budget was set at €1.46 billion but after multiple rounds of renegotiating, it had risen to €3.1 billion by 2009. It was originally intended to be completed by 2011, but latest estimates now suggest it might be completed by 2017. The Big Dig in Boston has suffered similar problems leading to over $1 billion in renegotiated payments to the construction firms involved. The Boston Globe maintains a repository of information on the history of the Big Dig at http://www.boston.com/news/specials/big_dig_problems/

in which there is no Winner's Curse problem to deal with. In a common or affiliated values environment, the theoretical argument for how renegotiation could resolve the problem is less clear. The goal for a buyer would be to use renegotiation to help sellers who accidentally bid too low avoid default. By doing so the buyer hopes to avoid additional recontracting costs associated with finding another seller to complete the project but this results in obvious risks to the buyer.

While renegotiation is used in practice as a remedy against potential bankruptcy from underestimated costs, there is no clear theoretical case for why this might be effective. The reason is that allowing for renegotiation may actually serve to exacerbate the low-bids problem it is intended to solve. If sellers believe that renegotiation is a viable option and that they can obtain favorable prices through renegotiation then this changes the incentives of the auction stage. Instead of competitively bidding to make a margin of profit based on the price obtained in the auction, sellers will competitively bid for the right to renegotiate the final price one-on-one with the buyer. This will distort the efficiency of the allocation as the winning seller may be the firm with the most bargaining power and not necessarily the firm with the lowest cost. It may also lead to very high final prices for the buyer.

The purpose of this study is to determine how the addition of a renegotiation option may change the outcome of a procurement process; in particular we are interested in whether renegotiation can help to overcome bidder problems caused by the Winner's Curse in first price auctions. Given that the renegotiation mechanism is commonly in use in conjunction with first price auctions, it seems that some must believe the combination to be effective and so our aim is to fully evaluate if it can be successful. We provide both theoretical and empirical analyses of this issue. The empirical analysis is based on an economic experiment constructed to allow us to identify the degree to which the theoretical predictions of allowing a post auction renegotiation stage match with actual behavior. We use an experiment rather than field data due to the difficulty of controlling for whether renegotiation would be a possibility in field auctions.

Our theoretical analysis of renegotiation begins by simply adding a post-auction renegotiation phase onto a standard low price auction. In doing so, it becomes immediately apparent that in order for renegotiation to have any impact on the process, from a theoretical perspective, the winning seller must have bargaining power in the renegotiation phase. We model this bargaining power as a credible commitment to abandon the project after learning the true cost. If a seller lacks such a credible commitment, then the buyer will have no incentive to approve a higher price through renegotiation and a seller clearly would never request a lower price. We therefore augment

the model by examining situations in which sellers have varied degrees of credible commitment to default. The commitment to default is modeled on the premise that firms differ according to their size and borrowing/loss constraints. One can well imagine large firms with deep pockets, who have the ability to sustain losses on individual projects, choosing to absorb those losses rather than abandon the project. This is because such firms might face substantial costs from defaulting on one bad deal due to the possibility that such an action could have adverse consequences on multiple other projects in which the firm is engaged. On the other hand, smaller firms may be unable to complete projects whose price is lower than the costs due to liquidity constraints. For ease of analysis, we restrict our theoretical environment to two types of sellers: strong sellers who can (and will) accommodate any loss without defaulting and weak sellers who cannot absorb any loss and so default.

The results from our theoretical analysis demonstrate that simply adding on a renegotiation phase to a standard low price auction should not be expected to alleviate any problems arising from the Winner's Curse. When we add in sellers with credible commitment to default in the event of a loss, our results show that they can use this to exploit the buyer for high prices in renegotiation. These results indicate that renegotiation is unlikely to be a successful mechanism in overcoming problems from the Winner's Curse.

While theoretically the case in favor of renegotiation is not strong, we note that the mechanism is still in common use in the field and so it might be the case that the mechanism performs better in practice than expected. There are many reasons one might expect renegotiation to be more successful than predicted in our base model. There is substantial prior experimental research showing that individuals possess some potential preference for fairness or equity in bargaining situations (e.g. Güth, Schmittberger, and Schwarze (1982); Kahneman, Knetsch, and Thaler (1986); Rabin (1993); Fehr and Schmidt (1999); Bolton and Ockenfels (2000)) that leads them to be willing to split surplus in a manner more equitable than predicted by subgame perfection. This could lead to buyers being willing to accept price increases even when sellers have no bargaining power. If subjects possess such preferences and they are active in this environment then renegotiation might actually allow for post-auction price corrections that allow sellers to escape the consequences of the Winner's Curse. Similarly, sellers with bargaining power may choose not to exploit that bargaining power over the buyer so that the buyer does not end up as disadvantaged when facing such a seller in renegotiation. The experiment was therefore designed to determine if subjects would respond to the strategic incentives in the situation or if they would exhibit any sort of other regarding concerns

that might produce different results, perhaps more favorable for the prospect of renegotiation.

We find strong evidence that both buyers and sellers integrate the strategic consequences of renegotiation into their behavior and that this dominates any fairness considerations. Consistent with our theoretical predictions, when all sellers are strong (no credible commitment to default), renegotiation is not successful as buyers reject almost all offers knowing the strong sellers can absorb losses. Consequently, bidding behavior and the level of losses at the auction stage are not affected by the presence of renegotiation. In contrast, when wealth constrained bidders with a credible commitment to default are added to the population of sellers, these weak sellers win the majority of all projects at very low prices, and then use the credibility to default as leverage to negotiate a much more favorable price.

We do note that renegotiation is not the only mechanism that has been considered to solve this problem. It is often found that the Winner's Curse is less severe in ascending auctions and one might consider using them to diminish the problem, Levin and Reiss (2013). Another common approach often used in practice involves an auctioneer removing bids from consideration that are deemed too low.[3] Harstad and Rothkopf (1995) examine a similar option in a standard common value auction where winning bidders can withdraw bids. Chen, Xu, and Whinston (2010) suggests that the way to deal with the problem involves contingent contracting. Calveras, Ganuza, and Hauk (2004) propose a solution to the problem of sellers going bankrupt midway through a project which consists of a surety bond that the winning seller must pay for that will compensate the buyer in the event that the seller fails to complete the project. Chang, Chen, and Salmon (2015) and Decarolis (2013) examine the use of an average bid auction in which the winning bidder is the one who bids closest to the average. While there are many methods other than renegotiation in use, in order to determine whether one wants to use one of them instead, it is important to determine the degree to which renegotiation can solve the problem. That is why our goal is to determine the answer to that question as part of the larger agenda in the literature of evaluating mechanisms to find ways to minimize the consequences of the Winner's Curse.

There are also prior studies that examine some of the aspects of the problem addressed in this paper, but none deal directly with the issue of whether renegotiation helps to correct for losses associated with the Winner's Curse in procurement auctions. Waehrer (1995), Roelofs (2002), and

[3] For example, the US Office of Defense for Acquisition, Technology, and Logistics, which is responsible for supervising acquisition for the Department of Defense, has the following policy for unreasonable low bids (FAR 14.407- 3(g)): If the evidence does not reasonably support the existence of a mistake and the contracting officer has determined that the bid price is unreasonable, the contracting officer must reject the bid as unreasonable.
Retrieved from: http://www.acq.osd.mil/dpap/cpf/docs/contract_pricing_finance_guide/vol1_ch8.pdf

Parlane (2003) all examine auctions with default options used by bidders to examine the impact of limited liability on auction outcomes. Shachat and Tan (2014) investigate procurement auctions with renegotiation when sellers have independent private costs, which excludes the possibility of the Winner's Curse. Wealth constrained bidders in the procurement context were examined theoretically by Burguet, Ganuza, and Hauk (2012) and empirically by Krasnokutskaya and Seim (2011).

In the next section we will explain the theoretical environment and the predictions for bidding and renegotiation behavior. Section 3 will describe the design of the experiment. Section 4 will present the results and we conclude in Section 5 with a discussion of what the results from this study suggest regarding the use of renegotiations following procurement auctions.

2 Theory

We analyze a procurement auction with n sellers bidding for a project offered by one buyer. The model is constructed around experimental treatments which investigate the effects of post-auction renegotiation between the winning seller and buyer, and bidders with wealth asymmetry. The environment we use is based on the private plus common value structure introduced in Goeree and Offerman (2003) and Goeree and Offerman (2002). Translated into the procurement setting, this structure assumes that the total cost facing a seller, c_i, is composed of both private and common costs. We assume that the private portion of a seller's cost, t_i, is drawn independently for each of the n bidders from a common distribution. Each bidder is also assumed to receive a signal, τ_i, which is independently drawn from another common distribution related to the common cost element. A bidder's total cost is equal to their private cost plus the common cost, τ, which is equal to the average of all of individual signals. This can be represented as

$$c_i = t_i + \sum_j \frac{\tau_j}{n} \tag{1}$$

For the experiments we assume $n = 4$, while t_i is drawn from a uniform distribution over the range $[0, 50]$ and τ_i is drawn from a uniform distribution on the range $[0, 100]$. We further assume that the value of the buyer is high enough such that they would be willing to buy under any cost realization, or $v_b = 150$.

The auction structure we investigate is the standard low price procurement auction in which

sellers simultaneously submit a bid, b_i, and the lowest bid wins and sets the auction price, $p^* = \min_j\{b_j\}$. The earnings of the winner are equal to the auction price less the true cost or $p^* - c_i$, and the buyer earns $v_b - p^*$. In some treatments we will add a renegotiation stage which follows the auction. In the renegotiation stage we allow the winning seller to observe the true common cost, $\sum_j \frac{\tau_j}{n}$, and then make a take it or leave it offer of a new transaction price, f, to the buyer. If the buyer accepts the offer then that offer becomes the new transaction price and the buyer will earn $v_b - f$ while the winning seller receives earnings of $f - c_i$. If the buyer rejects, then the transaction price remains the price which resulted from the auction. When making his decision, the buyer will not be able to observe the total cost of the seller, but he will know the distribution of possible costs.

We will assume the existence of two types of bidders - strong and weak. Strong bidders are assumed to be able to absorb losses and so any losses they make are simply represented as negative earnings. In contrast, weak bidders will be assumed to default upon suffering a loss of any size. This default will result in a penalty to the seller, k, and recontracting costs, r, to the buyer. For the experiment we assume that the default penalty suffered by the seller is $k = 50$ and the recontracting cost for the buyer is $r = 25$. During the auction, the number of weak and strong sellers is common knowledge and in the event of a renegotiation phase, a buyer knows whether or not they face a strong or weak winning seller.

In order to understand the effects of allowing renegotiation we must analyze this model under conditions where renegotiation is and is not allowed. Similarly, in order to understand how the asymmetric liability of sellers affects outcomes we also need to analyze the model in the presence and absence of the weak sellers. This leads to 4 different scenarios which we address in turn to provide a characterization of the bidding behavior in each case.

2.1 Baseline

The baseline condition refers to the case in which only strong sellers are included in the auction and renegotiation is not allowed. The bidding behavior in this environment can be derived in a straightforward manner based on Goeree and Offerman (2003) and we will mimic the notation used in Milgrom and Weber (1982). We first construct a combined cost signal of $x_i = t_i + \tau_i/4$ which represents the information available to seller i about his final cost when placing his bid. As all bidders are symmetric strong bidders, in this environment we will focus on bidder 1 who has a combined cost signal $x_1 = t_1 + \tau_1/4$. If we let X_1 be equal to the expected minimum combined

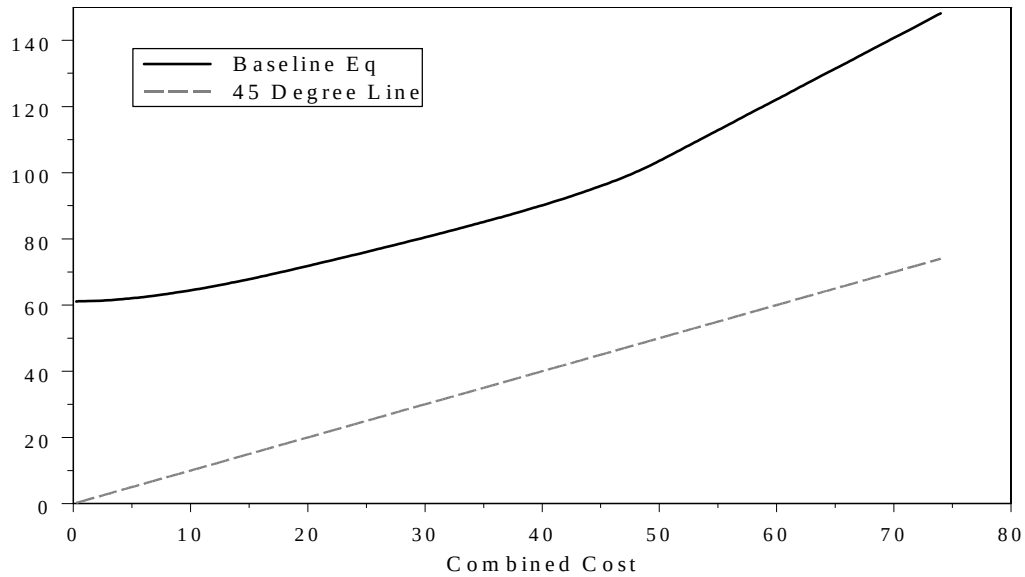

Figure 1: Equilibrium bid function for Baseline and Renegotiation treatments.

signal out of n draws and Y_1 be equal to the expected minimum combined signal out of $n-1$ draws that are above X_1, then the equilibrium bid function can be expressed as

$$b(x_i) = E(\tau + t_i | x_1 = x_i, X_1 = x_i) + E(Y_1 - X_1 | x_1 = x_i, X_1 = x_i). \tag{2}$$

The first component is an estimate of the cost to bidder i, assuming his signal is the lowest, while the second represents an estimate of the difference between the expected lowest and second lowest combined signals given that bidder i has the lowest signal.

Figure 1 graphically illustrates the baseline bid function, given the distributional assumptions used in the experiment.[4]

2.2 Renegotiation

We now add the possibility of renegotiation after the auction between the winning seller and buyer. We maintain the assumption of the baseline that all sellers are strong types who can absorb losses. As this is a two stage game with imperfect information, we use backward induction to solve for a weak perfect Bayesian Nash equilibrium. In the second stage, we can assume that the auction has concluded yielding some price, p^*. While the buyer does not observe the seller's cost, he knows that all sellers are strong which implies that the probability of default for any winning seller is 0 at any price. Thus, any belief structure over costs yields equivalent decisions by the buyer. If the seller

[4] The full derivation of the baseline bid function with the distributional assumptions can be found in the appendix.

requests any price $f > p^*$, then the buyer can accept and pay a price of f or reject and pay the lower price of p^*. It is clearly a best response to reject any offer of $f > p^*$. A seller could in principle request some $f < p^*$, which the buyer would accept, but this is clearly dominated behavior for the seller. The result of this analysis is that we observe there to be an infinite set of equilibria which could involve the sellers making any offer $f \geq p^*$ or never making an offer but in all cases the buyer rejects any offer above the auction price. In all possible equilibria, the final transaction price stays at the auction price, p^*.

Since the final transaction price will always stay at the level set by the auction, there is no change in the incentives during the auction stage. The equilibrium bidding behavior is therefore the same between the baseline and renegotiation cases. This is an important result because it demonstrates that the only way renegotiation can affect the auction outcome is when the final transaction price might change during the bargaining stage which requires the winning seller to have some form of bargaining power.

2.3 Wealth

Our third case removes the option of renegotiation and introduces asymmetric bidders. Half of the bidders are strong sellers who can absorb all losses, and the remaining half are weak sellers who will default if they make losses in the auction stage. Defaulting forces the weak seller to accept a penalty of k and imposes recontracting costs, r, on the buyer. We assume that k is larger than the loss from the auction, so defaulting does not serve as a form of limited liability.[5] In this case, default by the weak seller substantially increases the potential damage from underestimating the cost of the project as any bid, even slightly under the true cost, will lead to a substantial loss. In contrast, strong sellers do not face the threat of additional costs when the true cost is misjudged. This asymmetry in liability constraints leads to a change in the incentives underlying the bidding behavior. Weak sellers essentially bid as if they possess a loss averse utility function, as they must bid to avoid any losses, while strong sellers do not.

Solving for closed form bid functions for both strong and weak sellers is not tractable for this case. We have chosen instead to construct the bid functions using a computational method. This approach involves calculating the expected utility to a seller from following a specified function for their type using a simulations method assuming that the other three sellers are using the proposed

[5] The default penalty, k, can be interpreted as the closure of the firm, or the present value of reputation loss in future transactions as a firm who defaulted.

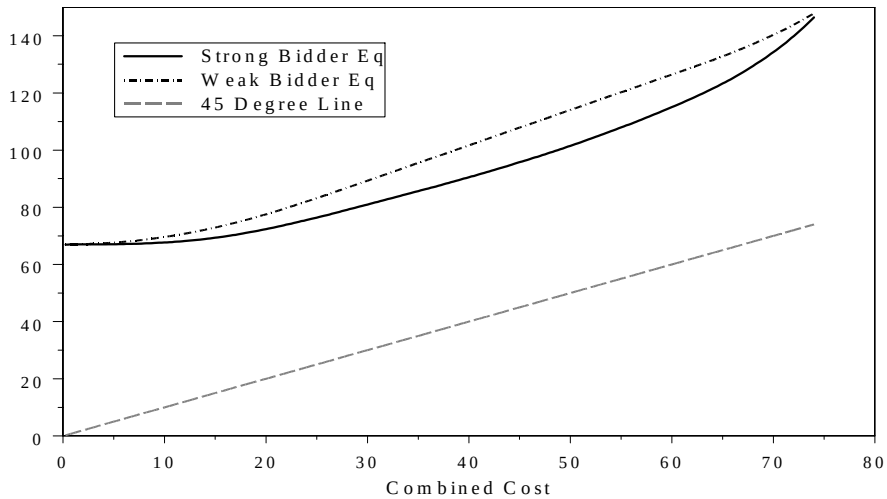

Figure 2: Equilibrium bid functions for strong and weak bidders in Wealth treatment.

bid functions for their type. We then determine whether there are profitable deviations from the proposed function for a seller of either type by calculating the best response bid function for both types. If profitable deviation for either type is possible, a new function can be proposed and tested using the same criteria.[6] This process iterates until we find a fixed point for both functions, i.e. a set of functions that are best responses to themselves.[7]

Figure 2 graphically displays the bid functions for this case. The qualitative nature of these bid functions is quite intuitive. The weak sellers bid less aggressively than the strong sellers to avoid the possibility of even small losses. The strong sellers are able to bid more aggressively than the weak bidders as they do not suffer from the same serious consequences related to small losses. The strong bidders here are bidding less aggressively than the strong bidders in the baseline case where there are no weak bidders but they bid more aggressively than the weak bidders. The effect of the difference in behavior between strong and weak bidders is that at equivalent cost draws strong sellers will be advantaged over weak, allowing them to win more often.

[6] Bid functions were approximated using splines over the point estimate of the best response that resulted from the simulation process. While the spline process was relatively flexible, we did impose shape primitives on the process which are based on the theoretical assumptions for these bid functions. Specifically, we estimated a monotonically increasing polynomial where the bid for the maximum combined cost signal is equal to 150, the highest possible cost, $b(\overline{x}) = \overline{c}$. A detailed algorithm is available from the authors upon request.

[7] This approach differs from a more standard computational approach used in asymmetric auctions that involves numerical solutions to the systems of ordinary differential equations (ODE) defined by the first order conditions. This was necessary because the penalty in the case of default is constant leading to a discontinuous utility function and so we are only able to isolate $b'_w(x_i)$ for the case of no default (positive earnings).

2.4 Wealth and Renegotiation

The final environment combines both renegotiation and asymmetric bidders. The presence of weak sellers changes the outcome from what is observed under renegotiation without asymmetry due to the fact that the weak bidders have a fully credible commitment to default if they end up losing money at the conclusion of all final transactions. The equilibrium concept necessary for this environment is a weak perfect Bayesian Nash equilibrium as we must specify the beliefs held by the buyer regarding the probability with which the seller they are facing will default at the transaction price from the auction.[8] Let $h \in \{w, s\}$ define the type of the seller and let $\rho(p^*, f, h)$ be the belief on the part of the buyer regarding the probability with which the seller they are facing will go bankrupt if the final price is the winning bid observed in the auction, p^*, given the price requested in renegotiation, f, and the type of the seller, h. We can easily construct one very broad class of equilibria in which the equilibrium strategies for the bidders can be characterized by:

$$b_w^*(x_i) = 0 \tag{3}$$

$$f_w^*(c_i, b_i) = 175 \tag{4}$$

$$b_s^*(x_i) > 0 \tag{5}$$

$$f_s^*(c_i, b_i) \geq b_i \tag{6}$$

Note that the offers made by the sellers in the renegotiation stage can be conditioned on their true cost, c_i, rather than just their combined cost signal, x_i, due to the fact that the sellers learn their true cost at the conclusion of the auction and prior to making the renegotiation offer. The choice for the buyers in the renegotiation stage, $D \in \{Accept, Reject\}$, must be accompanied by an appropriate set of beliefs regarding the probability with which they believe a seller will default

[8] Technically the beliefs should be over the type or cost of that seller with those beliefs about type determining the beliefs about the probability of default. It is notationally more convenient to make the beliefs directly about the default probability and so we have chosen to represent them that way. Since the beliefs about type and default probability have a direct relationship between them, this choice has no important consequences.

after a rejection. Both can be characterized as follows:

$$D(p^*, f, h) = \begin{cases} Reject & if \quad h = s \ \& \ f > p^* \\ Accept & if \quad h = s \ \& \ f \leq p^* \\ Accept & if \quad h = w \ \& \ p^* = 0 \ \& \ f \leq 175 \\ Reject & if \quad h = w \ \& \ p^* > 0 \ \& \ f > 0 \end{cases} \tag{7}$$

$$\rho(p^*, f, h) = \begin{cases} 0 & if \quad h = s \\ 1 & if \quad h = w \ \& \ p^* = 0 \\ 0 & if \quad h = w \ \& \ p^* > 0 \end{cases} \tag{8}$$

The justification for this set of strategies and beliefs being an equilibrium is straightforward. In equilibrium, all weak sellers should bid 0 to maximize their chances of winning in the auction stage regardless of their observed combined signal. In the renegotiation stage, the buyer will expect that a weak seller who has bid 0 will default with probability 1 if the offer is rejected.[9] Recall that in the event a seller defaults, the recontracting cost for the buyer is $r = 25$ and the value of the completed project to the buyer is $v_b = 150$. Consequently a buyer will expect to receive utility of -25 for rejecting an offer of the weak seller or $150 - f$ for accepting. A rational buyer accepts so long as $150 - f \geq -25$ or $f \leq 175$. A weak seller therefore maximizes his earnings by setting $f = 175$.

We must define beliefs and actions for the buyer off the equilibrium path, i.e. for when a weak seller is observed to bid a positive amount. Since it is impossible for a weak seller to ever do better than a price of 175 it matters little what these off equilibrium path beliefs and actions are so we will assume that the buyer assumes the bid was perfectly revealing of true cost and that the seller will never go bankrupt leading to a reaction of always rejecting. Any alternative belief structure works as well, including one in which the buyer accepts any offer less than or equal to 175, regardless of the auction price.

Given that weak sellers will bid 0, a strong seller can never bid in such a way so as to profitably win the auction. The best response by the strong seller is to bid anything greater than 0 to ensure losing. In the off-equilibrium case that the strong seller wins, again any renegotiation offer they request will be rejected by the buyer if it is above the auction price as the buyer knows with

[9] It is technically possible, given the discrete distributions used in the experiment, for a subject to have a true cost equal to 0 and therefore not default at that price. This involves receiving $t_i = 0$ and then $\tau_j = 0$ for all j. The probability of that event is equal to $\frac{1}{51} * \frac{1}{101}^4 = 1.88 \times 10^{-10}$. We feel justified in the choice to round this to 0 for convenience.

certainty that the strong seller will not default. Of course the buyer will happily accept an offer below the auction price though the seller should never offer that. Consequently the strong seller can request anything greater than or equal to the auction price and it will be rejected by the buyer.

While the environment we have set up with weak sellers that default on any loss is an extreme case, the general nature of the equilibrium structure will hold for different ways of specifying the default threshold. As a weak seller can demonstrate a credible commitment to default with a very low winning bid, he will have leverage in post-auction renegotiations. This is something that should be possible for most alternative specifications of weak sellers since realistic cost distributions should be bounded away from zero. In the conclusion we will return to a discussion of these issues including discussing the degree to which practices like excluding abnormally low bids can eliminate this possibility. Before determining if such practices might work we must first determine if a problem exists in their absence and so we specify and test this more extreme environment.

3 Experimental Design

The experimental design consists of four treatments which correspond to each of the theoretical environments described above: Baseline (B), Renegotiation (R), Wealth (W), and Wealth & Renegotiation (W&R). Each session of the experiment consisted of a single treatment. At the beginning of a session the subjects were randomly assigned to either the role of buyer or seller and that role assignment remained constant over the entire experiment.

As already noted, the cost environment used a direct translation of the private plus common value structure of Goeree and Offerman (2003) to the procurement setting. The cost draws and rules of the mechanisms are implemented exactly as described in the last section.

We conducted 3 sessions for each treatment yielding a total of 12 sessions. Each session consisted of 30 rounds. To ensure the least amount of changes between sessions, we used the exact same signal draws across all treatments. In the sessions without renegotiation we had 16 subjects in the role of sellers with no humans in the role of the (completely passive) buyers. In the sessions with renegotiation we had 16 sellers and 4 buyers.[10] Sellers and buyers were randomly matched into new groups each period. This matching protocol minimized the potential for reputation building, replicating the static theoretical environment.

As subjects could and are expected to make losses in these experiments, we endowed subjects

[10] One session in the Renegotiation treatment had only 3 buyers and 12 sellers due to a smaller than expected number of subjects showing up that day.

	B	R	W	W&R
Buyer		\$52.35		\$38.53
Seller	\$23.07	\$23.45	\$22.33	\$27.11

Table 1: Average earnings by treatments.

with initial balances to hopefully ensure that they would not have negative cumulative earnings at any point in time. We set the seller endowment at 375 ECUs while the buyer endowment was set at 150 ECUs. As it was still possible, even with those endowments, for individuals to achieve a negative overall balance, we employed some standard rules for dealing with the bankruptcy prospect. The rules allowed for subjects who go bankrupt a single time to be re-initialized with a new endowment. Any subject who went bankrupt a second time was removed from the session and received only their participation fee. In the end, eight subjects went bankrupt once and one subject went bankrupt twice.

Due to the fact that buyers were expected to earn substantially more ECUs than sellers, we used different conversion rates for converting ECUs into USD for the two roles. For those in the role of sellers, 1 ECU=\$0.04 and for those in the role of buyers, 1 ECU=\$0.015. This lead to overall earnings of \$45.73 for the buyers and \$23.93 for the sellers which includes the \$10 participation fee. Table 1 shows the earnings broken down by type and treatment.

3.1 Hypotheses

The preceding theoretical analysis in conjunction with the experimental design sets up a series of hypotheses that can be empirically tested. In this section, we briefly state our primary hypotheses of interest to guide the data analysis in the results section.

Hypothesis 1 *There will be no observed difference in bidding behavior or auction outcomes between the B and R treatments.*

This first hypothesis proceeds from the fact that all renegotiation requests from strong sellers are expected to be rejected. It is certainly also possible though that a different behavioral norm could emerge based on fairness. Other prior studies also suggest that some individuals have a bias towards honesty in their behavior (Charness and Dufwenberg (2006); López-Pérez and Spiegelman (2013); Fischbacher and Föllmi-Heusi (2013)). If these behaviors transfer into this environment, then it is possible that sellers bid in good faith during the R treatment (i.e. they bid no differently from the B treatment) and then only request renegotiation when they find out that they are losing

money. The buyers might then decide, since they value the welfare of the other, to accept the renegotiation offer. The development of such norms is what would be necessary for renegotiation to help mitigate losses for the winning seller stemming from the Winner's Curse.

Hypothesis 2 *Strong sellers will bid more aggressively than Weak sellers in the W treatment and this will be reversed in the W&R treatment.*

Hypothesis 3 *Strong sellers will win more often than weak sellers in the W treatment while weak sellers will win more often in the W&R treatment.*

The support for both of these hypotheses falls directly out of the predictions of bidding behavior in the two treatments.

Hypothesis 4 *All renegotiation offers by Strong sellers will be rejected in both R and W&R treatments. Renegotiation offers by Weak sellers will be more likely to be accepted at lower auction closing prices.*

According to the equilibrium structure, all renegotiation requests by strong sellers should be rejected. All weak sellers should bid 0, request a renegotiated price of 175, and then the buyer should accept. In regard to the offers by strong sellers, there is no self-interested rationale for why buyers would ever accept requests above the auction price. However, if weak sellers deviate from the precise equilibrium, buyers may have reason to reject offers. The reason is that if an auction closes at a non-zero price, it is not guaranteed that the seller would default should the renegotiation request be rejected. The higher the close price is above 0, the more likely it is that a renegotiation request might be safely rejected. Therefore, in our hypothesis we have accounted for how buyers may respond if weak sellers do not make bids/offers in exact accordance with the theoretical prediction.

Hypothesis 5 *Final transaction prices will be rank ordered such that $P_{W\&R} > P_W$ and $P_B = P_R$. Seller losses should be ordered in the reverse.*

Of ultimate importance to those involved in these types of auctions is this last hypothesis which ranks the final transaction prices paid by buyers and the expected losses to sellers. According to the equilibria derived above, the W&R treatment clearly delivers the highest price and the B and R treatments are expected to be equivalent. According to the equilibria, there should be no losses

in expectation among the sellers. Of course, from many prior experiments with common/affiliated values environments (e.g. Kagel and Levin (1986); Kagel and Richard (2001)) we know that individuals are highly unlikely to bid exactly as theory predicts and that losses are likely. If the price ranking remains robust to the new bidding behavior and if there are losses, they should be inversely related to the final prices.

4 Results

4.1 Summary Statistics

We begin by displaying summary statistics of the data from the experiments. The hypotheses stated above will be formally tested in the next sections in a series of regressions but prior to seeing the regression analysis, it is informative to examine the summary statistics to get an intuitive idea for how the treatments compare.

Table 2 contains summary statistics on revenue, efficiency, and buyer surplus. For revenue we provide observed and predicted revenue based on the actual cost draws, rather than the unconditional expected revenue, as a way of determining whether the minor differences in observed revenue could just be due to similar minor differences in actual cost draws. We include average auction prices in addition to average final prices for the R and W&R treatments to demonstrate how much of the change in observed revenue is due to renegotiation. Auction and final prices are the same for the B and W sessions since renegotiation is not possible.[11] Also for the B treatment, the Auction and Final surplus are the same since the sellers cannot default in the event of a loss. This is not so for the W treatment.

By hypothesis 1, auction prices are expected to be equivalent between the B and R treatments because renegotiation should never be successful. We see that the average auction price in R (59.38) is approximately equivalent to B (61.35). Auction prices in the W&R treatment are theoretically predicted to equal 0, and while we do not observe prices at the zero level, we do find that auction prices in this treatment are substantially lower than the W treatment and all other treatments.

Comparing auction prices to final prices for the treatments with renegotiation, there appears to be little difference between the auction price and the final price for the R treatment. For the W&R treatment we see a large difference between the average auction price (27.95) and the average

[11] In the case of default, the final price is recorded as the last price from the auction if there was no renegotiation, or if renegotiation was refused.

final price (79.10) and that average final price is higher than the average final price in all other treatments, even though it is still lower than the theoretical prediction of 175. This suggests that renegotiations were more often successful in W&R than in R.

Table 2 also includes measures of efficiency. The possibility of seller default in this environment generates two types of efficiency - auction efficiency which is the equal to the fraction of the total possible surplus that was realized at the end of the auction, and final efficiency which is measured in an identical manner, but accounts for situations where the weak seller defaulted and no surplus is realized.[12] The auction efficiencies in the B, R, and W treatments are all similar, scoring about 90% efficient. This is higher than 82% which is what would have been achieved had the auction been won by a randomly selected seller.[13] The auction efficiency achieved by the W&R treatment (87%) appears lower than the other treatments, though it is still higher than what would be generated by random assignment.

In the treatments with asymmetric bidders (W and W&R), weak sellers default in the event of a loss which may lead to a different final efficiency because under default the realized surplus and resulting final efficiency are zero. In the W treatment, renegotiation cannot help the weak seller overcome auction losses and as a result this drastically lowers efficiency to 74% from 91% achieved at the auction stage. While renegotiation was possible in the W&R treatment, the final efficiency was still lower than the auction efficiency due to the fact that some renegotiation offers were rejected leading to default.[14]

How well the buyers fared in each treatment is summarized by buyer surplus. At the auction stage this is measured as the value of the buyer less the price paid in the auction. The buyer's final surplus accounts for changes due to renegotiation and default. If a new price was renegotiated, the surplus is updated using the new transaction price and in the case of default the buyer's surplus is set to -25, the penalty imposed on the buyer. Of course, the buyer's surplus from the auction is inversely related to the auction price. The final surplus is of interest on top of the auction price because it incorporates both renegotiated changes in price, and default. It is evident that the final surplus of the buyer is substantially lower with asymmetric wealth constraints (due to occasional

[12] The buyer's value is 150, so the maximum possible surplus is equal to 150 minus the lowest cost realization in the bidding group. The surplus achieved is equal to 150 minus the cost of the winning seller. In the case of default, the calculation of final efficiency assumes the achieved surplus is 0.

[13] This random efficiency is high because of the substantial common value element to the value structure.

[14] Final efficiency changes based on the frequency of default for weak sellers. Out of 240 auctions in the W&R treatment where weak won, 27 (11.25%) resulted in losses and default after renegotiation. In the W treatment, weak won in a 139 cases, 72 (52%) of which resulted in losses and default. Even though default was only 11% of outcomes under W&R, the complete loss of any realized surplus is substantial enough to dramatically decrease efficiency.

Treatment	Price (Predicted)		Efficiency			Buyer Surplus	
	Auction	Final	Auction	Final	Random	Auction	Final
B	61.35 (77.31)		0.90		.82	88.65	
R	59.38 (77.56)	60.79 (77.56)	0.90		.82	90.62	89.21
W	62.35 (76.82)		0.91	0.74	.82	87.65	64.24
W&R	27.95 (0)	79.10 (175)	0.87	0.80	.82	122.05	58.47

Table 2: Average price, efficiency and buyer surplus by treatment.

Treatment (n auctions)	Strong Wins	Weak Wins	Strong Offer	Weak Offer	Strong Accepted Offer	Weak Accepted Offer
R (n=356)	100% (356)	-	79.62 (331)	-	67.73 (83)	-
W (n=360)	61% (221)	39% (139)	-	-	-	-
W&R (n=330)	27% (90)	73% (240)	81.13 (89)	102.36 (239)	74.33 (12)	100.05 (219)

Table 3: Propensity of strong/weak sellers to win auction. Average offers made and accepted offers by type and treatment (frequencies in parentheses below).

defaults) and there is no indication that renegotiation is helpful to the buyer.

Table 3 provides summary statistics on the propensity of each type of seller to win in each treatment as well as average offers and average accepted offers. As predicted in hypothesis 3, strong sellers win more often than weak sellers in the W treatment (221 versus 139, or strong sellers win 61% of the time) while the weak sellers win more often in the W&R treatment (90 versus 240, or weak sellers win 73% of the time). Turning to renegotiation, in the R treatment we note that the average offer made by strong sellers, 79.62, is higher than the average accepted offer, 67.73. The average accepted offer, 67.73, is not much higher than the average auction price for B, 61.35, presented in table 2. In contrast, the average offer made by weak sellers in the W&R treatment, 102.36, is roughly equivalent to the average accepted weak offer and much higher than the average auction price (27.95). By hypothesis 4, all strong offers should be rejected while weak offers are more likely to be accepted. In the W&R treatment, buyers accepted 92% of the offers made by weak sellers, whereas only 13% of the offers (12 out of 89) were accepted from strong sellers. In R, strong sellers are slightly more successful in getting their offers accepted as 25% (83 out of 331) of strong offers are accepted by the buyer.

Losses may be excessive at the auction stage in B, R, and W, if bidders bid too low, failing to account for the adverse selection issue associated with winning. In the W&R treatment, we also expect losses at auction prices except that in this case it's not because of the Winner's Curse,

		Auction Prices		Final Prices	
		Strong	Weak	Strong	Weak
Earnings	B	-8.72	-	-8.72	-
	R	-11.31	-	-9.91	-
	W	-7.27	-5.56	-7.27	-19.27
	W&R	-31.92	-49.93	-27.48	19.96
Losses	B	-20.58	-	-20.58	-
	R	-23.63	-	-25.47	-
	W	-21.02	-23.54	-21.02	-50
	W&R	-35.01	-52.67	-34.66	-50

Table 4: Average earnings and losses by winning seller.

		# Losses at Auction Prices		# Losses at Final Prices	
		Strong	Weak	Strong	Weak
All Periods	B	229 (64%)	-	229 (64%)	-
	R	236 (66%)	-	216 (61%)	-
	W	129 (58%)	72 (52%)	129 (58%)	72 (52%)
	W&R	83 (92%)	229 (95%)	76 (84%)	27 (11%)

Table 5: Number of times winning bidder made losses by treatment. (Percentage of auctions in which winning bidder made losses.)

but because bidders are competing to win the right to renegotiate with no regard to their actual cost. To learn how severe the losses associated with low strategic bidding and the Winner's Curse are in the data, table 4 shows the average earnings per round, broken out by seller type and treatment. It contains average total earnings as well as the average size of loss which is conditioned on negative earnings. Table 5 then shows the propensity of sellers to make losses again broken down by seller type and treatment. The very clear picture is that the problem of seller losses is severe in this environment. In all treatments both types of sellers have negative earnings conditional upon winning an auction. In the W&R treatment winning bidders are making losses at auction prices at least 92% of the time while in the other treatments the winning sellers lose money approximately 50-60% of the time. This clearly demonstrates that a viable renegotiation opportunity generated substantially higher losses at the auction stage. Turning to final prices, the only class of sellers that have positive expected earnings are weak sellers in the W&R treatment, who make losses 11% of the time.

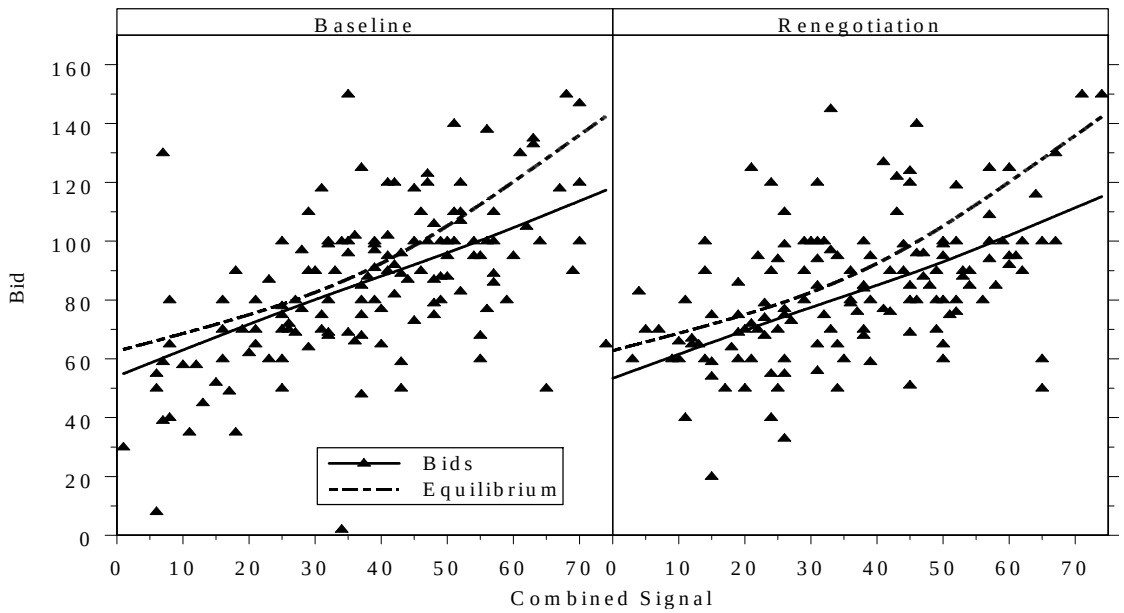

Figure 3: Bids plotted against combined signals for strong sellers in B and R treatments.

4.2 Analysis of Bidding Behavior

We begin the analysis of sellers' bidding behavior by examining Figures 3 and 4, which are scatterplots of bids versus the combined cost signals. The figures include regression plots through the observed bids and lines indicating the relevant equilibrium bidding behavior. Figure 3 provides the plots for the strong sellers in the B and R treatments. Theoretically, the bids are identical for strong sellers between these two treatments and visually the behavior is quite similar. However, in both treatments, sellers are bidding more aggressively than predicted.

Figure 4 examines the behavior in the asymmetric wealth treatments, W and W&R, for both strong and weak sellers. Theoretically there are significant differences predicted in bidding behavior between the treatments. In the W treatment, strong sellers are predicted to be more aggressive than weak and indeed the regression line through those points is lower than that for the weak sellers. The opposite is predicted in the W&R treatment and we observe the predicted switch in regression lines. We do not represent the equilibrium in the W&R panel as there is no distinct prediction for the strong sellers; the strong sellers are simply predicted to bid above the weak sellers. The weak sellers are predicted to bid 0 which clearly only a few do.

While the figures help provide an intuitive idea for how the sellers submitted bids, statistical tests for the properties noted above can be found in Tables 6 and 7. Table 6 contains three regressions based on the data from the B and R treatments to examine the differences between

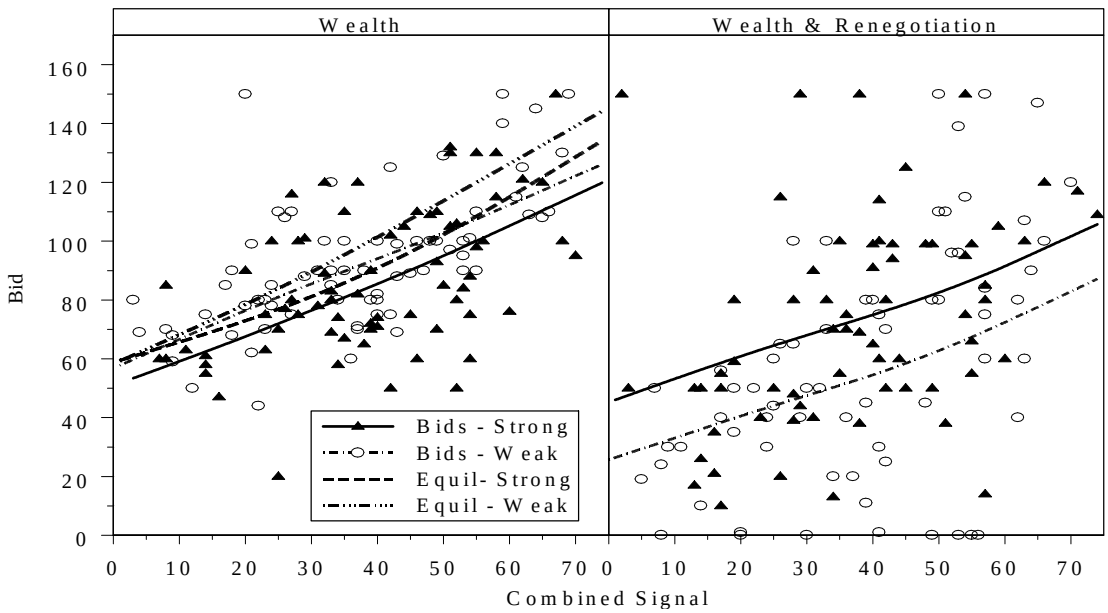

Figure 4: Plots of bids versus combined signals for weak and strong sellers in the W and W&R treatments.

those two treatments. All regressions use the bid submitted by the seller (all are strong in those treatments) in a round as the dependent variable and have standard errors clustered by subject. The first two specifications are random effects panel regressions and test whether there is a difference in behavior between when renegotiation is and is not present. We also include a dummy variable, L, to capture situations in which sellers have low balances, i.e. earnings balance less than 50, to see if they bid differently from cases when they do not have low balances. The last specification presents a test of whether observed bidding behavior differs from equilibrium behavior in a pooled OLS regression, used to eliminate the constant from the regression.

Result 1 *There is no statistically significant difference in bidding behavior between the B and R treatments.*

The support for the first result is provided by the fact that in models 1 and 2 the dummy variable for the renegotiation treatment, R, is insignificant. We also have provided several interactions between that and other variables and these too are mostly insignificant. We therefore fail to reject hypothesis 1.[15]

[15] We have re-run models 1 and 2 from Table 6 with the data restricted to the last 15 periods to analyze any learning effects. We find no qualitative differences in the restricted models; we also do not observe large changes in the coefficients, except for the case of $R * L$ which decreases to -69.63 ($p = 0.007$), indicating that the low balance effect is less prominent in later periods.

	(1)	(2)		(3)
Combined cost signal, x_i	0.817***	0.819***	Equilibrium Bid, b_i	0.914***
	(0.034)	(0.046)		(0.019)
Renegotiation, R	-2.595	-2.251	$b_i * R$	-0.027
	(2.347)	(3.498)		(0.024)
$R * x_i$		-0.002		
		(0.068)		
Balance<50, L		14.48		
		(17.66)		
$R * L$		-81.93***		
		(21.97)		
Constant	55.52***	55.35***		
	(2.236)	(2.516)		
Observations	2,864	2,864		2,864
Clusters	96	96		96

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

Table 6: Random effects regressions with bid as dependent variable using data from B and R treatments only.

Result 2 *Bidding behavior in both the B and R treatments is more aggressive than predicted.*

We present model 3 to test the relationship between theory and what sellers actually bid. We use the theoretically predicted bid as an independent variable rather than the signal and include an interaction with the treatment. If subjects were bidding as theory predicted, the coefficient on the variable b_i would be equal to 1 and the interaction would be equal to 0. While the coefficient on the interaction is not significantly different from 0, the coefficient on b_i is less than 1 and is significantly different from 1 ($p < 0.001$) as is the linear combination of the two which captures the total effect for the Renegotiation treatment ($p < 0.001$). This provides support for result 2.

We predict in hypothesis 2 that bidding behavior should shift between the W and W&R treatments. In the W treatment, weak bidders should bid less aggressively than strong. In the W&R treatment, this relationship is predicted to reverse. We provide tests of the comparative static predictions and also a general test of how well the theory explains the behavior in the W treatment.

Result 3 *In the W treatment, strong sellers bid more aggressively than weak sellers. In the W&R treatment, weak sellers bid more aggressively than strong sellers.*

The statistical support for result 3 is derived from the first two columns in Table 7. Column 1 provides a regression specification looking at how bids in the W treatment depend on the combined cost signal and on type. Column 2 provides the same specification for the W&R treatment. We

find that the dummy variable for the strong seller is negative and significant in the W treatment specification and positive and significant in the W&R treatment. This is the comparative static predicted by hypothesis 2.[16]

Result 4 *In the W treatment, sellers of both types bid more aggressively than predicted.*

To examine whether we observe behavior corresponding to the specific predictions of the theory, for the W&R treatment it is informative enough to examine the right side panel in Figure 4. We have no specific prediction for strong sellers in that treatment but the weak sellers are predicted to bid 0. They clearly do not, though as we just saw, they do bid more aggressively than the strong sellers. So while the specific prediction fails, the comparative static still holds. For the W treatment we do have specific predictions of bid functions for both types and regressions testing conformance with those bid functions are contained in columns 3 and 4 of Table 7. These are again pooled OLS regressions with the standard errors clustered on the subjects without an intercept term. Similar to what was observed in the B and R treatments, we find that bidders of both types bid more aggressively than predicted as the coefficients on the predicted bid are significantly different than 1 (for both $p < 0.001$). This provides the support for result 4.

4.3 Analysis of Renegotiation Behavior

The renegotiation stage involves a take-it-or-leave-it offer between the winning seller and the buyer. Theoretically, all offers made by strong sellers should be rejected as these sellers can absorb any losses made in the auction. In contrast, when renegotiation is allowed, weak bidders are predicted to bid zero which almost always implies losses at the auction stage and as a result any offer made by weak sellers less than or equal to 175 will be accepted by buyers to avoid default by the weak seller. As seen in the previous analysis of bidding behavior, zero bids by weak sellers are infrequent, and so as stated in hypothesis 4, we expect that a weaker property will hold which is that a buyer should be more likely to accept an offer if the auction price is lower.

A first look at the accept/reject decisions is seen in Figure 5 which provides histograms of all offers by seller type indicating which offers are accepted versus rejected. It is clearly the case

[16] We have again restricted the data to the last 15 periods to re-run models 1 and 2 from Table 7. Similar to the case of the B and R treatments, we find that the qualitative results in the W and W&R treatments hold across time with slight changes in the direction of the theory. In model 1, the difference between average bidding of strong and weak increases - the intercept (average weak bids) increases to 58.59 ($p < 0.001$) and the coefficient on strong shifts to -7.94 ($p = 0.021$). In the restricted version of model 2, the slope coefficient on x_i becomes less steep (.70, $p < 0.001$) and the coefficient on strong increases (23.17, $p = 0.008$).

	W Only	W&R Only	W Only	
			Strong Bidders	Weak Bidders
	(1)	(2)	(3)	(4)
Combined cost signal, x_i	0.937***	0.808***		
	(0.054)	(0.098)		
Strong seller, S	-6.591***	17.08**		
	(2.237)	(6.672)		
$S * x_i$	-0.027	0.053		
	(0.055)	(0.113)		
Equilibrium bid, *bid*			0.926***	0.918***
			(0.020)	(0.019)
Constant	56.42***	23.93***		
	(2.820)	(4.487)		
Observations	1,433	1,320	715	718
Clusters	48	44	48	48

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

Table 7: Random effects regressions with bid as dependent variable using data from noted treatment and bidder type only.

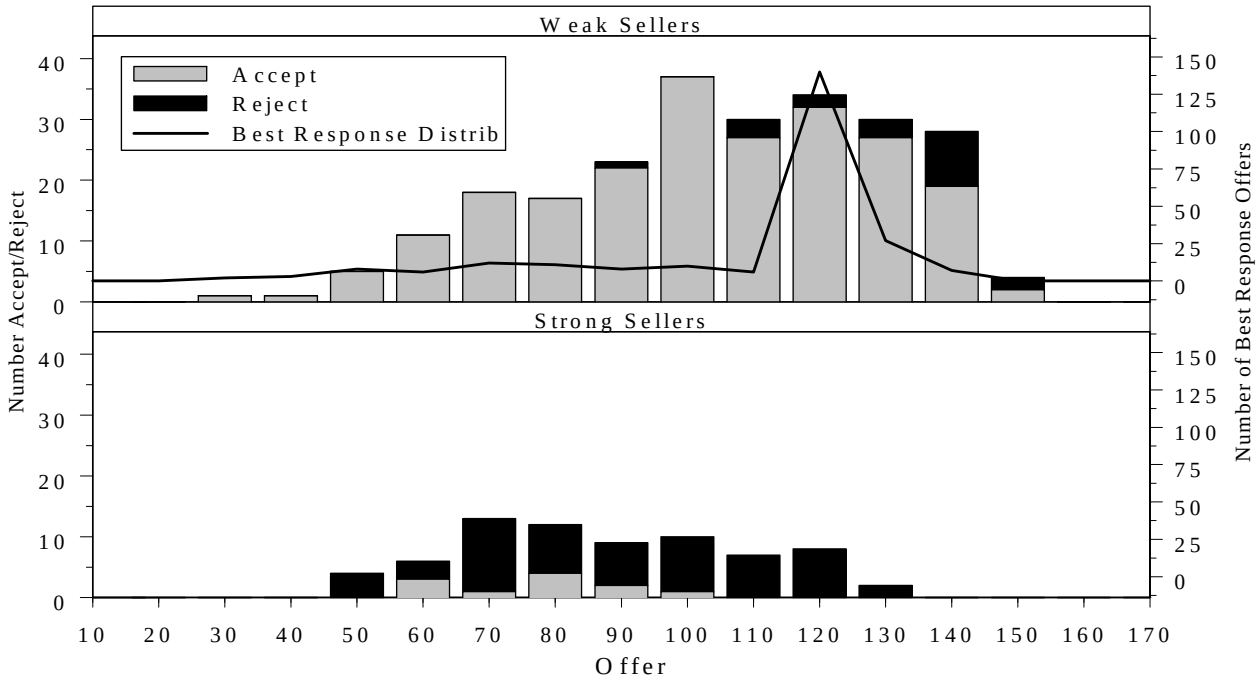

Figure 5: Histograms of offers by seller type.

	(1)	(2)		(3)
Auction price, p^*	0.003	0.008*	Best Response Offer	0.907
	(0.003)	(0.005)		(0.049)
Renegotiation offer, f	-0.011***	-0.014***		
	(0.002)	(0.005)		
Weak seller, WS	0.902***	0.976***		
	(0.036)	(0.027)		
$WS * p^*$		-0.011**		
		(0.006)		
$WS * f$		-0.003		
		(0.005)		
Wealth & Renegotiation, $W\&R$	-0.059	0.023		
	(0.198)	(0.204)		
$f < p^* + 10$	0.197**	0.131		
	(0.075)	(0.094)		
Clusters	23	23		39
Observations	659	659		239

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

Table 8: Marginal effects from probit regressions with acceptance of renegotiation offer as the dependent variable.

that offers by strong sellers are almost always rejected, as predicted, and offers by weak sellers are typically accepted. There is also an indication that offers by weak sellers at high levels do get rejected.

Columns 1 and 2 in Table 8 provide marginal effects from probit regressions with agreement to a new offer as the dependent variable. Standard errors are clustered at the subject level for sellers. In both models we include the auction price, p^*, the offer made by the seller, f, and we include a dummy, $f < p^* + 10$, to account for offers that are near or below the auction price. We also include dummy variables for the treatment and whether or not the proposing seller was weak (WS). The second model includes interactions of the weak seller dummy with auction prices and offers to account for variations in the probability of acceptance due to the type of proposer.

Result 5 *Buyers are more willing to accept offers made by weak sellers than strong.*

Result 6 *Buyers willingness to accept is decreasing in the level of the offer.*

Support for result 5 is derived from the positive coefficient on the dummy for a winning weak seller which is significant across both specifications. As what is shown in the table are the marginal effects, the fact that the coefficient is above 0.90 indicates that this is by far the most substantial determinant in whether an offer is accepted or rejected. This provides robust verification that

buyers are much more willing to accept the offers of weak sellers than they are strong sellers. The support for result 6 is found in the negative and significant coefficient on the offer variable. Thus both parts of hypothesis 4 on renegotiation behavior have failed to be rejected.[17]

The offer/reject behavior we observe is consistent with the comparative statics of the model but there are important deviations from the theoretical predictions. The average offer made by weak sellers is 102.36 (see table 3), which is relatively far from the predicted offer of 175, or even the value of the buyer 150. This is a substantial deviation from the theoretical prediction and it is worth investigating to determine why sellers are making such low requests. The key question is whether the sellers are making requests more generous (i.e. lower) than they have to, perhaps due to fairness considerations, or are the buyers rejecting more high offers than they should according to the subgame perfect prediction which causes the sellers to best respond with lower price requests. Figure 5 does show that high offers tend to be rejected even from weak sellers and they might be doing this with high enough propensity to make such offers unprofitable.

In order to answer this question we need to understand how a rational seller would best respond to the buyers, given how buyers tend to make accept/reject decisions. Due to the multi-dimensional nature of the problem, Figure 5 does not offer a detailed enough perspective of the rejection probabilities as the buyer sees not just the offer but also the auction close price - both of which may figure into the accept/reject decision. Due to the size and sparseness of the full grid, we cannot directly observe acceptance propensities from the data. We can, however, estimate acceptance probabilities over the entire space.[18] We have done so and present the estimated probability of acceptance for any combination of auction price and offer (assuming a weak seller) in Figure 6. What we see is that despite the subgame perfect prediction, offers over 100 have a rapidly declining probability of being accepted.[19]

Using this probability of acceptance, we can calculate the expected utility of any offer from

[17] Re-running models 1 and 2 in Table 8 with restricted data (last 15 periods), we find that the coefficient on auction price, p^* in model 1 becomes significant ($p = 0.004$) in later periods, but the magnitude of the effect is relatively small (coefficient = 0.008). We also find that the probability a weak proposer is accepted increases to .95 ($p < 0.001$). The restricted results for model 2 are virtually identical to the unrestricted results presented, with one exception - the effect of $WS * p^*$ becomes insignificant ($p = 0.663$).

[18] The logit model is $Accept = \beta_0 + \beta_1 * price + \beta_2 * price^2 + \beta_3 * price^3 + offer$, where $price$ is the winning low bid from the auction, and $offer$ is the proposed offer of the weak seller in renegotiation. The dependent variable was a binary indicating whether or not the buyer accepted the offer. Standard errors were clustered at the individual level.

[19] Of course since these accept/reject decisions were not elicited using the strategy method it is not necessarily the case that this estimated acceptance probability represents the true average strategy of the buyers. The problem is that certain very high and very low offers are not observed with much frequency. On the other hand we do observe a great deal of heterogeneity in the offers which should be enough to indicate that for the bulk of the offer range we do have enough information to make reliable claims over the average strategy of the buyers.

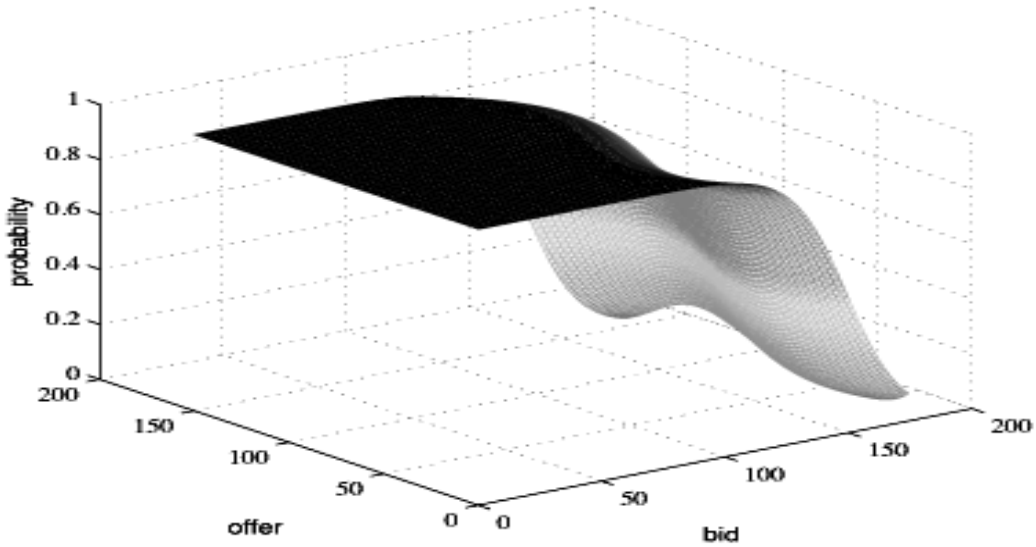

Figure 6: Estimated probability of buyer accepting offer given a closing bid.

the point of view of a seller at a given auction close price and cost realization to determine the empirical best response offer. In Figure 5 we provide one way of summarizing these best response offers as we include a distribution plot of what the optimal offers would have been for the range offers made by weak sellers in the experiments. What we observe is that there is a spike in this distribution indicating that in most cases, the optimal offer is in the range of 120-130 which is well below 175. This spike also roughly corresponds to the peak of the distribution of actual offers, though the "peak" in observed offers is more of a plateau as the bulk of the actual offers, 60%, are in the range of 100-140. In the third column of Table 8 we provide a pooled OLS regression with standard errors clustered on the subject of the actual offer on the best response offer. We find that the actual offer is on average 90% of the best response offer and the difference is marginally significant (test of whether the coefficient is equal to 1 yields a $p-$value of 0.068). While sellers make offers that are less aggressive than the theory predicts, this behavior is roughly consistent with the sellers forecasting the rejection behavior of the buyers and approximately best responding.

4.4 Analysis of Outcomes

In this section, we formally analyze price, efficiency, and buyer surplus outcomes at the auction stage and after any default or renegotiation. Table 9 presents six pooled OLS regressions with standard errors clustered at the session level. The first two specifications focus on auction price

	Price		Efficiency		Buyer Surplus	
	Auction	Final	Auction	Final	Auction	Final
	(1)	(2)	(3)	(4)	(5)	(6)
Lowest Combined Signal	0.344***	0.314***	-0.0007***	-0.003***	-0.344***	-0.602***
	(0.037)	(0.030)	(<0.001)	(<0.001)	(0.037)	(0.104)
Random Efficiency			0.512***	0.404***		
			(0.102)	(0.105)		
Wealth, W	1.048	1.044	0.013	-0.162***	-1.048	-24.49***
	(2.572)	(2.564)	(0.011)	(0.024)	(2.572)	(2.361)
Renegotiation, R	-2.083	-0.665	-0.002	-0.001	2.083	0.767
	(1.751)	(2.101)	(0.013)	(0.012)	(1.751)	(2.246)
Wealth & Renegotiation, $W\&R$	-33.49***	17.67***	-0.030**	-0.098***	33.488***	-30.018***
	(5.847)	(3.790)	(0.012)	(0.016)	(5.847)	(5.183)
Constant	40.22***	42.05***	0.521***	0.74***	109.78***	125.697***
	(3.112)	(2.812)	(0.091)	(0.118)	(3.112)	(6.717)
Clusters	12	12	12	12	12	12
Observations	1,406	1,406	1,406	1,406	1,046	1,046

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1

Table 9: Pooled OLS regressions on noted variables clustered at session level.

and final price. In theory, the auction and final prices should be equivalent between the B and R treatments while the W&R treatment should result in the lowest auction price and the highest final price across all treatments.

Result 7 *Auction and final prices are equivalent between the B and R treatments. Auction prices are lowest and final prices highest in the W&R treatment.*

Result 7 is in favor of the theoretical price predictions given by hypothesis 5 and is supported by the coefficient on the treatment dummy, R, which is found to be not significantly different from the baseline treatment (B) in both the auction price and final price regressions. The W&R treatment is the only treatment which is significantly different from the others which matches the prediction that the final prices from it would be higher than in the other treatments. Prices in W&R are significantly lower at the auction stage than all other treatments and significantly higher after renegotiation than all other treatments.

Models 3 and 4 examine efficiency at the auction stage and final efficiency, respectively. Recall that final efficiency only differs from auction efficiency when bidders default and final efficiency is set to zero. Similar to what was observed for auction prices, we find no significant differences between the B, R, and W treatments for auction efficiency. On the other hand, auction efficiency is significantly lower in the W&R treatment. Default was possible in the W and W&R treatments and

as a result final efficiency changed from the auction stage for these treatments. The final efficiency levels of the W treatment are significantly lower than either the B or R treatments. Renegotiation did not help overcome the efficiency loss from default as the coefficient on the W&R treatment is still negative, significant, and a larger magnitude that what was observed in the auction efficiency regression.

Models 5 and 6 examine the buyer's surplus, which is measured (at the auction stage) as the difference between the buyer's value and the transaction price. The buyer's surplus may change if the winning seller defaults, in which case the surplus is set to -25, or if renegotiation took place and the surplus calculation uses the new transaction price. The buyer's surplus is an important indicator for whether or not the buyer should allow renegotiation in the procurement process. At the auction stage, significant differences only exist between the B and W&R treatments, with W&R yielding much higher potential buyer surplus due to the very low auction prices. In the final surplus regression we find that there is still no difference between the surplus achieved in B and R treatments, as expected, while there is a substantial reduction in the W and W&R treatments below what is achievable in the B treatment. This is of course due to the existence of weak sellers who occasionally default on projects. We can test the difference in the coefficients from model 6 on the W and W&R treatments to find a lack of a significant difference ($p = 0.319$). The indication here is that despite the higher prices paid in the W&R treatment, there is no drop in surplus to the buyers as there is also a drop in the frequency of defaulting sellers. We hesitate to over interpret this finding too much as an indication that renegotiation fails to harm buyers. This finding is to a great extent dependent on how we parameterized the recontracting costs. By changing them we could have likely ended up with a different result here. The important take away is that even with the hefty price increase stemming from renegotiation, the buyer is not necessarily worse off as there is the important trade-off regarding default probability to consider. In this case we see that renegotiation may not hurt buyers as much as indicated by the prices, but neither is it helping as the buyer surplus is not increasing.[20]

[20] Re-running all models in table 9 with restricted data, we find only two alternations to the results with neither being important. In model 3, the already small coefficient on Lowest Combined Signal becomes insignificant ($p = 0.174$), and in model 4, random efficiency as a predictor variable also becomes insignificant ($p = 0.133$).

5 Conclusion

The use of renegotiation following competitive bidding in procurement settings is commonly justified as a way to help alleviate losses and prevent default by winning sellers who were unable to anticipate the full cost when placing bids. However, if bidders anticipate successful renegotiation, there is an incentive to strategically lower bids below standard levels to win which then provides a credible commitment of post-auction default and confers upon the seller a significant advantage in post-auction bargaining. Due to the potential of post-auction renegotiation changing incentives in the auction, the effect of renegotiation should be analyzed carefully to understand whether it will be effective in solving the intended problems.

We use a combination of theory and laboratory experiments to analyze the effects of post-auction renegotiation in procurement auctions with and without asymmetric bidders. Asymmetry is introduced through wealth constraints where sellers differ in their ability to sustain losses without bankruptcy. Weak sellers will default in the event of a loss, imposing a penalty on the buyer and incurring a default penalty. Strong sellers will never default, absorbing all losses.

Our experimental results provide strong support for the comparative static predictions of the theory. First, in an environment with no wealth constraints, renegotiation is almost never successful as the majority of offers made by strong sellers are rejected by buyers. As a result, bids placed by strong sellers do not adjust to the presence or absence of a renegotiation option. In contrast, when wealth constraints are added to the environment, weak sellers who cannot absorb losses are shown to shift their behavior substantially between situations in which renegotiation is versus is not present. As predicted, buyers accept the majority of renegotiation offer requests by weak sellers. This high acceptance rate encourages weak sellers to bid low to win and then use their credibility of default to obtain substantial price concessions from the buyer. The end result is that weak sellers are substantially advantaged in situations in which renegotiation is allowed. The upside for buyers is that the sellers default substantially less often and depending on the costs of default that might be enough to offset the price increase. Such a determination depends on the utility function of the buyer. On the other hand, all price formation is essentially relegated to the bargaining phase making the auction phase somewhat superfluous.

The important aspect of these results is that we find that both buyers and sellers understand and respond to the strategic incentives of renegotiation and as a result, we find little evidence that renegotiation is used to correct for any problems stemming from the Winner's Curse. The primary

effect of renegotiation is to give those with the credibility to default an incentive to distort their bidding behavior during the auction so that they can win and then bargain for advantageous prices with the buyer. While renegotiation almost incidentally diminishes the number of defaults, it does so in a way that does not deal directly with the cause of the defaults and can cause other negative consequences through the resulting distortions in bidding behavior and the effect on the bargaining stage.

There are a number of differences in the experimental environment specified here and many external procurement cases. One key point is that the experimental design clearly specified whether renegotiation would or would not be allowed at the conclusion of an auction. In the field, many buyers may be less clear regarding whether renegotiation would be considered or not. While this is certainly a difference between our environment and the field, our results are still informative in regard to these situations. If renegotiation is uncertain, the theoretical implication is that a seller will form a belief about the possibility of renegotiation and will adjust their bid in a similar way based on that expectation. Our experimental results suggest that one would expect actual sellers to continue to respond as theory predicts and the main implication would be that buyers should discourage the belief that they would be willing to consider renegotiating.

It is of course important to note that there are practices commonly in use that may counter the adverse consequences of renegotiation. One common practice is to discard "low" bids from an auction. While used in practice, implementing this mechanism properly encounters multiple concerns. The first difficult is to define what "low" means. Some may automatically throw out the lowest bid or a certain number of low bids. Other rules involve discarding bids that are a certain distance below the average. In each case, the idea is to remove offers which appear too low and are more likely to lead to default. In certain cases, practices of this sort may be effective but that requires a careful determination of which bids to discard. If the cutoff is too high, a buyer could end up excluding valid bids and paying more than necessary. If the cutoff is too low, there may still be bids allowed which are low enough that a seller who wins can establish a credible commitment to default and gain leverage for post auction renegotiation. Further, this practice diminishes the pressure on sellers to bid low prices and so this may lead to shifts in bidding behavior adverse to the buyers' interest. The likelihood of success of this method depends on the procurer setting the threshold correctly. For a procurer who is well informed about the likely costs of the project, this may be possible, but an environment in which costs are known up-front is not one in which the Winner's Curse would be a concern.

There are of course other practices in use as well. One such practice is the use of a pre-screening mechanism prior to an auction to exclude financially constrained or undesirable bidders. A related mechanism is a multi-attribute mechanism in which the buyer factors in many aspects of a seller other than just the price offered in a pre-determined way to arrive at the winner. Some buyers also use what is referred to as a buyer-determined auction in which sellers submit bids but the winner determination rule is not made clear to bidders prior to the auction, allowing the buyer to select bidders based on his or her subjective preferences. Engelbrecht-Wiggans, Haruvy, and Katok (2007) discuss these methods in detail and describe how all of these methods allow a buyer to take into account aspects of a seller other than just price when determining who to contract with. The results from our paper provide strong support to the idea that buyers should look to practices like these rather than relying on price-only mechanisms along with the possibility of renegotiating after an auction has concluded.

While our analysis doesn't explicitly deal with long run dynamics of a market, our results do provide some insights regarding what sorts of things might occur in the long run should buyers be known to allow for renegotiation. Given that only sellers with some degree of bargaining power will be able to take advantage of the renegotiation prospect, one would expect that strong sellers would try to gain bargaining power. Exactly how they would do this would depend crucially on how buyers engage in renegotiation and in the nature of any pre-screening process buyers choose to use. There is little doubt, though, that sellers would try to put themselves in a position to take advantage of any renegotiation opportunities and there is also little doubt that such behavior would not be beneficial to the buyer. There could also be some interaction between these attempts to gain bargaining power and how strong firms might exert anti-competitive pressure on weaker firms to drive them from the market.

There are also possible positive consequences in the long-term as well. With repeated interactions, a norm may develop with sellers asking for renegotiated prices only when necessary to prevent default and buyers accepting when the request appears credible. Cooperation could be sustained through some form of trigger strategy. While such mutual cooperation is feasible, if a buyer has a seller with whom they have such a cooperative relationship, that buyer is unlikely to be using auctions as a procurement mechanism. Instead the buyer could and likely should just negotiate directly with the cooperative contractor which would be consistent with the argument that auctions can stifle the communication and cooperation necessary to complete a project (Gil and Oudot (2009); Bajari, McMillan, and Tadelis (2009)). Alternatively, long-run interactions could also lead to the

formation of reputations in which sellers that are notorious for bidding low just to gain an option for renegotiation could be excluded from future auctions through the pre-screening mechanisms we mentioned above. While this might decrease the negative consequences of renegotiation, there is still no clear argument that renegotiation can help once you have excluded the sellers most likely to take advantage of it.

While it isn't clear exactly what would occur in the long run, the short run implications of our results are clear. Both sides of the procurement relationship will act according to their immediate self-interest which leads to the renegotiation primarily serving to advantage sellers with a credible commitment to default. After winning an auction, these firms will use their strategic leverage to obtain very favorable prices upon renegotiation. While this may lessen the likelihood of default, it is not occurring in such a way as to compensate for the Winner's Curse. Consequently, in environments that are a reasonable match with the one studied here, low-bid procurement auctions with the possibility of renegotiation should be avoided in favor of other contracting mechanisms such as contingent contracting, performance bonds, or average bid auctions.

References

Bajari, P. and A. Hortacsu (2003). Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions. *Rand Journal of Economics 3*(2), 329–355.

Bajari, P., S. Houghton, and S. Tadelis (2014). Bidding for incomplete contracts: An empirical analysis of adaption costs. *The American Economic Review 104*(4), 1288–1319.

Bajari, P., R. McMillan, and S. Tadelis (2009). Auctions versus negotiations in procurement: an empirical analysis. *Journal of Law, Economics, and Organization 25(2)*, 372–399.

Bolton, G. and A. Ockenfels (2000). ERC: A Theory of Equity, Reciprocity and Competition. *American Economic Review 90*(1), 166–193.

Burguet, R., J.-J. Ganuza, and E. Hauk (2012). Limited liability and mechanism design in procurement. *Games and Economic Behavior 76*(1), 15–15.

Calveras, A., J.-J. Ganuza, and E. Hauk (2004). Wild bids. gambling for resurrection in procurement contracts. *Journal of Regulatory Economics 26*(1), 41–68.

Chang, W.-S., B. Chen, and T. C. Salmon (2015). An investigation of the average bid mechanism for procurement auctions. *Management Science 61*(6), 1237–1254.

Charness, G. and M. Dufwenberg (2006). Promises and partnership. *Econometrica 74*(6), 1579–1601.

Charness, G. and D. Levin (2009). The origin of the winner's curse: a laboratory study. *American Economic Journal: Microeconomics 1(1)*, 207–236.

Chen, J., L. Xu, and A. Whinston (2010). Managing project failure risk through contingent contracts in procurement auctions. *Decision Analysis 7*(1), 23–39.

Decarolis, F. (2013). Comparing public procurement auctions. Working Paper Boston University.

Decarolis, F. (2014). Awarding price, contract performance and bids screening: Evidence from procurement auctions. *American Economic Journal: Applied Economics 6*(1), 108–132.

Engelbrecht-Wiggans, R., E. Haruvy, and E. Katok (2007). A comparison of buyer-determined and price based multiattribute mechanisms. *Marketing Science 26*(5), 629–641.

Fehr, E. and K. M. Schmidt (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics 114*(3), 817–868.

Fischbacher, U. and F. Föllmi-Heusi (2013). Lies in disguiseŮan experimental study on cheating. *Journal of the European Economic Association 11*(3), 525–547.

Gil, R. and J.-M. Oudot (2009). Competitive bidding, renegotiation and relational contracting: Evidence from french defense. Working Paper.

Goeree, J. and T. Offerman (2002). Efficiency in auctions with private and common values: An experimental study. *American Economic Review 92*(3), 625–643.

Goeree, J. K. and T. Offerman (2003). Competitive Bidding in Auctions with Private and Common Values. *Economic Journal 113*(489), 598–613.

Güth, W., R. Schmittberger, and B. Schwarze (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organisation 3*(4), 367–388.

Harstad, R. M. and M. H. Rothkopf (1995). Withdrawable bids as winner's curse insurance. *Operations Research 43*(6), 983–994.

Kagel, J. H. and D. Levin (1986). The winner's curse and public information in common value auctions. *American Economic Review 76*(5), 894–920.

Kagel, J. H. and J.-F. Richard (2001). Super-experienced bidders in first-price common-value auctions: Rules of thumb, nash equilibrium bidding, and the winner's curse. *The Review of Economics and Statistics 83*(3), 408–419.

Kahneman, D., J. L. Knetsch, and R. H. Thaler (1986). Fairness and the assumption of economics. *Journal of Business 59*(4), 285–300.

Krasnokutskaya, E. and K. Seim (2011). Bid preference programs and participation in highway procurement auctions. *The American Economic Review 101(6)*, 2653–2686.

Levin, D. and P. Reiss (2013). Could we overcome the winner's curse by (behavioral) auction design? Working paper.

López-Pérez, R. and E. Spiegelman (2013). Why do people tell the truth? experimental evidence for pure lie aversion. *Experimental Economics 16*(3), 233–247.

Milgrom, P. R. and R. J. Weber (1982). A Theory of Auctions and Competitive Bidding. *Econometrica 50*(5), 1089–1122.

Parlane, S. (2003). Procurement contracts under limited liability. *The Economic and Social Review 34*(1), 1–21.

Rabin, M. (1993). Incorporating fairness into game theory and economics. *American Economic Review 83*(5), 1281–1302.

Roelofs, M. R. (2002). Common value auctions with default: An experimental approach. *Experimental Economics 5*(3), 233–252.

Shachat, J. and L. Tan (2014). An experimental investigation of auctions and bargaining in procurement. *Management Science 61*(5), 1036–1051.

Waehrer, K. (1995). A model of auction contracts with liquidated damages. *Journal of Economic Theory 67*(2), 531–555.

Wang, R. (2000). Bidding and renegotation in procurement auctions. *European Economic Review 44*(8), 1577–1597.