Received: 12 April 2023 | Revised: 30 August 2024 | Accepted: 2 September 2024

DOI: 10.1002/joom.1333

RESEARCH ARTICLE

WILEY

Helping hampered bidders—Do subsidy auctions work as intended?

Sanghoon Cho[1] | **Joel O. Wooten**[2] | **Timothy D. Fry**[2]

[1]Information Systems and Supply Chain Management Department, TCU, Fort Worth, Texas, USA

[2]University of South Carolina, Management Science, Columbia, South Carolina, USA

Correspondence
Joel O. Wooten, University of South Carolina.
Email: joel.wooten@moore.sc.edu

Handling Editor: Xiaosong (David) Peng

Abstract

The use of preference or discriminatory auctions, where one class of bidders is offered favored treatment over another class, has received mixed attention in the literature. Research has shown there is often an economic benefit of such policies in procurement auctions thanks to lower costs for buyers as incentives are offered to disadvantaged sellers. In this paper, we study one type of preferential procurement auction: the subsidy. Using a set of controlled experiments, we compare actual bidder behavior to what is predicted in equilibrium and find consistent (but overly aggressive) patterns overall. By testing a common bid strategy assumption, we also identify a behavioral framing bias that may trap sellers in these suboptimal strategies. Finally, we compare subsidies to another common discriminatory mechanism—the price preference—and find evidence that buyers interested in increasing the welfare of disadvantaged sellers should use subsidies instead of price preference auctions, thanks to a surprising difference in outcomes between preference types. Due to the wide use of bid preference auctions to support both policy and social aims, our findings have both financial and societal implications.

KEYWORDS

behavioral experiment, price preference, procurement auction, subsidy

1 | INTRODUCTION

There are many situations where a procurement auction buyer may want to favor one type of bidder over another. Non-price characteristics (e.g., distance from buyer and supplier reputation) heavily impact a buyer's surplus in many auctions (Haruvy & Katok, 2013; Tunca et al., 2014) as *strategic* considerations. *Policy* considerations also come into play, especially in the public sector, where the buyer may wish to favor a particular class of bidders (e.g., small, domestic, veteran-, or minority-owned businesses).

Government procurement contracts are one example where billions of dollars are awarded to favored bidders. In the United States, federal agencies awarded 28.4% of their $630B in contracts to small businesses in 2023 as part of a prescribed effort, including $65B to minority-owned businesses, $31B to women-owned businesses, and $32B to disabled veteran-owned businesses (U.S. Small Business Association, 2024). Alternatively, consider the policies of the Buy-American Act, where the US government awards supply contracts after adjusting bids from foreign suppliers: adding 6% when facing large domestic suppliers, 12% when facing small domestic

suppliers, and 50% when competing for Department of Defense contracts (Manuel, 2016). Other well-known policy examples include renewable energy auctions—where worldwide transfer payments equaled $140B in 2016 (Matthäus, 2020)—and resource development auctions, such as energy capacity projects in Europe (e.g., France and Germany subsidizing wind and solar tender contracts) and the 1993 Federal Communications Commission (FCC) auction for slices of regional radio bands (which gave women- and minority-owned businesses a 40% financing credit) (Ayres & Cramton, 1996; Mir-Artigues & Del Río, 2014). These examples all favor certain seller types, making uncompetitive bids from those groups more competitive.

In practice, three auction mechanisms have been used to help disadvantaged sellers: set-asides (or quotas), price-preferences, and subsidies. However, economic theory has not settled on which approach is best in which situation (Athey et al., 2013). Recent work has attempted to empirically analyze the effectiveness of certain auction elements—such as how to set subsidy rates in energy auctions (Matthäus, 2020)—and to what degree such mechanisms impact the buyer (Krasnokutskaya & Seim, 2011; Nakabayashi, 2013). Why are these issues (in the well-developed auction domain) gaining attention now? First, the use of procurement auctions has exploded over the last couple of decades thanks to the prevalence of online marketplaces. In a 2003 survey, only 15% of purchasing professionals reported the use of procurement auctions in their purchasing activities; by 2019, that number had risen to over 50%—with 95% of procurement agents age 30–44 having used reverse auctions (Forde, 2019). Second, the use of preference mechanisms has grown. Consider that an explicit priority of the US federal government is to award at least 23% of its annual contracts to small businesses, and there are now similar programs in place for other disadvantaged groups (Athey et al., 2013) as well as initiatives at the state and local government level. The sheer scale of such programs and the associated economic impact are demanding renewed attention. And finally, the combination of the two trends above is creating new situations for practitioners and new opportunities for researchers to assess the complex interplay between bidder behavior, market efficiency, and policy outcomes. Subsidies are used to achieve various economic and social objectives and understanding the effect of different subsidy structures on bidding behavior and market outcomes is critical to assessing their market efficiency and whether they are meeting their intended policy goals. Whereas set-asides and price preference auctions have been studied experimentally, the subsidy preference has received much less attention, despite the widespread use of auctions with subsidies. The significance of that gap has been highlighted as a research area needing careful examination in terms of the benefits, assumptions, and general usage conditions (Aloysius et al., 2016; Athey et al., 2013; He & Chen, 2021).

The objective of this paper is to address some of these gaps in the literature by investigating how differences in the structure of a subsidy auction might affect buyer and supplier outcomes. A major challenge with asymmetric auction models (where sellers are different in some way) is that there are seldom closed-form strategies (Flambard & Perrigne, 2006), especially given the richness and difficulty of the auction landscape. Recently, a number of papers have turned to behavioral experiments to help bridge the gap between economic theory and the decisions and outcomes being observed in real markets. Examples include explorations of rank feedback (Elmaghraby et al., 2012), scoring rule transparency (Quiroga et al., 2021), competitive information (Grosskopf et al., 2018), asymmetric information (Wooten et al., 2020), split-awards (Chaturvedi et al., 2019), bid visibility (Carter & Stevens, 2007), and seller cost and number (Aloysius et al., 2016) in procurement auctions. This paper began from a similar place, as our experience with procurement auction practitioners highlighted that certain bidding behaviors did not seem consistent with normative predictions.

In this paper, we experimentally test a set of equilibrium bidding strategies derived for procurement auctions (based on Rothkopf et al., 2003) to understand how bidders might behave in the presence of subsidies. We explicitly test several implications that follow from that excellent Rothkopf et al. (2003) paper and show how the choices buyers make when setting up subsidies impact sellers and, ultimately, profit. Specifically, we focus on three research questions: (1) whether sellers respond to varying subsidy levels as predicted, (2) if the framing of the bid changes how sellers formulate their bid, and (3) whether the choice of preference type influences seller bidding. The answers to these contribute to the field of operations management in two important ways. First, our research experimentally examines procurement auctions to help understand and explain fundamental bidder behavior in the presence of subsidies—in the same way several other papers in the *Journal of Operations Management* (e.g., Carter & Stevens, 2007; Quiroga et al., 2021) have used rigorous laboratory experiments to uncover important insights into human behavior in the auction landscape. Second, by empirically testing assumptions used by theoretical models in the auction literature, we uncover a common assumption that does *not* hold up in the lab. This has important theoretical contributions for researchers and important practical

contributions for practitioners. Overall, our behavioral experiments present novel evidence that increases our understanding of this growing auction type.

Our findings show that bidders largely respond to varying subsidy levels as predicted (although with more aggressive bids), that their bidding strategies may be more complicated than what theory assumes and subject to some behavioral biases, and that subsidies may be preferred to price preference auctions in order to achieve the original goal of the preference. Importantly, we uncover that the foundation of many normative bidding models—a percentage markup bidding strategy—does not prevail as the dominant approach of most bidders in our experiments. Given the importance of procurement auctions in global sourcing, our results identify a number of important ways that auction design could be modified when a particular class of bidder is preferred in order to avoid suboptimal outcomes.

2 | BACKGROUND ON AUCTION PREFERENCES

In addition to potential political and social benefits resulting from favoring a disadvantaged class of bidders, research has shown that such preference mechanisms are often theoretically optimal (Myerson, 1981). Despite increasing inefficiency (which occurs when the auction is won by a seller with higher costs), buyer procurement costs can actually be *reduced* by offering incentives to weaker sellers (Ayres & Cramton, 1996; Corns & Schotter, 1999; Hellerstein et al., 2014; McAfee & McMillan, 1989; Myerson, 1981; Rothkopf et al., 2003). By offering preferences to disadvantaged sellers, those sellers become more akin to advantaged sellers, which increases competition and results in more aggressive bidding by the *advantaged* parties. To start, we introduce the three main preference mechanisms used in practice—set-asides (or quotas), price-preference auctions, and subsidies—and describe several insights around bidder dynamics for each mechanism.

First, quota systems set aside a number of items for weaker sellers to bid on, leading advantaged sellers to adopt more aggressive bidding strategies on the remaining items due to differentiation (Ayres & Cramton, 1996; Hellerstein et al., 2014). In a first-price-sealed-bid (FPSB) multi-item auction format with two sellers, Hellerstein et al. (2014) show that in the laboratory, savings from an imposed quota are greater than predicted by a numerical equilibrium approximation. Assumptions include no uncertainty by the sellers regarding their true costs, which are drawn from a known uniform distribution. The effective implementation of a quota system requires that: (1) the buyer be able to identify a set of sellers that warrant favorable treatment, and (2) although the buyer need not know precisely the differences in costs for all sellers, the buyer must be able to reasonably estimate the cost difference between the two classes of sellers (Ayres & Cramton, 1996).

Second, price-preference incentives exist in auctions like those governed by the Buy-American Act, which favors domestic sellers by 6% (or 12% if a small business or 50% if the buyer is the Department of Defense) (Manuel, 2016). This is known as a price-preference auction since the winner of the auction may not be the seller who submits the lowest bid. The preference is used to compare bids and determine the auction winner, but that seller is only paid their submitted bid amount. In terms of effectiveness, there is evidence that price preferences are a more efficient mechanism than set-asides (McAfee & McMillan, 1989). In a study by Corns and Schotter (1999), equilibrium solutions estimated via numerical approximation (since closed-form solutions may not exist) are compared to seller bidding strategies from a set of laboratory experiments. Results indicate that if the preference is set appropriately, then aggressive bidding by low-cost sellers leads to cheaper procurement (that more than offsets any increases paid to high-cost sellers).

A third type of preference, which is the focus of this paper, is one where the incentive is offered in the form of a subsidy—with the winning seller receiving an additional payment if they belong to the favored group. If a disadvantaged seller submits the lowest bid, they are declared the winner and paid their bid amount plus a predetermined subsidy. If an advantaged seller wins the contract, they are paid only their bid amount. Such auctions have been used in a variety of settings: carbon emission reduction (He & Chen, 2021; Lopomo et al., 2011), telecommunications (Sorana, 2000), broadband service areas (Rosston & Wallsten, 2011), energy generation markets (Brown, 2018), and agricultural land rental rates (Kirwan, 2005). In terms of effectiveness, there is evidence that subsidies are a more efficient mechanism than set-asides if the goal is to ensure that the favored group of sellers wins sufficiently often (Pai & Vohra, 2012). Athey et al. (2013) suggest that even a small subsidy of 6% is sufficient to insure that disadvantaged sellers win at a level similar to a set-aside policy—and a subsidy up to 20% reduces buyer procurement costs relative to the set-aside policy. But it is less clear how subsidies compare to price preferences. Similar to the quota and the price-preference examples above, setting an appropriate subsidy requires some prior knowledge about the difference in cost between the high- and low-cost sellers.

This paper makes several specific contributions to this literature.[1] First, it establishes how sellers may actually behave in the presence of a procurement auction subsidy. Comparing behavioral results to theoretical results permits auction designers to know what to expect from sellers and what their actual procurement costs might be (which is especially important since some subsidies are politically or socially controversial). Since it is well known that bidders generally fail to bid rationally, a subsidy may not be a viable option to those buyers wishing to favor one type of bidder over another if the original intent of the preference auction is not realized. This directly responds to the call from Jap (2002) to understand different auction designs because of the big impact on buyer–seller relationships. Second, determining the appropriate level of incentive to set remains an open research question. Previous research agrees that calculating the optimal incentive level requires some understanding from the buyer of the degree of disadvantage between the advantaged and disadvantaged sellers (Ayres & Cramton, 1996; Corns & Schotter, 1999; Matthäus, 2020). If the preference is set too high, disadvantaged sellers may win more but with increased procurement cost for the buyer. Further, advantaged sellers may choose to not participate in the auction (thereby reducing the supplier base) if the subsidy is felt to be too generous to the disadvantaged sellers (Brown, 2018). If the preference is set too low, the goals of the preference may not be met. We examine a range of subsidy levels in our realistic setup to help with this question. Third, we show that a common assumption around bid strategy and the form it takes in many analytical auction models does not hold up in the lab. This applies to both how sellers develop their bid and also the applicability of bidding strategies to projects of varying sizes. As a final contribution, we recommend an increased use of subsidy auctions thanks to a difference in their theoretical and practical effectiveness. We show how sellers change their behavior based on the preference mechanism, and we offer prescriptive guidance to help buyers and policy setters more effectively use subsidies in procurement auctions.

3 | RESEARCH QUESTIONS

In this section, we highlight three questions related to the equilibrium bidding strategies predicted in the seminal paper by Rothkopf et al. (2003). In that paper, an analytical closed-form solution is presented in which disadvantaged sellers are offered a subsidy in the form of additional payment. For this type of asymmetric procurement auction, Rothkopf (1969) develops a foundational model in which low- and high-cost sellers utilize a multiplicative strategy such that their bids are simple scalar multiples of their estimated cost. The true cost for seller i, denoted by c_i, is unknown, and seller i prepares its estimate, $c_i' = \rho c_i$, where ρ is a random variable that follows a single, two-parameter Weibull distribution such that all sellers face equal uncertainty regarding their ability to estimate their own costs. Most studies on bid preference auctions assume the only seller uncertainty relates to a *competitor's* value of the good or service (and treat the seller's own cost to provide as known). Given that sellers are rarely certain about their actual costs in advance (Elmaghraby, 2007; Milgrom, 1981; Wooten et al., 2022), this cost uncertainty formulation closely approximates the environment that sellers and buyers face in real life.

If c_L and c_H denote the costs for the low- and high-cost sellers, respectively, where $c_L < c_H$, then the ratio of $\frac{c_H}{c_L}$ represents the relative degree of cost disadvantage between the two sellers, d. The equilibrium bidding strategies for the two sellers, P_L and P_H, are:

$$P_L = \frac{mX(d)[1+X(d)]^{\frac{1}{m}}}{[mX(d)-1]},$$

and

$$P_H = \frac{m[1+1/X(d)]^{\frac{1}{m}}}{[m-X(d)]},$$

where m represents a shape parameter in the two-parameter Weibull distribution, influencing the slope of the function, and

$$X(d) = \frac{\left\{m(1-d)+\left[m^2(1-d)^2+4d\right]^{\frac{1}{2}}\right\}}{2}.$$

Rothkopf et al. (2003) build on the multiplicative model presented by Rothkopf (1969) to account for a preferential auction where the buyer includes an incentive for a high-cost seller by paying a subsidy, s, in addition to her bid amount. In this situation, rational sellers should bid as if the cost is actually $c_h = \frac{c_H}{(1+s)}$ and the effective disadvantage ratio is $d_h = \frac{c_H}{[c_L(1+s)]}$. Since the probability that the disadvantaged seller wins is $\frac{X(d_h)}{[1+X(d_h)]}$, the expected amount paid is then calculated as $\frac{(c_H-c_L)X(d_h)}{[1+X(d_h)]}$, which leads to an expected procurement cost of $Z = \frac{md_h}{[m-X(d_h)]} + \frac{(c_H-c_h)X(d_h)}{[1+X(d_h)]}$. Based on the analytical models, optimal bidding strategies for each type of seller, corresponding expected profits, and the buyer's expected procurement cost are derived for several levels of cost disadvantage and cost estimation uncertainty.

Using a set of controlled experiments, we investigate whether the predictions from Rothkopf et al. (2003) can be observed in a laboratory setting. In that model, each seller has their own private value; however, because the authors assume there is uncertainty in that value for each seller, we also anticipate seeing some of the common auction literature behavioral deviations. While somewhat interesting (such differences will be noted when observed), we are more concerned with the unique results related to subsidies and propose three research questions around unique implications from the theoretical results.

3.1 | Convergence effect

Our first research question addresses one of the most basic results from Rothkopf et al. (2003)—how sellers respond to increased subsidy levels. When sellers are noticeably different, as in the bid preference auctions studied in this paper, the auction is said to be asymmetric. It is well known that asymmetric auctions often lead to inefficient outcomes in the sense that it is possible for a seller with a higher cost to win the auction over a seller who has a lower cost. The purpose of a subsidy is to enhance the ability of a favored class of seller who are generally thought to have higher costs to compete with low-cost sellers, thereby increasing the level of competition in the auctions. In the absence of subsidies, a disadvantaged seller must resort to very aggressive bidding in order to win at all (and will realize low profits), while the advantaged seller exploits their cost advantage and realizes high expected profits. As predicted by Rothkopf et al. (2003), a subsidy increases the level of competition between the two types of sellers, which results in the natural reaction of the advantaged (low-cost) seller submitting more aggressive bids. While the cost-disadvantaged seller will still bid more aggressively than the cost-advantaged seller (and realize decreased expected profits as the relative cost disadvantage increases), their level of aggressive bidding is less than when no subsidy is offered. Other forms of price preference yield similar results in laboratory settings. Hellerstein et al. (2014) investigate the effect of a quota system where several items are set-aside for disadvantaged sellers; in their lab experiments, advantaged sellers adopt more aggressive bidding strategies as available demand decreases. Corns and Schotter (1999) study a price-preference system where bids submitted by disadvantaged sellers are discounted by the buyer. Advantaged sellers again demonstrate more aggressive bidding in the lab as price-preferences are introduced. It is noteworthy that disadvantaged sellers *decrease* the aggressiveness of their bids with price-preferences—since the buyer adjusts bids post-submission, the price-preference allows disadvantaged sellers to increase their bids and remain competitive (Corns & Schotter, 1999). In the case of subsidies, disadvantaged sellers must account for the subsidy in their bid. There may certainly be differences in how sellers respond to the varying mechanisms, but as a first-order effect, we expect bidders to recognize and respond to the changing level of competition as subsidies are introduced. Moving from no subsidy to full subsidy should equalize any cost disadvantage. Thus, we expect bidding strategies to converge as subsidies increase. Our first research question examines this important implication from the theoretical models.

> CONVERGENCE EFFECT: Does increasing the level of subsidy cause advantaged sellers to bid more aggressively and disadvantaged sellers to bid less aggressively, narrowing the difference in bidding strategies, as predicted?

3.2 | Framing effect

Understanding bidding strategies in procurement auctions is important since sellers often face uncertainty as to their true costs, are subject to the winner's curse, and can lose money with a poor bid. In describing procurement auctions, analytical models often represent bidding strategies using a multiplicative strategy (i.e., real scalar multiples) in order to improve model tractability.[2] Framing bidder decisions as a percent markup makes sense if sellers are focused exclusively on profit margin and are not employing more complicated bidding strategies. This also has the benefit of allowing bid strategies to be easily described, readily compared across scenarios, and indifferent to project size. A seller employing such as strategy will use the selected markup whether the cost estimate is $10,000 or $100,000. This approach is defended as "often reasonably realistic from a behavioral point of view" (Rothkopf et al., 2003). However, the question of whether sellers do, in fact, utilize a multiplicative strategy for bidding—or if, instead, they use other policies—remains open and has implications for auction policy. As an example, Wooten et al. (2020) demonstrate that previous procurement auction outcomes myopically influence subsequent bids, so there is some evidence that sellers are not employing straightforward markups. Recognizing that heterogeneity likely exists in bidding strategies, our goal is not to perfectly describe all strategies, but rather to understand if a multiplicative strategy is realistic and what may hinder such strategies (especially in the case of subsidies, which complicates the decisions).

We devise a laboratory experiment to gain insight along this dimension and into the employed strategies. If the analytical models are correct, then we should observe similar markups regardless of whether sellers enter bids as a percentage markup (as suggested by the multiplicative strategy) or as a dollar amount. This follows from the principle of "descriptive invariance"—which posits that simply describing a decision with different but equivalent statements should not change choices (Tversky & Kahneman, 1992). We should also observe similar markups regardless of the size of the contract, as percentage markup strategies are indifferent to magnitude.

There is, of course, a rich body of work on biases in behavioral economics that shows—in violation of descriptive invariance—that humans behave differently based on how information is presented in an astounding array of contexts. Classic framing biases highlight information in either positive or negative terms and distinguish between valence effects in risky choice framing, attribute framing, or goal framing (Levin et al., 1998). In operations management, framing has been examined in broader situations. Consider two recent papers in the *Journal of Operations Management*—in Soundararajan and Brammer (2018), sub-suppliers were found to perceive sustainability requirements differently; in Kwasnitschka et al. (2024), feedback framing and psychological distance were used in conjunction with performance feedback delivered via smartwatches to explain behaviors in a manufacturing shop. While we are not aware of any studies exploring framing in the context of choosing bid strategies, there are some relevant works in accounting and marketing focused on eliciting choices with respect to number format. Hite and Roberts (1991) find that people's responses to ideal tax levels depend on whether the input is asked for as a percentage or dollar amount—the mean responses when entered as dollars are almost uniformly lower. In marketing contexts, charitable donation amounts framed in dollars (as opposed to percentages) are more effective for low-priced products, while the opposite is true for high-priced products (Chang, 2008). Thus, we see examples of varying decision-making strategies and do not expect descriptive invariance to hold in procurement bidding. Furthermore, psychology studies show that when presented with absolute numbers, people tend to process information more intuitively, whereas when presented with percentages, people tend to process information more analytically (Denes-Raj & Epstein, 1994). This leads people to behave more rationally and less spontaneously when dealing with percentages.

Our second research question examines whether sellers actually adopt a multiplicative bidding strategy, as employed by theoretical models. We tease this apart by exploring the framing of the bid entry and contract amount.

> FRAMING EFFECT: Does framing the bid input as a percentage markup or as a dollar amount cause sellers to bid differently in subsidized auctions—and does the size of the contract affect seller bidding strategies?

3.3 | Subsidy effect

Finally, we examine the choice of subsidy as the preference mechanism. Since the purpose of a preference auction is to favor one class of sellers over another, it is important that the procurement agent understand how sellers in an auction using different preference mechanisms behave. Without this understanding, it is possible that the intent of the auction (greater participation of the disadvantaged sellers) will not be realized. Not only has prior research not looked at how real auction behavior deviates from predicted subsidy models, but subsidy behavior relative to other preference mechanisms has been woefully understudied.

We start by comparing subsidies to another mechanism for which identical scenarios can be constructed, the price preference auction. The difference in these two mechanisms lies in how bidder parity is achieved. In price preference auctions, the buyer takes the advantaged seller's bid and adjusts it upwards before comparing it to the bid from the disadvantaged seller. In subsidy auctions, the disadvantaged seller knows that they will get an extra payment if they win, but they must factor that information into their bid. Thus, the buyer takes care of the preference mechanism for the advantaged seller, but the disadvantaged seller must perform that same task for themselves. The economics literature highlights that higher cognitive load reduces numeracy, increases the likelihood of anchoring, and leads to more risk-averse behavior (Benjamin et al., 2013; Deck & Jahedi, 2015). Each of these would lead disadvantaged sellers to not properly discount their bid enough in the case of subsidies. Our third research question follows this line of reasoning and explores how markups will differ based on the auction mechanism—and the associated implications on procurement costs for the *buyer*.

> SUBSIDY EFFECT: Will sellers in procurement auctions bid differently based on the type of price preference mechanism employed?

4 | EXPERIMENTAL DESIGN

We constructed a series of experiments to explore our three research questions around subsidies in asymmetric procurement auctions—organized as three studies. Lab experiments are the ideal method to explore these questions, as we can exactly replicate the parameters and setup in theoretical benchmarks, exercise a high degree of control—which reduces threats to internal validity via clean and powerful tests, and establish some first-order understanding about how individuals will behave. We followed best design practices (including published guidelines from the *Journal of Operations Management*) related to using student samples, validating with other subject pools, constructing incentives utilizing Induced Value Theory, and random assignment (Eckerd et al., 2021; Katok, 2018; Lonati et al., 2018).

All experiments were conducted in the z-Tree software platform (Fischbacher, 2007) assuming a first-price sealed-bid procurement auction format. We recruited 180 undergraduate students from an Operations Management core class at a major state university. Students were used as the subject pool because of the feasibility of controlling for unwanted confounds and offering meaningful incentives for performance. Additionally, there is evidence from the operations literature that students behave similarly to professionals (Lonati et al., 2018) and perform like experienced managers in profit-maximizing contexts (Bolton et al., 2012), which is the setting for our experiments. Despite this, students are still a sample of convenience, so we took care to increase the external validity where possible and pay attention to how we generalized results. For robustness along this dimension, we also recruited 30 experienced professionals as a second sample (details in Section 8).

For participating, each subject was paid $10 and exempted from one class assignment. In addition, participants were rewarded for how well they did in the auctions. One challenge with our auction format is that participants are subject to the winner's curse and can lose money with a bid that is too aggressive. It is not practical to take real money away from subjects, but some mechanism is needed to provide appropriate incentives.[3] Here, we provide an initial endowment and a rule for profit payouts. Each subject starts the experiment with a simulated balance of $500,000. At the end, the computer takes their balance and converts it to US dollars at the rate of 2 cents per $1000 for any amount over $400,000 ($2 for every $100,000). Thus, each seller has some cushion to prevent losses from spiraling out of control but also an incentive to perform well. This was explained to all subjects prior to the experiment and was consistent between sessions. On average, each participant received $16.04 in total (fee plus payout) for participating.

Participants assumed the role of contractors attempting to win construction contracts for 30 highway projects from the Department of Transportation (DOT). Upon arrival, each student was randomly assigned to one of the two seller types—cost-advantaged or cost-disadvantaged—and given a set of instructions according to their type. An administrator read aloud summary instructions and answered questions. Participants from the cost-advantaged group were randomly matched with participants from the cost-disadvantaged group each round, and each pair would then compete for a contract. The contracts were awarded via first-price sealed-bid auctions with private values—the first five auctions were trial runs to become familiar with the auction process and the remaining 25 affected their profit.

The underlying cost difference between sellers is calculated as follows. Each round, the advantaged seller's cost to complete the project (c_1) is randomly drawn from a uniform distribution between $100,000 and $400,000. This range comes from actual bid data from a set of 3960 state DOT projects from 2003 to 2015 and represents nearly 40% of the projects in that set. Then, the cost for the disadvantaged seller to complete the same project (c_2) is calculated by multiplying the advantaged seller's cost by the disadvantage ratio d for that treatment, $c_2 = c_1 \times d$.

In practice, a seller seldom knows their true project cost beforehand and thus must rely on an estimate of cost in preparing their bid. In the experiment, such an estimate is provided—derived in the same manner as Rothkopf et al. (2003)—by multiplying a seller's actual (but unknown) cost by an uncertainty adjustment drawn from a two-parameter Weibull distribution. All sellers face the same level of uncertainty in their cost estimates, which ends up being ∼12% (∼10%) relative standard deviation around the cost estimate for advantaged (disadvantaged) sellers.

The auction proceeds as follows. Each seller is provided with a cost estimate, unique to themselves, for their randomly drawn project. Sellers know the relative degree of cost disadvantage, the uncertainty for their estimate, and that their competitor faces the same uncertainty. Using their individual cost estimates, each seller prepares and submits a bid, and the lowest bidder is declared the winner. After each auction, all bids are posted, and each seller is informed as to the auction outcome. Given that project cost is unknown, the profit for a winning seller can be positive or negative. This is the well-known adverse selection problem known as the

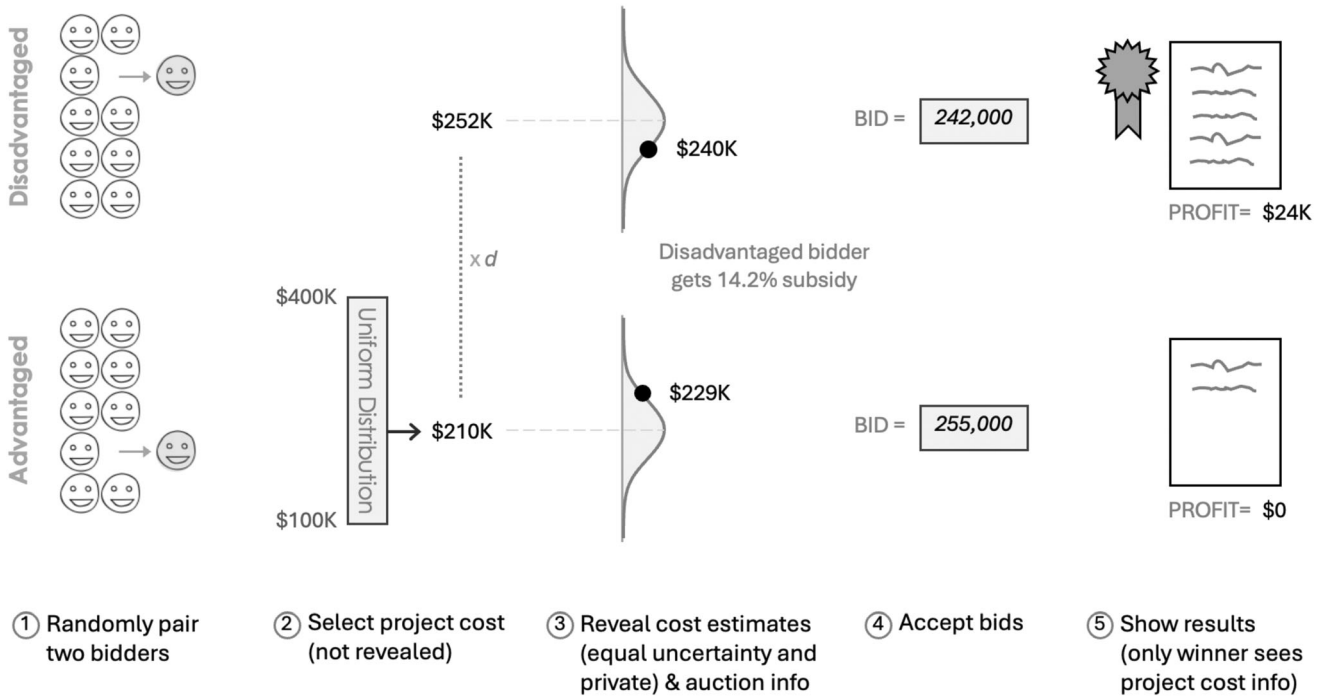

FIGURE 1 Example of each round of bidding.

winner's curse. If a seller wins the auction, the actual cost is revealed, profit is calculated (including the subsidy, if appropriate), and their cash balance is updated. For the losing seller, only their unchanged cash balance is revealed. An illustration of each round of bidding is shown in Figure 1. Example participant instructions and screenshots from experiments are given in Appendix A and B, respectively.

After completing all lab sessions, we aggregated the data. The first five (trial) auctions were excluded from the analysis. We also omit the last five rounds of data because two of our sessions ended early due to student scheduling conflicts.[4] For all analyses that follow, we use a consistent set of 20 rounds. We inspected the data and concluded that all subjects demonstrated an excellent understanding of the instructions and followed a reasonable bidding strategy. However, there are a handful of unusual bids that we flagged as mistakes. For example, one participant entered a bid with a markup of 284% when their other nineteen markups averaged 24%. It seems likely that an entry error occurred (one extra digit). We handled these few instances by locating within-subject outliers more than three standard deviations away from each subject's average bidding strategy. In these 21 instances (out of 3600 bids), the results from that round (including the results of their auction partners in that round) were removed from our analysis. Next are the details for our three studies.

5 | STUDY 1 (CONVERGENCE EFFECT)

5.1 | Treatment

For our first study, we look at whether increasing levels of subsidy support lead to converging bids from sellers, as theory predicts. We specify three subsidy levels (none, optimal, and full—which is 100% of the cost disadvantage) for two different cost disadvantage levels ($d = 1.20$ and $d = 1.05$). These three subsidy levels and two levels of cost disadvantage mimic those presented as analytical results in Rothkopf et al. (2003), allowing us to compare our experimental results to those predicted by the Rothkopf model. These scenarios are shown in Table 1 along with the number of participants for each treatment.

5.2 | Results and discussion

To begin, we report our empirical results next to the theoretical predictions from Rothkopf et al. (2003). Table 2 shows the normative bidding strategies and seller profits alongside those observed in Study 1 for both cost-advantaged and cost-disadvantaged sellers, along with the percent difference between the predicted and observed values. Columns 1–2 show the parameter levels, Columns 3a–3f present the seller bidding data, and

TABLE 1 Experimental design—Convergence effect.

Treatment	No. of auctions	No. of participants	Cost ratio (d)	Preference mechanism	Preference amt (%)	Bid input
1	177	18	1.2	Subsidy	**0**	$
2	214	22	1.2	Subsidy	**71**	$
3	140	14	1.2	Subsidy	**100**	$
4	138	14	1.05	Subsidy	**0**	$
5	140	14	1.05	Subsidy	**72**	$
6	139	14	1.05	Subsidy	**100**	$

Note: Preference amounts represent none (0%), optimal, and full (100%) subsidy levels.

Columns 4a–4f present seller profit. To account for the magnitude in project costs, we calculate percent profit as (bid amount—actual cost)/actual cost. Buyer procurement costs are shown in Columns 5a–5c, while the percent of auctions won by the cost-disadvantaged seller is shown in Columns 6a–6c. The values presented in Table 2 represent aggregated averages across individual treatments and specific bidder type groups.

5.2.1 | Seller bidding strategies

Looking at Columns 3a and 3d reveals the insights from the normative model for two bidders. Theory predicts that (1) cost-advantaged sellers should bid more aggressively (lower markups) as the subsidy increases, (2) disadvantaged sellers should bid less aggressively (higher markups) as the subsidy increases, and (3) bidding strategies should converge as the subsidy increases from 0% to 100% (and cost-differentiated bidders become identical).

Columns 3b and 3e show the bidding strategies observed in our experiment. Several things are immediately apparent. First, both seller types submit bids that are very aggressive compared with theoretical predictions in all treatments. In other words, both types of sellers exhibit overbidding, which has been shown to be a common occurrence in the experimental research. If we look at the % deviation for markups (Columns 3c and 3f), we see that advantaged sellers submit bids 22% and 36% lower than predicted on average (for cost disadvantages of 20% and 5%, respectively) and disadvantaged sellers submit bids 33% and 50% lower on average. These differences are statistically significant (t-test $p < .001$) for every comparison in Columns 3a/b and 3d/e.

For ease of discussion, Figure 2 graphically displays the bidding strategies data. Looking at the panels ($d = 1.2$ and 1.05), the normative model insights seem to hold, but with more aggressive markups across the board: as the subsidy increases, advantaged sellers decrease their markup and disadvantaged sellers increase their markup. In both cases, bidding strategies converge with a full subsidy, at which point both seller types adopt a statistically indistinguishable strategy ($p = .792$ and .691 for $d = 1.2$ and 1.05, respectively), bidding as if their costs were the same. This answers our convergence effect question and shows high-level bidding behavior trending in line with established theory, with increased subsidies bringing seller bids closer together.

One of the issues observed in the above preference auction experiments is seller bids that are lower overall than theory would predict (Corns & Schotter, 1999; Hellerstein et al., 2014). A large body of literature has demonstrated that bidder strategies often deviate from predicted levels in this manner—with laboratory bids that are more aggressive than normative predictions (e.g., Dyer et al., 1989). One explanation for this behavior that has emerged is learning—bidders get better as they experience an auction setting. This is true for common value English auctions, where a unique stable equilibrium is ultimately reached when the game is repeated many times, although some naïve players can still be found (Deltas & Engelbrecht-Wiggans, 2005). And it is also true for second-price sealed bid auctions, where inexperienced bidders are the main source of irrationality (Zeithammer & Adams, 2010). Srinivasan and Wang (2010) show that inexperienced bidders adjust their bidding strategy toward rational bidding over time. Even in independent private value auctions, repetition allows bidders to acquire useful information (Hon-Snir et al., 1998; Jeitschko, 1998), with bidders adopting the Nash equilibrium strategy in the long run even without feedback on the auction results (Neugebauer & Perote, 2008). We expect bids in subsidy auctions to more closely align with theoretical predictions through repetition, leading to pronounced learning effects. The results are given in Appendix D.

TABLE 2 Experimental results versus theoretical predictions.

1	2	3a	3b	3c	3d	3e	3f	4a	4b	4c	4d	4e	4f	5a	5b	5c	6a	6b	6c
		Bidding strategy (%)						**Profits (%)**											
		Advantaged seller			**Disadvantaged seller**			**Advantaged seller**			**Disadvantaged seller**			**Procurement cost (%)**			**Inefficient outcome (%)**		
d	Subsidy	Pred	Obs	Δ%	Pred	Obs	Δ%	Pred	Obs	Δ%	Pred	Obs	Δ%	Pred	Obs	Δ%	Pred	Obs	Δ
1.2	0	31.2	23.6	−24%	17.5	10.2	−42%	17.6	12.8	−28%	2.0	−0.8	−138%	26.1	19.2	−27%	32.6	36.5	3.9
	71	21.2	17.9	−16%	17.9	12.2	−32%	7.9	5.1	−36%	4.1	0.4	−90%	21.5	13.7	−36%	44.5	47.2	2.7
	100	19.1	14.3	−25%	19.1	14.0	−27%	5.6	4.5	−19%	5.6	3.7	−33%	22.2	19.0	−15%	50.0	50.0	0.0
Average		**23.8**	**18.6**	**−22%**	**18.2**	**12.1**	**−33%**	**10.4**	**7.5**	**−28%**	**3.9**	**1.1**	**−71%**	**23.3**	**17.3**	**−26%**	**42.4**	**44.6**	**2.2**
1.05	0	21.1	15.4	−27%	17.9	7.3	−59%	7.9	2.7	−66%	4.1	1.4	−66%	14.2	4.6	−67%	44.6	52.9	8.3
	72	19.6	13.2	−33%	18.7	9.9	−47%	6.1	0.6	−90%	5.1	1.3	−75%	13.8	6.7	−51%	48.5	50.0	1.5
	100	19.1	10.0	−48%	19.1	10.5	−45%	5.6	0.1	−99%	5.6	1.4	−75%	13.9	4.0	−71%	50.0	49.6	−0.4
Average		**19.9**	**12.8**	**−36%**	**18.6**	**9.2**	**−50%**	**6.5**	**1.1**	**−83%**	**4.9**	**1.4**	**−73%**	**14.0**	**5.1**	**−63%**	**47.7**	**50.8**	**3.1**

Note: For the observed disadvantaged bidding strategy, we use adjusted numbers to match those reported by Rothkopf et al. (2003), where markup is calculated from a cost estimate that is reduced proportionally to the subsidy level. With no subsidy, markup is just calculated relative to the cost estimate. Disadvantaged Bidding Strategy (Markup) $= \frac{\text{Bid Amount} - \left(\frac{\text{Cost Estimate}}{1+s}\right)}{\left(\frac{\text{Cost Estimate}}{1+s}\right)}$ Comparing predicted versus observed values for bidding, profit, and procurement with *t*-tests revealed significance levels <0.01 for every cell. Results shown for experimental treatments 1–6.

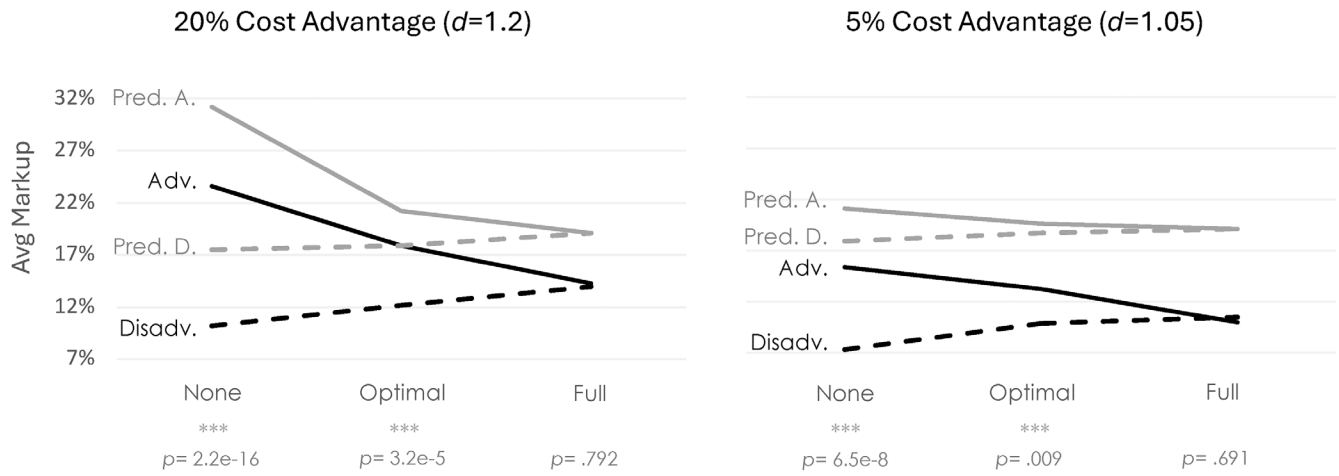

FIGURE 2 Bidding strategy—Observed versus predicted. Significance *t*-tests shown between observed types (Adv vs. Dis) for each subsidy level. ***$p < .001$, **$p < .01$, *$p < .05$.

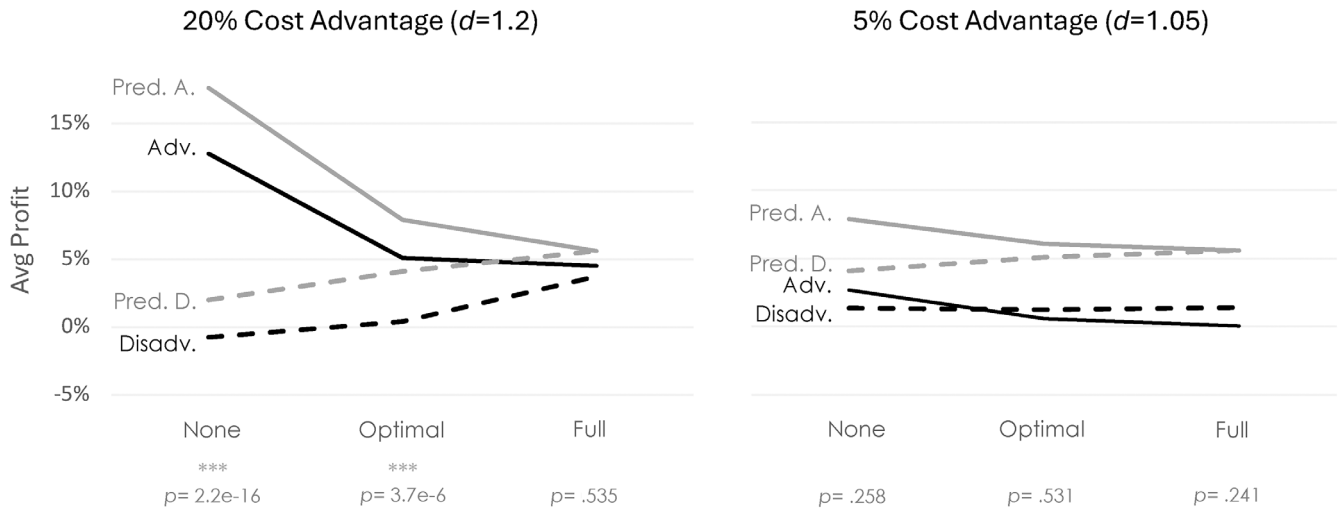

FIGURE 3 Profit—Observed versus predicted. Significance *t*-tests shown between observed types (Adv vs. Dis) for each subsidy level. ***$p < .001$, **$p < .01$, *$p < .05$.

5.2.2 | Seller profits

In terms of profit, the patterns from the normative model correspond to the patterns in the normative bidding strategy (Rothkopf et al., 2003), as one would expect. As the level of subsidy increases, profits for the advantaged seller decrease, and those for the disadvantaged seller increase—which is the original purpose of the preference.

Figure 3 presents a graphical illustration of the profits for the two seller types by subsidy and cost disadvantage. Our experimental results follow the predicted patterns quite well, suggesting that bidders are behaving as expected in many ways—such that the convergence effect demonstrated in bidding strategies leads to similar patterns in profit. When a full subsidy is used, the profits for both sellers are statistically indistinguishable for both cost disadvantages ($p = .535$ and .241 for $d = 1.2$ and 1.05, respectively), though always less than the predicted profits (due to the aforementioned overbidding). Rothkopf et al. (2003) suggest that a cost-advantaged seller should never bid aggressively enough to shut out the disadvantaged seller from winning (as it reduces long-run profits substantially) and instead should adopt a conservative bidding strategy. This tactic was not heeded by our subjects, as the disadvantaged seller often realized negative profits, which in a real setting could not have been sustained.

TABLE 3 Experimental design—Framing effect.

Treatment	No. of auctions	No. of participants	Cost ratio (*d*)	Preference mechanism	Preference amt (%)	Bid input
1	178	18	1.2	Subsidy	**0**	**$**
2	214	22	1.2	Subsidy	**71**	**$**
3	140	14	1.2	Subsidy	**100**	**$**
7	180	18	1.2	Subsidy	**0**	**%**
8	160	16	1.2	Subsidy	**71**	**%**
9	140	14	1.2	Subsidy	**100**	**%**

Note: Preference amounts represent none (0%), optimal, and full (100%) subsidy levels.

5.2.3 | Buyer procurement costs

Rothkopf et al. (2003) suggest that with no subsidy, the disadvantaged seller presents little competition for the advantaged seller. Consequently, the advantaged seller should rationally adopt a bidding strategy that allows the disadvantaged seller a reasonable probability of winning a bid—such a strategy maximizes the profits of the advantaged seller and results in an increased buyer procurement cost. Table 2 (Columns 5a–c) shows these data. Important to the buyer is the fact that a subsidy can actually decrease her procurement cost (as bidding becomes more competitive). In our lab scenarios with a 20% cost disadvantage, procurement costs average 26% less than predicted; with a 5% disadvantage, procurement costs for the buyer average 63% less than predicted.

Interestingly, the probability of an inefficient outcome—where the high-cost seller wins the auction (Columns 6a–c in Table 2)—is relatively close to predicted level. So, while the number of auctions won by the disadvantaged seller is close to that predicted, due to the aggressive bidding by both types of sellers, the final cost to the buyer is much less than predicted. Another notable observation is that buyer procurement costs in one particular case do not mirror theory. Theory predicts that an optimal subsidy will reduce buyer procurement cost a bit more than a full subsidy. Here, we see one case (5% cost difference) where buyer cost goes *up* and another (20% cost difference) where buyer cost is dramatically reduced. This suggests that the bidding dynamics are playing out differently than predicted. We explore some of these possible behavioral issues in the next experiments.

6 | STUDY 2 (FRAMING EFFECT)

6.1 | Treatment

In our second study, we turn to our research question on framing to help understand how bids are formulated. Theory suggests that sellers use a percentage-based markup strategy. Specifically, we test if bids are different depending on how bidders are prompted to enter their bids (either as dollar bids or percentage markups), which has implications in practice as well as for interpreting theory. We use our baseline treatments (1–3) from the previous study and include a set of treatments (7–9) that require subjects to enter their bids as percentage markups instead of dollar bids—but are otherwise identical (Table 3).

6.2 | RESULTS AND DISCUSSION

Figure 4 shows the average markups observed (all reported as a %), along with the 95% confidence intervals. We observe that sellers bid differently based on bid framing—and that it differs by bidder type. Disadvantaged sellers bid more aggressively using dollar-bids than percentage-bids, while advantaged sellers bid more aggressively using percentage-bids than dollar-bids. These differences are statistically significant except for the full subsidy treatment.

To investigate further, we look at whether the size of the contract affects the bidding strategy. We compare projects with lower costs to those with higher costs. Since project cost for our advantaged seller is taken from a uniform distribution between $100,000 and $400,000, we take the midpoint of that range as the breakpoint.[5] Figure 5 shows average markup by the size of the contract. We again observe differences, but with an interesting pattern. It turns out that all the differences observed in Figure 4 are driven by a particular type of project. For disadvantaged sellers, there is no difference in their bidding strategies for small contracts; only for large contracts do they adopt different markups based on the input framing of the bid ($p = .006$ and 2.9e−4). One might think that this is simply an amplification of a numeracy problem—that as the magnitude of the numbers gets larger, analytical mistakes are more obvious and lead to

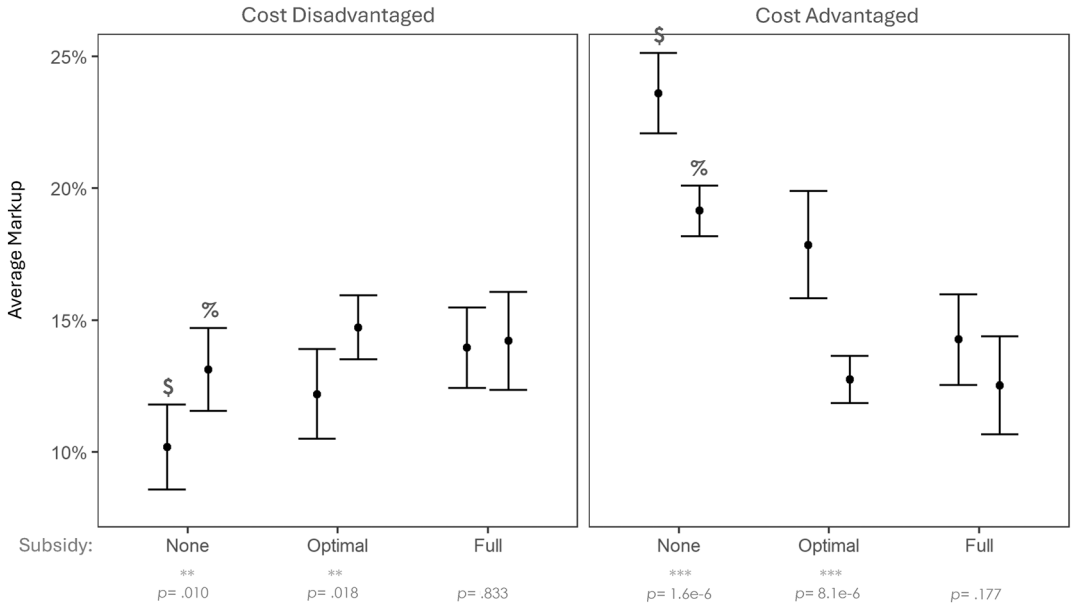

FIGURE 4 Bidding strategy based on bid input framing ($ vs. %). Significance *t*-tests shown for percent markup (%) versus dollar amount ($) bid framing; Error bars = 95% CI. ***$p < .001$, **$p < .01$, *$p < .05$.

this result. However, advantaged sellers have the *opposite* issue. For advantaged sellers, there is no difference in their bids for large contracts; only for small contracts do they adopt different markups based on the input framing of the bid ($p = 5.0e{-}9$ and $3.2e{-}6$). Thus, it is unlikely that we are simply seeing some sort of calculation problem. Because of the stark difference in how these differences manifest for different sellers, the better explanation is a behavioral one. A poker analogy may help illuminate this behavior. Just like chip stack and the size of the pot matter in poker, cost advantage and size of contract influence bid behavior. In procurement auctions, advantaged sellers may bid more for "small pots" when focused on the dollar amount because the small dollar amount does not feel big enough compared to their relative advantage. This fits with the psychology literature findings around more spontaneous/intuitive information processing when dealing with absolute numbers (Denes-Raj & Epstein, 1994). In a similar way, disadvantaged sellers may bid less for "big pots" when focused on the dollar amount because any profit is good from where they sit. In both cases, this is a form of anchoring, where bidders have a minimum dollar amount they expect from "winning a hand" given the difference in power—and that manifests as higher bids by advantaged sellers in small contracts and lower bids by disadvantaged sellers in large contracts. This pattern does not hold in the full subsidy case, but in that scenario, sellers have been placed on equal footing so there is no "chip stack" imbalance. This suggests that there is a framing effect based on the size of the contract. Importantly, we show that the assumption in analytical models that bidder strategies can be described in terms of simple percentage markups does work in some scenarios—but fails to account for bidder behavior in many others.

In addition to *when* sellers deviate, we also explore *how* they deviate. Specifically, we test if the sellers behave more rationally (less spontaneously) when dealing with percentages compared to dollar amounts. To do so, we use the coefficient of variation (CV) statistics to compare variations from two different input types. Specifically, we utilize the asymptotic test for the equality of CV (Feltz & Miller, 1996). We also reduce every seller's data to the standard deviation of their bids and show simple t-tests for the different framing conditions. Table 4 shows that bidders are more consistent with their markup strategy when they enter percentage bids, as might be expected. This is particularly stark in the advantaged seller case, where the variability in bid markup is 94% bigger (9.5% vs. 4.9%, $p = 8.0e{-}4$).

7 | STUDY 3 (SUBSIDY EFFECT)

7.1 | Treatment

For our final study, we examine whether the preference mechanism used in a procurement auction influences the outcomes. We compare subsidies to another widely used preference mechanism—the price preference. Ignoring the uninteresting case when no preference is given (none or 0%), we again devise a set of experiments (10–11) to

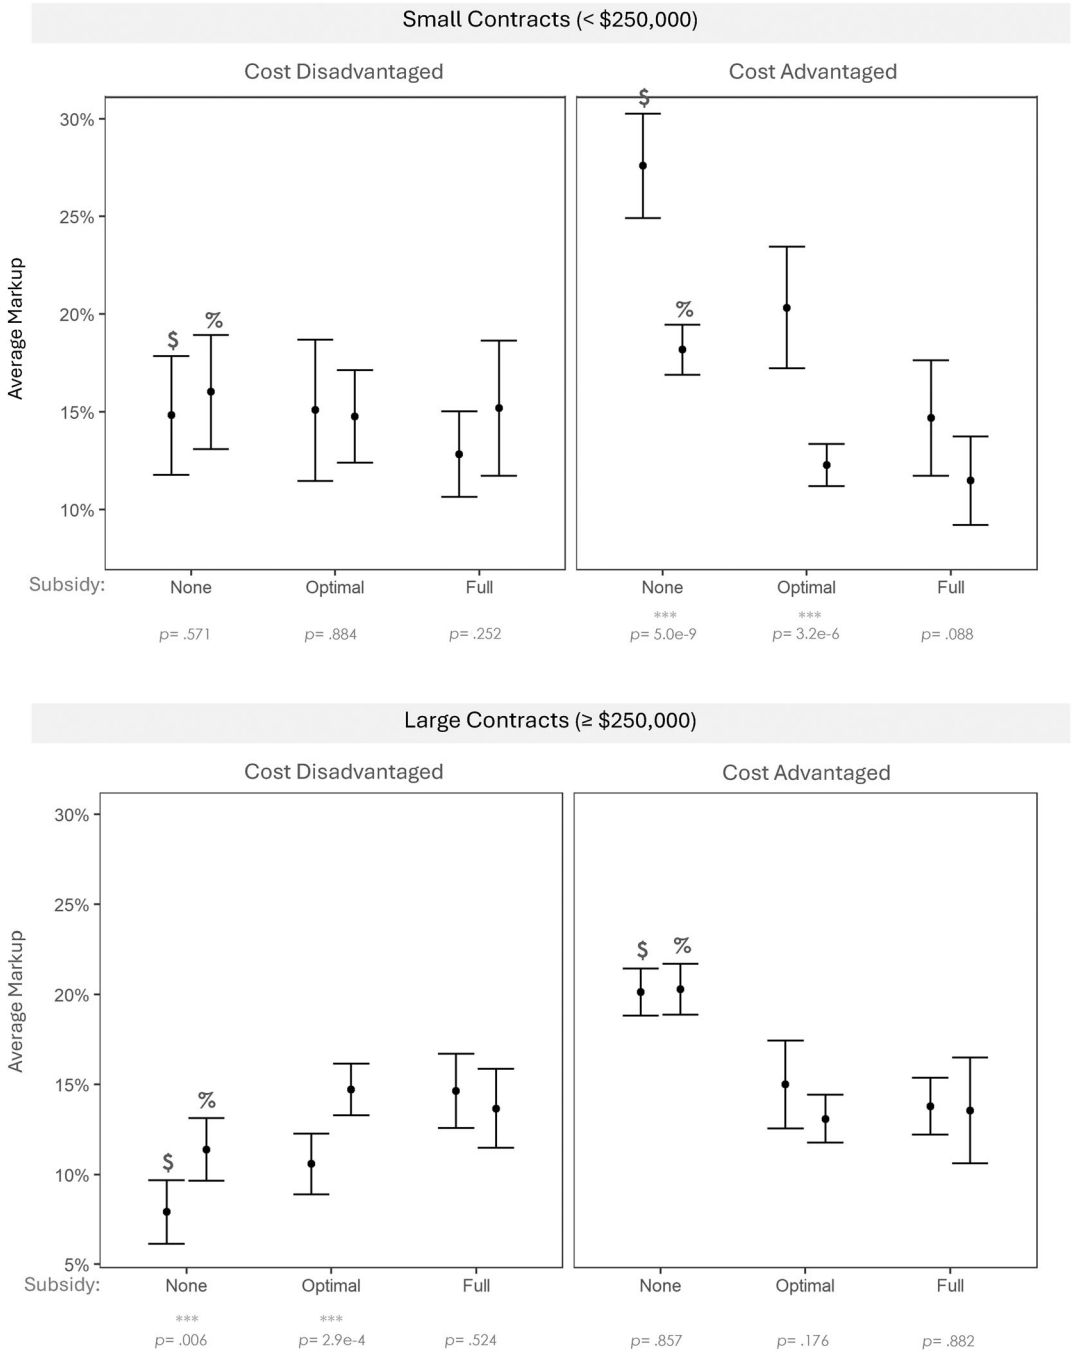

FIGURE 5 Bidding strategy and bid input framing ($ vs. %) by contract size. Significance *t*-tests shown for percent markup (%) versus dollar amount ($) auctions; Error bars = 95% CI. ***$p < .001$, **$p < .01$, *$p < .05$.

TABLE 4 Bidder consistency based on bid input framing ($ vs. %).

	SD of percent (%) bids	SD Of Dollar ($) bids	CV test	*t*-test
Disadvantaged	0.061	0.080	$p = 5.5e{-}5$	$p = .092$
Advantaged	0.049	0.095	$p = 8.0e{-}4$	$p = 2.4e{-}5$
Overall	0.055	0.087	$p = 4.1e{-}8$	$p = 7.0e{-}5$

Note: Significance *t*-tests are shown for comparison of markup standard deviation by bidder.

TABLE 5 Experimental design—Subsidy effect.

Treatment	No. of auctions	No. of participants	Cost ratio (*d*)	Preference mechanism	Preference amt (%)	Bid input
2	214	22	1.2	**Subsidy**	71	$
3	140	14	1.2	**Subsidy**	100	$
10	139	14	1.2	**Price Pref**	71	$
11	211	22	1.2	**Price Pref**	100	$

Note: Preference amounts represent optimal and full (100%) subsidy levels.

TABLE 6 Demonstration of equivalency—Subsidy versus price preference mechanism.

	Full subsidy s=20%, d=1.2			Full price preference s=20%, d=1.2	
	Adv.	Disadv.		Adv.	Disadv.
Cost estimate	100	120		100	120
Markup (19.1%)	19.1%	-0.7%	Seller_{DIS} bids as if Seller_{ADV}: $C=C_H/(1+s)$	19.1%	19.1%
Seller bid	119.1	119.1		119.1	142.9
Auction decision	**119.1**	**119.1**	Buyer adjusts Seller_{ADV} bid: (1+s)*Bid	**142.9**	**142.9**
Buyer cost	**119.1**	**142.9**	Buyer adjusts Seller_{DIS} payout: (1+s)*Bid	**119.1**	**142.9**

Note: Illustrative for $d = 1.2$, $s = 20\%$, with cost estimates 20% apart (as if centered on actual costs). Similar logic applies in the case of an optimal subsidy, which is detailed in Appendix C.

compare to our base case treatments (2–3) that are identical except in terms of the preference mechanism used (Table 5).

It is, of course, critical that these two mechanisms are comparable. The main difference in the two lies in how bidder parity is achieved. In subsidy auctions, the disadvantaged seller gets an extra payment if they win, information that must be factored into their bid decision. In price preference auctions, the buyer takes the advantaged seller's bid and adjusts it upwards before comparing it to the bid from the disadvantaged seller. Thus, the buyer takes care of the preference mechanism for the advantaged bidder in second case, but the disadvantaged bidder must perform that same task for themselves in the first. In order to compare the two preference mechanisms experimentally, however, we must choose scenarios for which the two are identical. Fortunately, the same properties that result in a fully proportional model in Rothkopf et al. (2003) help us in the price preference case. Table 6 shows that subsidy auctions are equivalent to price preference auctions in the full subsidy case.[6] The key comparisons are in the last two rows, which show the number that decides the winner and the realized cost for the buyer. In the case shown, the sellers look identical (with the subsidy and price preference cases—on the left and right—yielding equal comparisons within auctions and identical payouts across auctions). Regardless of the markup chosen by each seller type, the outcome (decision and buyer cost) of these two preference mechanisms will be equivalent.

7.2 | Results and discussion

Figure 6 shows the average markup (and 95% confidence interval) in our subsidy and price preference auctions in the presence of a preference. We immediately observe that the bidding strategy is significantly different between the preference types at the optimal preference level ($p = 3.5e{-}8$ and $2.2e{-}16$). This means that sellers in procurement auctions may bid differently based on the preference mechanism employed when the buyer sets a theoretically optimal preference level. Such behavior is particularly concerning because the observed markups in the subsidy treatments are already more aggressive than theory would predict (see Figure 2, Table 2)—so much so that many of those disadvantaged sellers lose money in an unsustainable way. These new trials show that when the burden of the preference computation is shifted to the advantaged bidders (in the price preference condition), they deviate wildly.

Finally, we check if procurement costs for the buyer will be affected by the choice of preference mechanism. We check for these with simple *t*-tests on the average responses. Figure 7 shows that the average buyer procurement costs are significantly higher for subsidies than

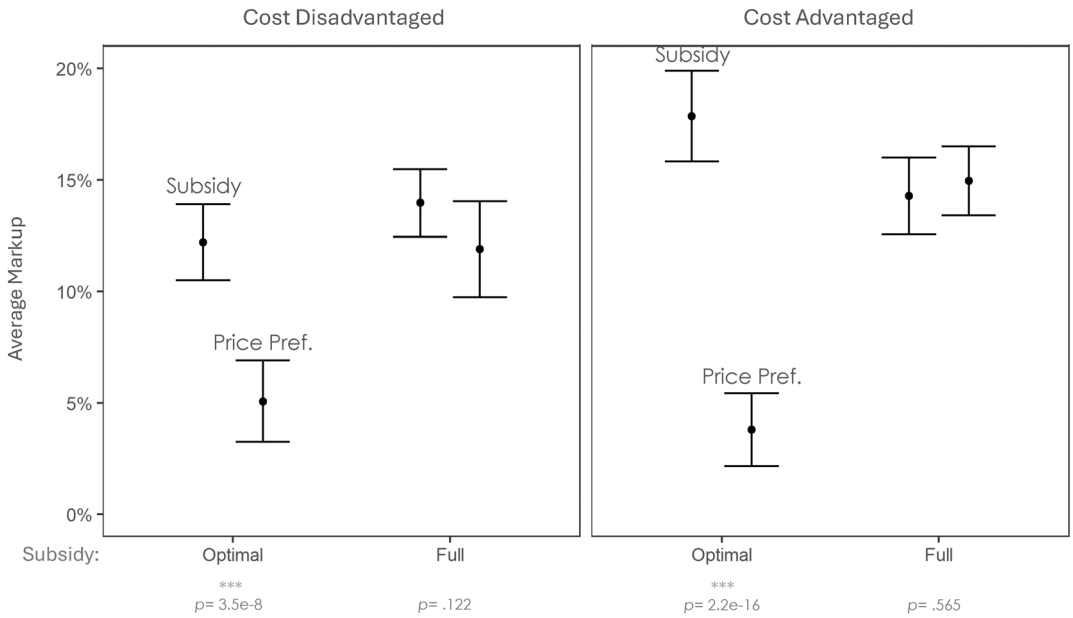

FIGURE 6 Bidding strategy and preference mechanism. Significance *t*-tests shown for subsidy versus price preference; Error bars = 95% CI. ***$p < .001$, **$p < .01$, *$p < .05$.

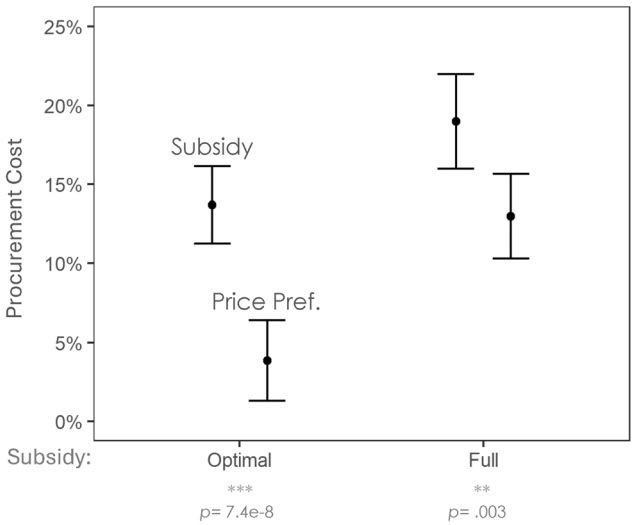

FIGURE 7 Procurement cost based on subsidy versus price preference. Significance *t*-tests shown for subsidy versus price preference auctions; Error bars = 95% CI. ***$p < .001$, **$p < .01$, *$p < .05$.

for price preferences at both preference levels ($p = 7.4e{-}8$ and .003 for optimal and full, respectively). This supports our conjecture that the preference mechanism will affect the bidding strategy of sellers, although it is worth mentioning that costs in both preference mechanisms are at or below predicted costs. Importantly, the reason we have preference auctions is to promote participation by disadvantaged sellers. We discuss in the final section the implication of these results on participation by (and success of) disadvantaged bidders.

8 | ROBUSTNESS

Prevailing wisdom suggests that using undergraduate students in laboratory experiments provides excellent data—with results that are equivalent (if not better) to using managers (Fréchette, 2012; Katok, 2018), along with ancillary benefits in access, control, and incentive compatibility (Katok, 2018). For example, Bolton et al. (2012) conclude that managers generally exhibit a *similar* pull-to-center bias as students and do *not* surpass students in applying information and task training. However, to make sure our experiments do not suffer from some domain-specific failing with respect to our sample population, we carry out a robustness test with experienced professionals in an identical laboratory setting. We recruited a set of working professionals from the executive education and professional MBA programs at the same major state university as before. These professionals were actively working full-time, had an average of 10.2 years of work experience (Q1 = 5, Q3 = 14), and represented a diverse set of fields (with the majority coming from operations, finance, and analyst roles). For replication with this smaller sample, we employed our baseline treatments (1–3) from before, as detailed in Table 7.

Due to the busy schedules of the professionals, our experiments were conducted with fewer rounds than the original treatments: 20 rounds for the first two treatments and 18 rounds for the final treatment. The first five rounds were trial rounds and are excluded from the results, as before. The observed outcomes, contrasting

TABLE 7 Experimental design—Robustness sample.

Treatment	No. of auctions	No. of participants	Cost ratio (d)	Preference mechanism	Preference amt (%)	Bid input
1*	90	12	1.2	Subsidy	0	$
2*	60	8	1.2	Subsidy	71	$
3*	65	10	1.2	Subsidy	100	$

Note: Preference amounts represent none (0%), optimal, and full (100%) subsidy levels.

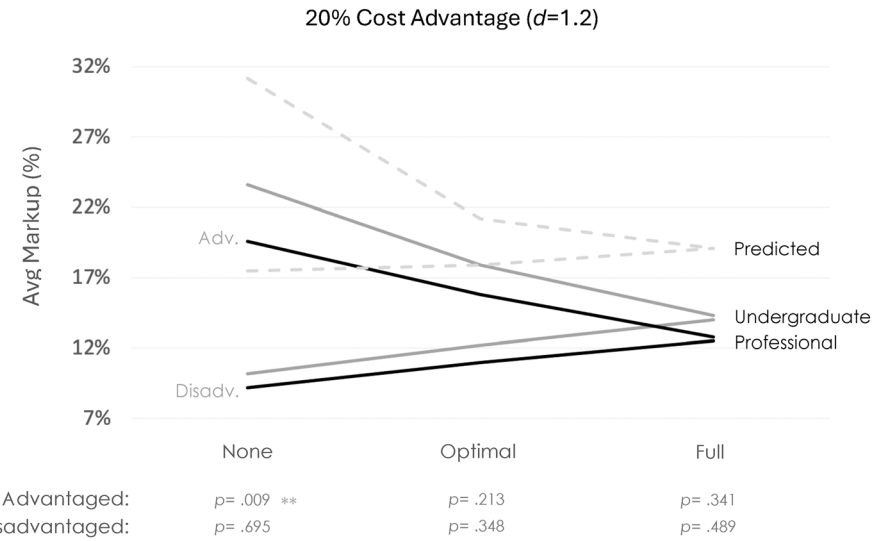

FIGURE 8 Bidding strategy—professionals versus undergraduates versus theory. Significance t-tests shown between cohorts (Undergraduate vs. Professional) for each subsidy. ***$p < .001$, **$p < .01$, *$p < .05$.

undergraduates and professionals (alongside the predicted values), are presented in Figure 8.

Even with a smaller number of observations and participants, our results are almost identical. The pattern of bidding behavior observed in the professionals is very consistent, with the same wedge shape as in our undergraduate bids (our Convergence Effect from Study 1). The only visual difference is that the professionals look slightly farther away from the predicted values than the undergraduates—similar to how managers performed slightly worse in some phases of the experiment in Bolton et al. (2012). Explicitly testing the difference in markups between the professionals and undergraduates, we see no statistically significant difference for most of the bidder-type/subsidy-level combinations; the one exception is the no subsidy condition for advantaged bidders (where the observed difference, 23.6% for undergraduates vs. 19.6% for professionals, is significant [$p = .009$] and farther away from the predicted value). The remarkable consistency observed between the two populations affirms that the bidding behavior between the two is comparable. Overall, this gives us great confidence in the results from Studies 1–3. It also reinforces that in a well-controlled lab setting, we observe fundamental behaviors that differ from what theory would predict. We anticipate that these same behaviors likely show up in real procurement auctions. An interesting next step might be measuring the size of these effects in a field experiment with active procurement auction bidders.

9 | CONCLUSION

The intent behind a preference auction is to help a weaker class of sellers compete against stronger sellers. A side benefit of preference auctions is that an increase in aggressive bidding by the stronger sellers, in response to the increased competition, can often offset the cost of to the preference shown to the weaker sellers. In this paper, we compare the predictions made for such subsidy auctions in Rothkopf et al. (2003) with experimental results—we analyze the convergence effects in bidding strategies, the effect of framing issues on bidding strategies, and how subsidy auction and price preference auction results compare.

Importantly, we do see profits improve for disadvantaged sellers in our experiments in most cases. So the transfer of wealth from the subsidy to the disadvantaged party seems to occur. Looking at all the results from our series of controlled experiments, bidding strategies (and thus realized profits) are slightly different from that predicted by theory. We find that sellers bid more

aggressively than optimal, sometimes to the point of realizing minimal (or no) profits. Cost-advantaged sellers seem to struggle more when the cost difference is less (d = 1.05), and cost-disadvantaged sellers seem to struggle more when the cost difference is more (d = 1.2). However, outside of that general aggressiveness, markup strategies respond to changes in the level of the subsidy and cost difference as predicted, in line with a convergence effect. As might be expected under overly aggressive bidding, we find that subsidies generally decrease the procurement cost for the auction buyer—the exception is our optimal (partial) subsidy case, which shows inconsistent effects. One plausible explanation is that a partial subsidy is more computationally challenging for sellers and that this extra burden resulted in different bidding behavior.

To understand the aggressive bidding and nuances of the subsidy mechanism versus other preference types, we run several additional experiments. We find framing effects for sellers that can lead sellers to suboptimal strategies. Disadvantaged sellers are more aggressive in bidding for large contracts when focused on the dollar amount, and advantaged sellers are more aggressive in bidding for small contracts when focused on percentage. We describe this anchoring behavior (similar to poker betting). This has implications in practice for two reasons: (1) virtually every auction asks for bids as a dollar amount and (2) auction models assume sellers bid based on stable percentage strategies and that bids are invariant to contract size. Here, we show that not to be the case. Importantly, if bidders are not using stable percentage bidding strategies in real life, this raises questions about the validity of auction results that rely on that assumption.

Finally, we show that equivalent preference mechanisms yield surprisingly different results. Sellers bid less aggressively when facing subsidies (especially partial) than they do when facing price preferences. The result of this is that price preference auctions result in lower profits for sellers (and lower costs for buyers). If the goal is to help disadvantaged classes of bidders, using price preference mechanisms seems to overstep. A particularly harsh critique of price preference procurement auctions might say they are designed to look equitable but are, in fact, too predatory once factoring in seller behavior. Our studies show that subsidies are likely a better mechanism to help disadvantage sellers (with buyer costs that are still lower than theory would predict)—but without crippling bidders with the (even more) aggressive bidding that accompanies price preference auctions. Some of the biggest preference auctions in the world (e.g., the Buy-American Act) use the price preference mechanism. Our experimental evidence suggests that while this may reduce the cost for the buyer, both advantaged and disadvantaged sellers may be hurt under this setup and, as a result of aggressive bidding by both parties, the favored class of bidders may not be helped as much as intended.

9.1 | Limitations

We attempt to overcome concerns around external validity by making our lab experiments mimic reality as much as possible. Our parameters and auction setup closely resemble actual DOT procurement auction data, but—as our results demonstrate—there may be differences as contract size and level of cost difference vary. Additional work could extend this work to other subsidy levels, settings, and setups. Second, we designed equivalent preference mechanisms to test in the lab, but it was necessary to construct these to match the theoretical model that served as the foundation for the comparisons. It may be that a more complete treatment could be modeled if it did not need to match the Rothkopf et al. (2003) benchmark. This would be an interesting additional direction in the future. Finally, we adopted a controlled experiment design using undergraduate students as our primary population. Even though we replicated treatments with a second sample of professionals for validation, extending the research to include different populations (e.g., procurement professionals) or field studies could be beneficial, especially for measuring specific effect sizes in particular industries or settings.

9.2 | Contributions

From a theoretical perspective, this research contributes to the field of supply chain management by providing a better understanding of procurement auctions and their efficiency in the context of preferential treatment for disadvantaged sellers during the procurement process. This applies not just to large government procurement processes. A survey of purchasing managers found that 31% of companies had adopted reverse auctions as part of their supply chain management strategies (Hannon, 2006). A key issue for the future integration of procurement auctions into supply chain management strategies is the study of efficient auction mechanisms and how they affect both suppliers and buyers (Chen et al., 2005).

Our empirical findings challenge established notions in operations management regarding equilibrium bidding strategies, particularly in asymmetric auction settings. By revealing that actual seller behavior deviates from the predictions of established economic models, this study

highlights the necessity to incorporate behavioral insights into operations management theories. Furthermore, the identification of a behavioral framing bias presents an important consideration for supply chain management theory, especially in the context of decision-making under uncertainty. This research also extends the theoretical discourse on the design of preference auctions by comparing two commonly used mechanisms (subsidies vs. price preferences) and their impact on auction outcomes. Such insights are crucial for developing more robust models in operations management that can accurately predict outcomes in procurement auctions, allowing for the design of auctions that are more aligned with actual bidder behavior and ensure a fair and competitive environment. The findings regarding subsidies versus price preferences are particularly salient, suggesting that subsidies might be a more effective mechanism for balancing the support of disadvantaged sellers with the maintenance of competitive procurement costs. This understanding enables managers to predict bidder responses more accurately, a crucial factor in strategic planning and decision-making. Additionally, these insights have implications for policy implementation in regulated sectors, guiding managers to achieve intended policy outcomes effectively. Finally, the nuanced understanding of auction dynamics can enhance vendor relationship management, particularly with disadvantaged suppliers, fostering a more diverse and resilient supplier base. Consequently, integrating these, insights can significantly contribute to achieving strategic goals of efficiency, cost-effectiveness, and equity in procurement processes. Overall, this study calls for a re-examination and potential refinement of existing supply chain management theories on procurement auctions to better align with empirical behaviors and practical realities.

9.3 | Recommendations

The big takeaway for managers (e.g., buyers in procurement auctions) is that using price preference auctions is even more challenging than previously thought. Choosing the level of subsidy has always been difficult for buyers since the true costs of the sellers are rarely known. Our findings suggest that sellers also encounter difficulties. This is especially true for partial (not 100%) subsidies. While prior literature suggests that subsidies less than 100% of the cost difference are optimal and can be more cost-effective for buyers while still improving the sellers' profit (relative to no subsidy), the same outcome did not show up in our experiments. Profit for disadvantaged bidders was near or even below the profit level with no subsidy (Table 2, Column 4e). The fact that disadvantaged bidders struggle with bids for large contracts when inputs are in dollar amounts (typical in real life)—often bidding much less because of our framing effect—could be a contributor.

Practically, we recommend several things. First, procurement auction buyers should be aware that preference mechanisms do impact sellers and change their behavior. Buyers should therefore be thoughtful and transparent in their policies. Second, even though partial subsidies yield optimal outcomes in theory, in real bidding (when sellers must make sense of a complicated bidding landscape and preference mechanism), full subsidies may result in better outcomes (especially in the case of larger cost differences). Third, buyers could reduce adverse auction outcomes by acknowledging predictable framing biases and designing systems that prevent disadvantaged sellers bidding too aggressively for extremely large contracts. And finally, if buyers truly want to help disadvantaged bidders, they should move away from price preference mechanisms—which seem to hamper sellers unnecessarily—and instead use subsidies to promote parity.

ORCID

Sanghoon Cho https://orcid.org/0000-0002-7480-885X
Joel O. Wooten https://orcid.org/0000-0002-9471-0817

ENDNOTES

[1] This paper improves on an early draft that was published as part of a conference proceeding (Cho et al. 2019). In that draft, the authors test various predictions of the Rothkopf et al. (2003) model, but the subsidy offered in the experiments deviates from that in the theoretical model. Here, we mimic the theoretical subsidy (as a percentage of the bid amount), which permits a more interesting set of comparisons for practice.

[2] See Rothkopf et al. (2003) for a nice overview of the models employing this representation.

[3] In practice, some markets have come up with mechanisms to allow bidders to escape the winner's curse. Dyer and Kagel (1996) document three such approaches, including withdrawn bids, relationship discounts, and change orders.

[4] Because of learning effects over time, balanced rounds are preferred. Rounds 26–30 exhibit consistent, expected trends in the experiments for which we have those data.

[5] In our experiments, there are 1549 small contracts and 2009 large contracts. There are more large contracts because disadvantaged costs are always larger than advantaged costs (the variable that is uniformly drawn from $100 to 400 k).

[6] The 0% subsidy case is not interesting, as neither seller has any preferential treatment from the buyer. The optimal subsidy case is interesting and can also be compared, thanks to equivalence in the mechanisms (detailed in Appendix C).

REFERENCES

Aloysius, J., Deck, C., Hao, L., & French, R. (2016). An experimental investigation of procurement auctions with asymmetric sellers. *Production and Operations Management*, *25*(10), 1763–1777.

Athey, S., Coey, D., & Levin, J. (2013). Set-asides and subsidies in auctions. *American Economic Journal: Microeconomics*, *5*(1), 1–27.

Ayres, I., & Cramton, P. (1996). Deficit reduction through diversity: How affirmative action at the FCC increased auction competition. *Stanford Law Review*, *48*, 761–815.

Benjamin, D. J., Brown, S. A., & Shapiro, J. M. (2013). Who is 'behavioral'? Cognitive ability and anomalous preferences. *Journal of the European Economic Association*, *11*(6), 1231–1255.

Bolton, G. E., & Katok, E. (2008). Learning by doing in the newsvendor problem: A laboratory investigation of the role of experience and feedback. *Manufacturing & Service Operations Management*, *10*(3), 519–538.

Bolton, G. E., Ockenfels, A., & Thonemann, U. W. (2012). Managers and students as newsvendors. *Management Science*, *58*(12), 2225–2233.

Brown, D. (2018). The effect of subsidized entry on capacity auctions and the long-run resource adequacy of electricity markets. *Energy Economics*, *70*, 205–232.

Carter, C. R., & Stevens, C. K. (2007). Electronic reverse auction configuration and its impact on buyer price and supplier perceptions of opportunism: A laboratory experiment. *Journal of Operations Management*, *25*(5), 1035–1054.

Chang, C. T. (2008). To donate or not to donate? Product characteristics and framing effects of cause-related marketing on consumer purchase behavior. *Psychology & Marketing*, *25*(12), 1089–1110.

Chaturvedi, A., Katok, E., & Beil, D. R. (2019). Split-award auctions: Insights from theory and experiments. *Management Science*, *65*(1), 71–89.

Chen, R. R., Roundy, R. O., Zhang, R. Q., & Janakiraman, G. (2005). Efficient auction mechanisms for supply chain procurement. *Management Science*, *51*(3), 467–482.

Cho, S., Donohue, J., Fry, T., & Wooten, J. (2019). The impact of subsidies in asymmetric procurement auctions: A comparison of theoretical and experimental results. *IFAC-PapersOnLine*, *52*(13), 1990–1995.

Corns, A., & Schotter, A. (1999). Can affirmative action Be cost effective? An Experimental Examination of Price-Preference Auctions. *American Economic Review*, *89*(1), 291–305.

Deck, C., & Jahedi, S. (2015). The effect of cognitive load on economic decision making: A survey and new experiments. *European Economic Review*, *78*, 97–119.

Deltas, G., & Engelbrecht-Wiggans, R. (2005). Naive bidding. *Management Science*, *51*(3), 328–338.

Denes-Raj, V., & Epstein, S. (1994). Conflict between intuitive and rational processing: When people behave against their better judgment. *Journal of Personality and Social Psychology*, *66*(5), 819–829.

Dyer, D., & Kagel, J. H. (1996). Bidding in common value auctions: How the commercial construction industry corrects for the Winner's curse. *Management Science*, *42*(10), 1463–1475.

Dyer, D., Kagel, J. H., & Levin, D. (1989). A comparison of naive and experienced bidders in common value offer auctions: A laboratory analysis. *The Economic Journal*, *99*(394), 108–115.

Eckerd, S., DuHadway, S., Bendoly, E., Carter, C. R., & Kaufmann, L. (2021). On making experimental design choices: Discussions on the use and challenges of demand effects, incentives, deception, samples, and vignettes. *Journal of Operations Management*, *67*(2), 261–275.

Elmaghraby, W. (2007). Auctions within E-sourcing events. *Production and Operations Management*, *16*(4), 409–422.

Elmaghraby, W. J., Katok, E., & Santamaría, N. (2012). A laboratory investigation of rank feedback in procurement auctions. *Manufacturing & Service Operations Management*, *14*(1), 128–144.

Feltz, C. J., & Miller, G. E. (1996). An asymptotic test for the equality of coefficients of variation from K populations. *Statistics in Medicine*, *15*(6), 647–658.

Fischbacher, U. (2007). Z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, *10*(2), 171–178.

Flambard, V., & Perrigne, I. (2006). Asymmetry in procurement auctions: Evidence from snow removal contracts. *The Economic Journal*, *116*(514), 1014–1036.

Forde, M. (2019). Reverse auctions on the rise among younger procurement professionals. In *Supply chain dive*. https://www.supplychaindive.com/news/reverse-auctions-rise-younger-procurement-professionals/558582

Fréchette, G. R. (2012). Laboratory experiments: Professionals versus students. In G. Frechette & A. Schotter (Eds.), *The methods of modern experimental economics*. Oxford University Press.

Grosskopf, B., Rentschler, L., & Sarin, R. (2018). An experiment on first-price common-value auctions with asymmetric information structures: The blessed winner. *Games and Economic Behavior*, *109*, 40–64.

Hannon, D. (2006). E-sourcing finds its place. *Purchasing magazine, November*, 2.

Haruvy, E., & Katok, E. (2013). Increasing revenue by decreasing information in procurement auctions. *Production and Operations Management*, *22*(1), 19–35.

He, H., & Chen, Y. (2021). Auction mechanisms for allocating subsidies for carbon emissions reduction: An experimental investigation. *Social Choice and Welfare*, *57*, 387–430.

Hellerstein, D., Higgins, N., & Roberts, M. J. (2014). *Using quotas to enhance competition in asymmetric auctions: A comparison of theoretical and experimental outcomes*. College of Social Sciences.

Hite, P. A., & Roberts, M. L. (1991). An experimental investigation of taxpayer judgements on rate structure in the individual income tax system. *Journal of the American Taxation Association*, *13*(2), 47–63.

Hon-Snir, S., Monderer, D., & Sela, A. (1998). A learning approach to auctions. *Journal of Economic Theory*, *82*(1), 65–88.

Jap, S. D. (2002). Online reverse auctions: Issues, themes, and prospects for the future. *Journal of the Academy of Marketing Science*, *30*(4), 506–525.

Jeitschko, T. D. (1998). Learning in sequential auctions. *Southern Economic Journal*, 65(1), 98–112.

Katok, E. (2018). Designing and conducting laboratory experiments. In *The handbook of behavioral operations* (pp. 1–33). Wiley Blackwell.

Kirwan, B. E. (2005). *Essays on US agricultural policy: Subsidies, crop insurance, and environmental auctions*. (Doctoral dissertation). Massachusetts Institute of Technology.

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

Kwasnitschka, D., Franke, H., & Netland, T. H. (2024). Effects of feedback in manufacturing: A field experiment using smartwatch technology. *Journal of Operations Management*, *70*, 933–956.

Levin, I. P., Schneider, S. L., & Gaeth, G. J. (1998). All frames are not created equal: A typology and critical analysis of framing effects. *Organizational Behavior and Human Decision Processes*, *76*(2), 149–188.

Lonati, S., Quiroga, B. F., Zehnder, C., & Antonakis, J. (2018). On doing relevant and rigorous experiments: Review and recommendations. *Journal of Operations Management*, *64*, 19–40.

Lopomo, G., Marx, L. M., McAdams, D., & Murray, B. (2011). Carbon allowance auction design: An evaluation of the current debate. *Review of Environmental Economics and Policy*, *5*(1), 25–43.

Manuel, K. M. (2016). The buy American act–preferences for" domestic" supplies. In *Brief* (pp. 7–5700). Congressional Research Service.

Matthäus, D. (2020). Designing effective auctions for renewable energy support. *Energy Policy*, *142*, 111462.

McAfee, R. P., & McMillan, J. (1989). Government procurement and international trade. *Journal of International Economics*, *26*(3–4), 291–308.

Milgrom, P. R. (1981). Rational expectations, information acquisition, and competitive bidding. *Econometrica: Journal of the Econometric Society*, *49*, 921–943.

Mir-Artigues, P., & Del Río, P. (2014). Combining tariffs, investment subsidies and soft loans in a renewable electricity deployment policy. *Energy Policy*, *69*, 430–442.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, *6*(1), 58–73.

Nakabayashi, J. (2013). Small business set-asides in procurement auctions: An empirical analysis. *Journal of Public Economics*, *100*, 28–44.

Neugebauer, T., & Perote, J. (2008). Bidding 'as if' risk neutral in experimental first Price auctions without information feedback. *Experimental Economics*, *11*(2), 190–202.

Pai, M. M., & Vohra, R. (2012). Auction design with fairness concerns: Subsidies vs. set-asides (No. 1548). Discussion Paper.

Quiroga, B. F., Moritz, B. B., & Guide, V. D. R., Jr. (2021). The role of transparency in procurement: Revealed versus concealed scoring rules in sealed bid a+ B auctions. *Journal of Operations Management*, *67*(1), 71–81.

Rosston, G. L., & Wallsten, S. J. (2011). The path to universal broadband: Why we should Grant low-income subsidies and use experiments and auctions to determine the specifics. *The Economists' Voice*, *8*(1).

Rothkopf, M. H. (1969). A model of rational competitive bidding. *Management Science*, *15*(7), 362–373.

Rothkopf, M. H., Harstad, R. M., & Fu, Y. (2003). Is subsidizing inefficient bidders actually costly? *Management Science*, *49*(1), 71–84.

Sorana, V. (2000). Auctions for universal services subsidies. *Journal of Regulatory Economics*, *18*(1), 33–58.

Soundararajan, V., & Brammer, S. (2018). Developing country sub-supplier responses to social sustainability requirements of intermediaries: Exploring the influence of framing on fairness perceptions and reciprocity. *Journal of Operations Management*, *58*, 42–58.

Srinivasan, K., & Wang, X. (2010). Commentary—Bidders' experience and learning in online auctions: Issues and implications. *Marketing Science*, *29*(6), 988–993.

Tunca, T. I., Wu, D. J., & Zhong, F. (2014). An empirical analysis of price, quality, and incumbency in procurement auctions. *Manufacturing & Service Operations Management*, *16*(3), 346–364.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*(4), 297–323.

U.S. Small Business Administration. (2024, April 29). *Biden-Harris Administration Awards Record-Breaking $178 Billion in Federal Procurement Opportunities to Small Businesses* [Press release]. https://www.sba.gov/article/2024/04/29/biden-harris-administration-awards-record-breaking-178-billion-federal-procurement-opportunities

Wooten, J. O., Cho, S., Fry, T. D., & Donohue, J. M. (2022). Overcoming cost disadvantages in procurement auctions. *Decision Sciences*, *53*(3), 486–513.

Wooten, J. O., Donohue, J. M., Fry, T. D., & Whitcomb, K. M. (2020). To thine own self Be true: Asymmetric information in procurement auctions. *Production and Operations Management*, *29*(7), 1679–1701.

Zeithammer, R., & Adams, C. (2010). The sealed-bid abstraction in online auctions. *Marketing Science*, *29*(6), 964–987.

How to cite this article: Cho, S., Wooten, J. O., & Fry, T. D. (2024). Helping hampered bidders—Do subsidy auctions work as intended? *Journal of Operations Management*, *70*(7), 1126–1154. https://doi.org/10.1002/joom.1333

APPENDIX A: SAMPLE PARTICIPANT INSTRUCTIONS

Treatment 2: Cost disadvantaged—Optimal subsidy (71%), $ bid.

A.1. | Introduction

This is an experiment on the economics of decision-making in auctions. Each participant will receive $10 for completing the experiment. In addition, you will be rewarded for how well you do in the auctions. You will start the experiment with a balance of $500,000. At the end, the computer will take your balance and convert it to US dollars at the rate of 2 cents per $1000 for any amount over $400,000 ($2 for every $100,000). The more money you make in the auctions, the more cash you get.

- It is important that you remain silent and do not look at other participants' work.
- If you need assistance of any kind during the experiment, please raise your hand.

A.2. | The experiment

There will be 30 auctions. In each auction, you will be randomly paired with another participant. Each of you will be acting in the role of a construction company competing for the rights to build a bridge for the Department of Transportation (DOT). Your task is to *submit a bid* indicating the amount that you will charge to complete the project. Neither bidder will know the exact cost to build the bridge, but each of you will receive *an estimate of the project cost* prepared by your company. The DOT will award the project to the lowest bidder and pay them their bid amount. Your bid should consist of an amount that covers the actual costs of bridge construction plus any profit you wish to make.

The first 5 rounds will be practice runs to get familiar with the process. After that, your balance will be affected by your decisions. Your goal is to increase your cash balance as much as possible.

There are strategic tradeoffs in each auction. The lower you bid, the more likely you will win the auction. However, if you win with a very low bid, the likelihood that you lose money increases. Similarly, the higher you bid, the more profit—but the less likely you are to win the auction. Keep in mind, there is no reward for winning with a very low bid; *your goal is to make as much profit as possible.*

A.3. | The details

There are important differences between you and your competitor:

1. **You are a small business and have higher construction costs.** You know your competitor uses cheaper labor and different machinery, so they can do a project for less than you can. Based on past-experience, you believe it will **cost you 20% more** than it will cost them to do the same work.
2. **If you win, the DOT will pay you a subsidy.** In an effort to promote small business, the DOT will award you additional payment if you win a contract. If you submit the minimum bid in a given round, the subsidy paid to you will equal **14.2%** of your bid. Thus, you would get your bid amount *plus* the extra subsidy payment upon winning.

In each round, you and the participant with whom you are paired for that auction will each receive an *estimate of the cost* to complete that project. The estimates will be randomly generated from a bell-shaped distribution that is centered on each company's actual cost to build that bridge. The cost estimate is equally likely to be *higher* or *lower* than your actual cost. Based on the bidding history of you and your competitor, you are confident that the accuracies of your estimates are the same; the relative standard deviation of your cost estimate is ∼**10%**.

For a hypothetical bridge project that costs you $120,000 to complete, the cost estimates could be in the ranges below:

	You	Competitor
Actual project cost	**$120,000**	**$100,000**
Profit (if winner)	Bid − actual cost + bid*(0.142)	Bid − actual cost
Cost estimate		
68% chance estimate between	$108,000 and 132,000	
95% chance estimate between	$96,000 and 144,000	
99+% chance estimate between	$84,000 and 156,000	
As an example, we illustrate how this hypothetical auction might play out		
- If you won with a bid of $115,000, your profit would be: ($115,000–120,000 + 16,330) = $11,330		
- If you won with a bid of $108,000, your profit would be: ($108,000–120,000 + 15,336) = $3,336		
- If you won with a bid of $100,000, your profit would be: ($100,000–120,000 + 14,200) = − $5,800		

The above values are for an example project that costs you $120,000. All of these values will change as the actual project costs change each round.

After each auction, the following results will be revealed:

After each auction, the following results will be revealed		
Participant	Cash balance	Information seen
Winning bidder	Profit added or loss subtracted	All bids, project cost, profit (or loss)
Losing bidder	No change	All bids

A.4. | Reminders

- You are trying to make as much profit as possible.
- You have higher costs than your competitor.
- You are being subsidized by the state DOT.
- The cost estimate you receive each round could be higher or lower than your actual cost.
- Your bid should cover the actual costs of bridge construction plus any profit you wish to make.
- When submitting bids, you want to be below your competitor but *NOT* so low that you lose money.

Treatment 10: Cost advantaged—Optimal price preference (71%), $ Bid.

A.5. | Introduction

This is an a-making in auctions. Each participant will receive $10 for completing the experiment. In addition, you will be rewarded for how well you do in the auctions. You will start the experiment with a balance of $500,000. At the end, the computer will take your balance and convert it to US dollars at the rate of 2 cents per $1000 for any amount over $400,000 ($2 for every $100,000). The more money you make in the auctions, the more cash you get.

- It is important that you remain silent and do not look at other participants' work.
- If you need assistance of any kind during the experiment, please raise your hand.

A.6. | The experiment

There will be 30 auctions. In each auction, you will be randomly paired with another participant. Each of you will be acting in the role of a construction company competing for the rights to build a bridge for the Department of Transportation (DOT). Your task is to *submit a bid* indicating the amount that you will charge to complete the project. Neither bidder will know the exact cost to build the bridge, but each of you will receive *an estimate of the project cost* prepared by your company. The DOT will award the project to one bidder and pay them their bid amount. Your bid should consist of an amount that covers the actual costs of bridge construction plus any profit you wish to make.

The first five rounds will be practice runs to get familiar with the process. After that, your balance will be affected by your decisions. Your goal is to increase your cash balance as much as possible.

There are strategic tradeoffs in each auction. The lower you bid, the more likely you will win the auction. However, if you win with a very low bid, the likelihood that you lose money increases. Similarly, the higher you bid, the more profit—but the less likely you are to win the auction. Keep in mind, there is no reward for winning with a very low bid; *your goal is to make as much profit as possible.*

A.7. | The details

There are important differences between you and your competitor:

1. **You are a large company and have lower construction costs.** You know your competitor uses more expensive labor and different machinery, so you can do a project for less than they can. Based on past-experience, you believe it will **cost them 20% more** than it will cost you to do the same work.
2. **The DOT is helping your competitor be more competitive.** In an effort to promote small business, the DOT will add **14.2%** to *your* bid before choosing the winner. The DOT will then compare your adjusted bid and your competitor's entered bid and pick the smaller of the two. The winning bidder will receive their entered bid amount. (e.g., If you bid $100,000 and your competitor bids $115,000, the DOT would compare $114,200 and $115,000, choose your bid, and pay you $100,000.)

In each round, you and the participant with whom you are paired for that auction will each receive an *estimate of the cost* to complete that project. The estimates will be randomly generated from a bell-shaped distribution that is centered on each company's actual cost to build that bridge. The cost estimate is equally likely to be *higher* or *lower* than your actual cost. Based on the bidding history of you and your competitor, you are confident that the accuracies of your estimates are the same; the relative standard deviation of your cost estimate is ∼**12%**.

For a hypothetical bridge project that costs you $100,000 to complete, your cost estimates could be in the ranges below:

	You	Competitor
Actual project cost	**$100,000**	**$120,000**
Decision-based on	Bid + bid*(0.142)	Bid
Profit (if winner)	Bid − actual cost	Bid − actual cost
Cost estimate		
68% chance estimate between $88,000 and 112,000		
95% chance estimate between $76,000 and 124,000		
99+% chance estimate between $64,000 and 136,000		
As an example, we illustrate how this hypothetical auction might play out:		
- If you won with a bid of $115,000, your profit would be: ($115,000–100,000) = $15,000		
- If you won with a bid of $105,000, your profit would be: ($105,000–100,000) = $5,000		
- If you won with a bid of $95,000, your profit would be: ($95,000–100,000) = − $5,000		
The above values are for an example project that costs you $100,000. All of these values will change as the actual project costs change each round.		

The above values are for an example project that costs you $100,000. All of these values will change as the actual project costs change each round.

After each auction, the following results will be revealed:

After each auction, the following results will be revealed:		
Participant	Cash balance	Information seen
Winning bidder	Profit added or loss subtracted	All bids, project cost, profit (or loss)
Losing bidder	No change	All bids

A.8. | Reminders

- You are trying to make as much profit as possible.
- You have lower costs than your competitor.
- You competitor is being helped with a price preference by the state DOT.
- The cost estimate you receive each round could be higher or lower than your actual cost.
- Your bid should cover the actual costs of bridge construction plus any profit you wish to make.
- When submitting bids, you want to be below your competitor but *NOT* so low that you lose money.

APPENDIX B: SAMPLE SCREENSHOTS FROM EXPERIMENT

Treatment 2: Cost disadvantaged—Optimal subsidy (71%), $ bid.

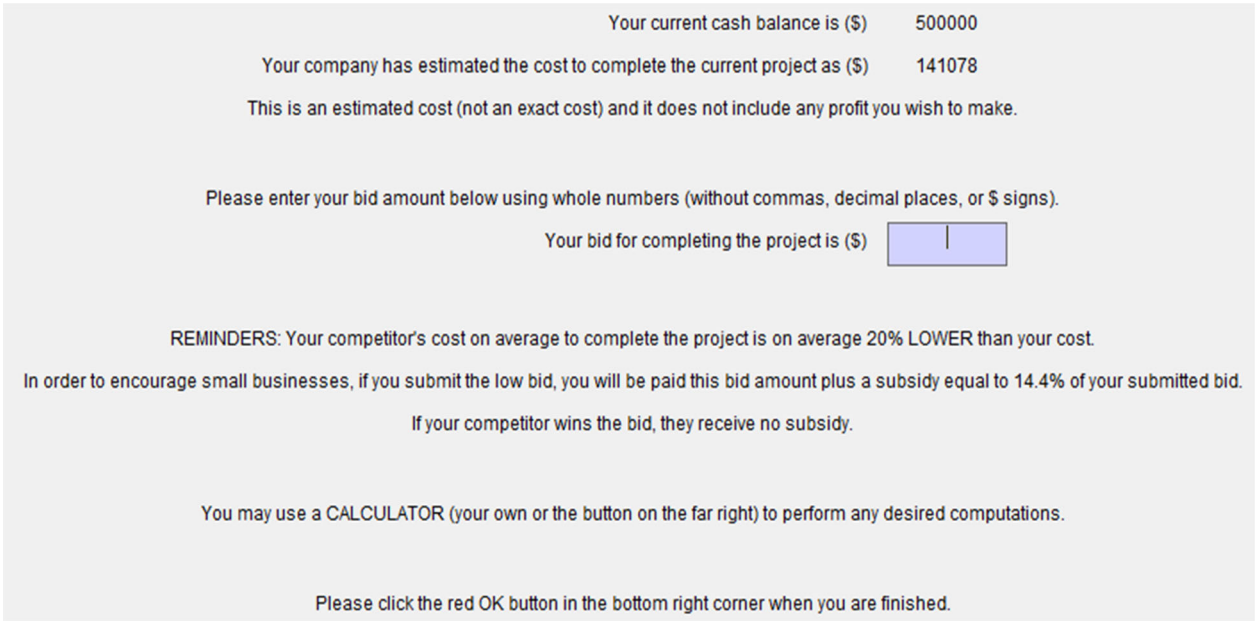

Results (disadvantaged winner).

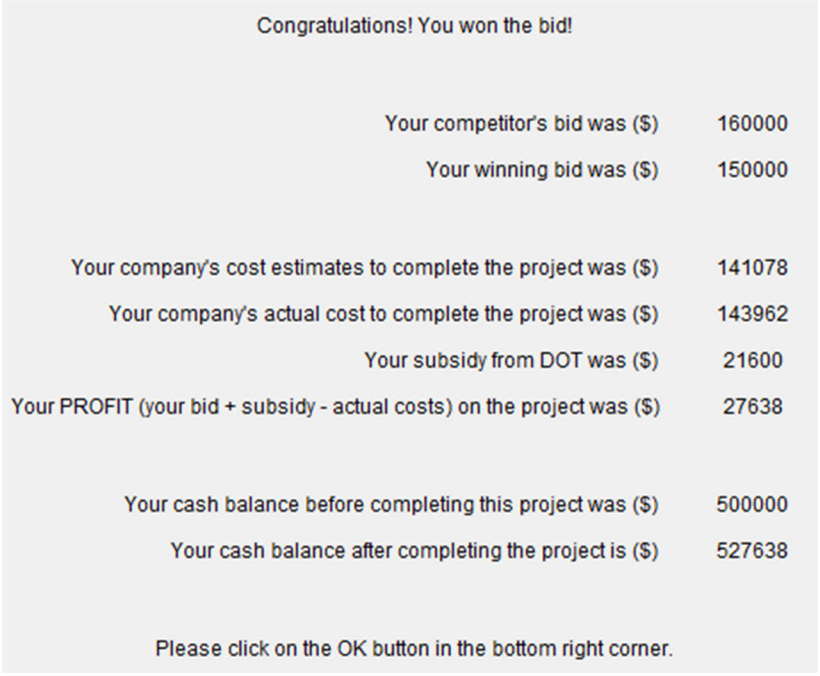

Treatment 10: Cost advantaged—Optimal price preference (71%), $ bid.

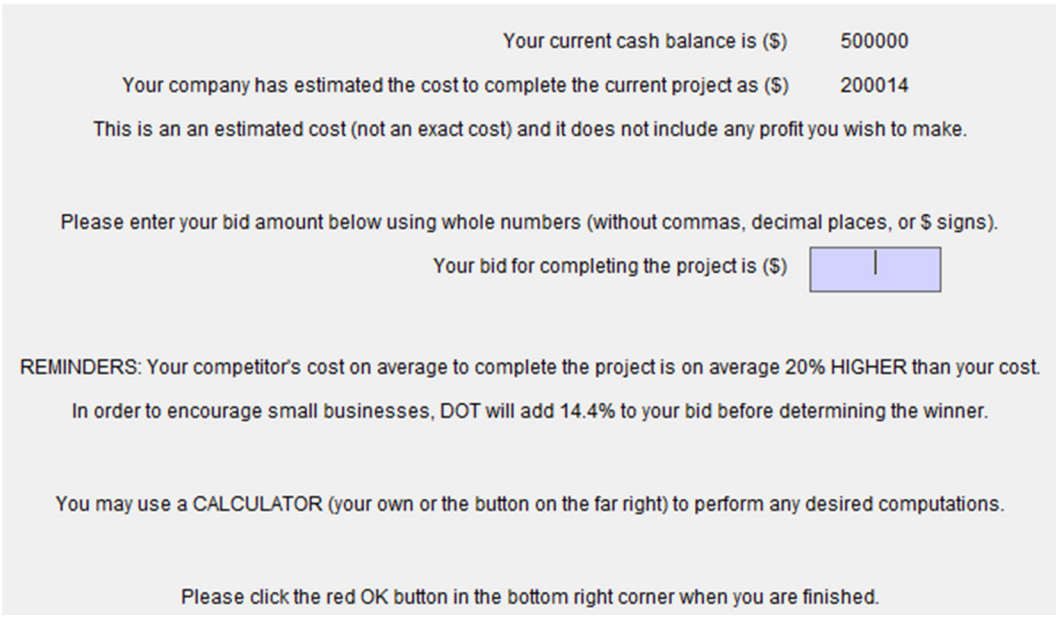

Results (advantaged winner).

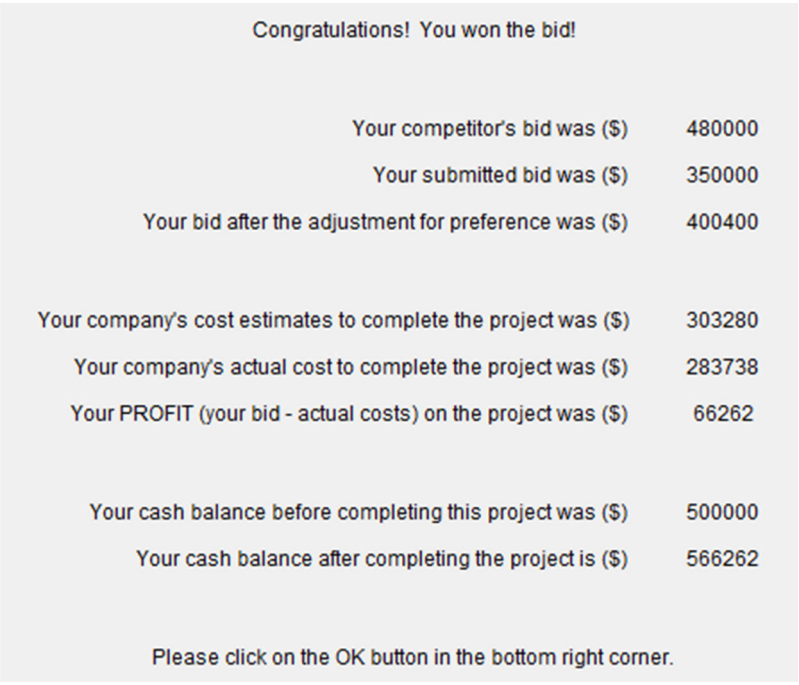

Treatment 8: Cost disadvantaged—Optimal subsidy (71%), % bid.

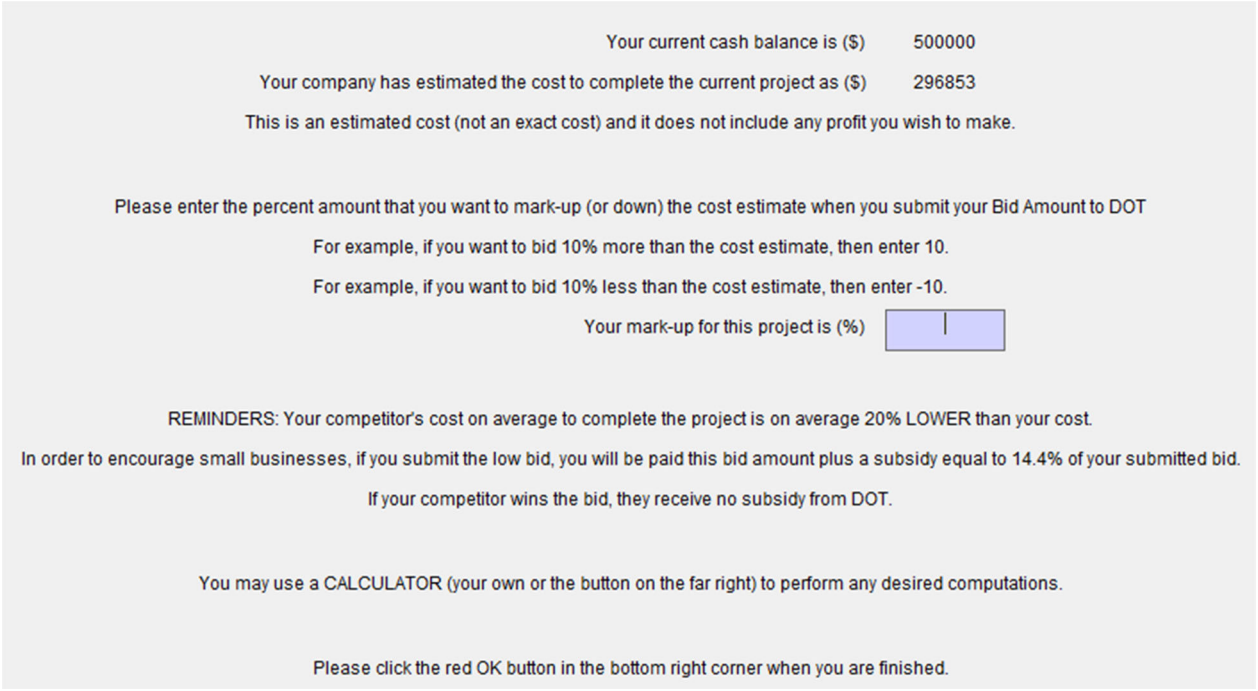

Results (disadvantaged winner).

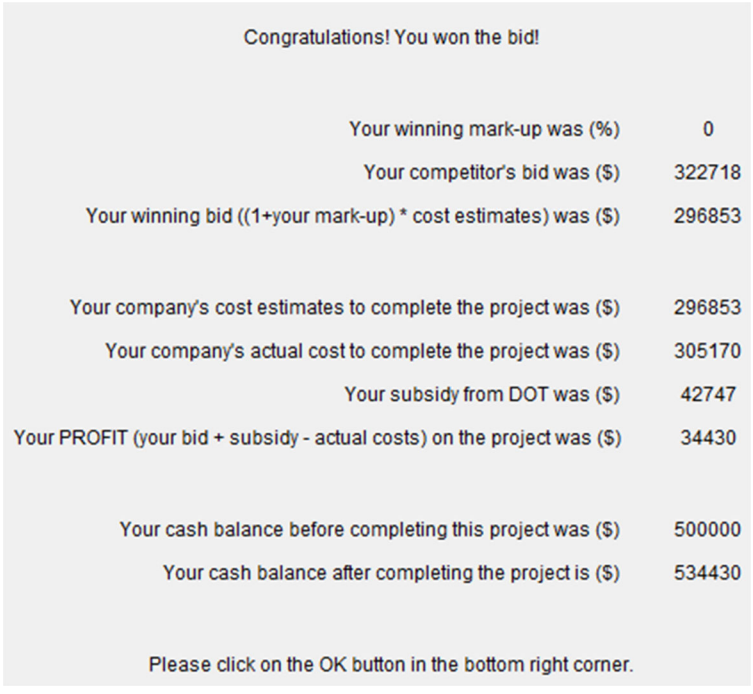

APPENDIX C: SUBSIDY VERSUS PRICE PREFERENCE MECHANISM—PARTIAL SUBSIDY EQUIVALENCE

Here, we show that subsidy auctions are equivalent to price preference actions in the optimal subsidy case. The most notable wrinkles (from Table 4) are that (1) different markup strategies are shown for the two bidders to match the theoretical optimal values and (2) the adjusted cost basis for the Disadvantaged seller is still proportionally reduced but the auction decision costs are not equal (but are equivalent in the two mechanisms)—the breakpoint, decisions, and costs are the same. Thus, regardless of the markup strategy chosen by each seller type, the buyer cost of these two preference mechanisms will be equivalent.

	Optimal subsidy s=14.2%, d=1.2			Optimal price preference s=14.2%, d=1.2	
	Adv.	Disadv.		Adv.	Disadv.
Cost estimate	100	120		100	120
Markup (21.2%, 17.9%)	21.2%	3.2%	$Seller_{DIS}$ bids as if $Seller_{ADV}$: $C=C_H/(1+s)$	21.2%	17.9%
Seller bid	121.2	123.9		121.2	141.5
Auction decision	**121.2**	**123.9**	Buyer adjusts $Seller_{ADV}$ bid: $(1+s)*Bid$	**138.4**	**141.5**
Buyer cost	**121.2**	**141.5**	Buyer adjusts $Seller_{DIS}$ payout: $(1+s)*Bid$	**121.2**	**141.5**

APPENDIX D: LEARNING EFFECTS

Figure D1 shows the average markup across the four subsidy levels by round for our scenarios. As already seen in the literature, sellers in the lab are more aggressive than the theory would predict. Over time, however, participants tend to become less aggressive, submitting markups that approach predictions as the experiment progresses, similar to prior research (e.g., Elmaghraby et al., 2012; Wooten et al., 2020). To check the statistical effect, we conduct the same test as Bolton and Katok (2008) and find a significant trend (i.e., slope of improvement) in the bidder's markups (ordinary least-squares, two-tailed $p < .01$) consistent with our expectation, except in the case of observed advantaged with $d = 1.2$.

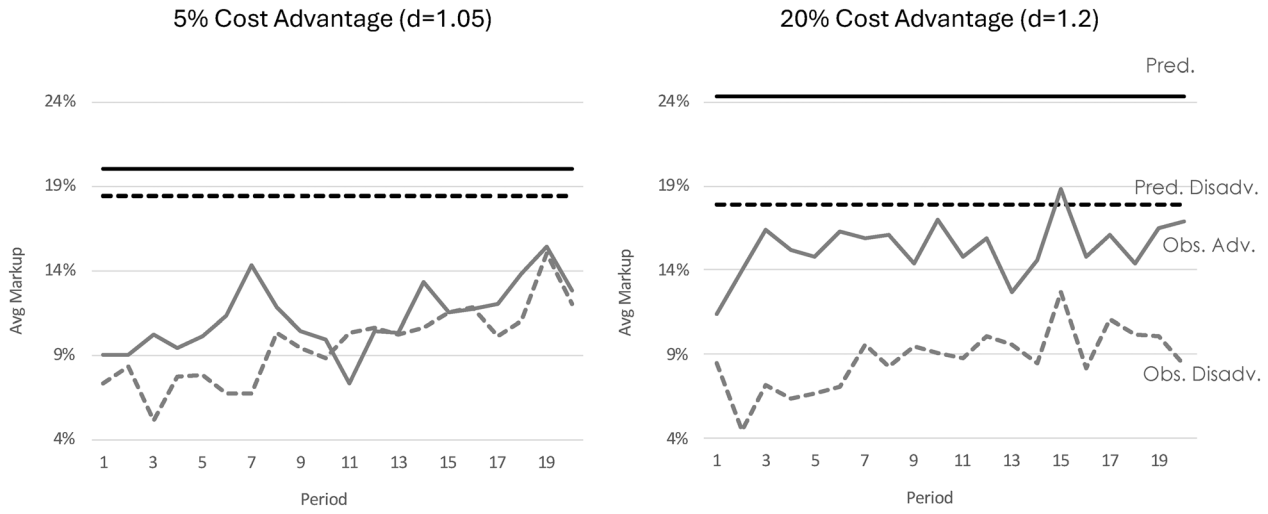

FIGURE D1 Observed versus predicted bidding strategies over time. Slope of improvement is significant (<.01) for all observed types except observed advantaged when $d = 1.2$.