The Role of Transparency in Multiattribute Sealed-bid Procurement Auctions

Bernardo F. Quiroga[*]
Clemson University

Brent B. Moritz, V. Daniel R. Guide, Jr.
The Pennsylvania State University at University Park

We study multiattribute procurement auctions, sometimes known as A+B auctions. Here, buyers request offers on attributes such as price and quality, and bidders choose the amounts of quality to offer and how much to charge for it. We analyze two sealed-bid scenarios: One where the scoring rule that weights quality and price is explicitly communicated to bidders before they submit their offers, and another where the rule is only known to the buyer and not to the bidders. In addition, we compare behavior where the scoring rule is made visible after the offers are submitted. Our experimental results show substantial losses to the buyer as a direct effect of transparency loss, while sellers see their profits increased in the same situation.

Keywords:
multi-dimensional procurement; auctions; transparency; market design.
History:
Original version, April 2015. This revision, May 2018. Please do not circulate or cite without the authors' explicit permission.

[*] Corresponding author, e-mail: bquirog@g.clemson.edu. Acknowledgements section pending end of peer review process.

1. Introduction

This article investigates the impact of a buyer's commitment to award contracts based on an announced scoring rule in a sealed-bid multiattribute procurement auction. In this setting, each supplier submits an offer for the requested good or service, including both the monetary price and a non-monetary contractible attribute valued by the buyer. Using a controlled laboratory experiment, we compare the performance differences when buyers either expressly communicate the weight placed on non-price attributes or when they hide this weight from bidders in advance of a procurement auction. When suppliers have to place bids that include both price and a non-monetary attribute, we analyze the effect of buyer transparency in the communication of the criterion used to evaluate bids on bidding decisions. Those decisions, in turn, impact the buyer's surplus and the suppliers' profits.

The question about the impact of using a clear and transparent scoring rule on auction performance originated during conversations with members of the research division of *ChileCompra*, the governmental procurement agency of Chile. Somewhat surprisingly, we found relatively little literature to assist in answering this question.

In many governmental procurement settings, buyers use sealed-bid multi-dimensional procurement mechanisms to purchase goods or services. Although nearly all auctions consider the purchase price as a means to rank offers, frequently other dimensions such as time to completion, supplier capability, product quality or firm reputation may be considered by the purchaser. Some of these factors are exogenous, such as geographic proximity, brand name or prior experience, whereas other factors, such as lead time, validation/inspection or local delivery, can be impacted by the choices made by suppliers in preparing the bid and fulfilling the contract. When buyers such as government procurement agencies issue a request for proposals (RFP), suppliers typically are asked to respond with a bid that includes price and other attributes. The conditions of the RFP list whether the buyer will decide based on a specific, publically-announced scoring rule based on the bids (a scoring auction), or if the final award of the contract remains at the discretion of the buyer.

Though prior literature has investigated procurement auctions, our work specifically investigates the impact of buyer transparency in multi-dimensional auctions, where supplier selection includes a non-price attribute with the submitted bid. Che (1993) was the first to characterize the optimal mechanism design of the US Department of Defense's (DoD) weapon system procurement practices as a multi-dimensional sealed-bid auction. By characterizing its analytical properties, Che (1993) helped popularize the DoD's elaborate scoring system that evaluated suppliers' bids, weighing price in combination with other key performance/quality metrics which resulted in a competitive sourcing process. Asker and Cantillon (2008) showed that fixing the quality characteristics at a level, and assigning according to prices only is worse for the buyer than using a score auction where the score reveals the buyer's preference for quality. Aside from the DoD, other agencies also have used similar procurement mechanisms. Several studies (e.g., Snir and Gupta 2011, Lewis and Bajari 2011, Bajari et al. 2014, Gupta et al. 2015) document how US state departments of transportation use scoring rules to rank multi-dimensional bids. These are often called "A+B bidding auctions", where "A" is a price bid, and "B" is some other performance metric such as accelerated time to complete the contracted product or service. For example, the quality dimension in Lewis and Bajari (2011) included factors such as the number of days a road was closed to complete construction, which was weighted along with the bid price. Our research focuses on the differences in outcome for the buyer, and provides specific recommendations on the value of developing and implementing a specific scoring rule.

This research sits at the intersection of operations and public policy (see Joglekar et al. 2016), because the auction mechanism design selected by the buyer has both cost and outcome implications for procurement agencies and for suppliers. Because it uses a controlled laboratory experiment, our research provides one way to mitigate the *Lucas critique* (Lucas 1976), as we explicitly acknowledge that the decision rules of economic agents vary systematically with policy changes. By creating parallel scenarios uncontaminated by one another, our research can inform policy-makers in their procurement design decisions.

From a policy perspective, transparent scoring and awarding of contracts may be preferred to reduce the appearance of corruption or improper influence, especially in public procurement. However, rather than committing to an ex-ante fixed scoring rule, buyers may prefer to have some discretion selecting the supplier for several reasons. Accurately specifying all contract attributes or factors and their weights ahead of receiving the offers is a time-consuming and sometimes expensive process to perform ahead of requesting bids. For example, before announcing a winner, buyers may prefer to consider if a winning supplier was perceived to be "easy to do business with" on past contracts, or if that supplier was able to successfully complete past contracts with a minimum of additional supervision by the buyer, without explicitly communicating that preference or precisely how that preference is measured. Additionally, a buyer may wish to give additional weight to a firm's commitment to hire local employees, subcontract or source from local businesses or give special consideration to firms owned by under-represented populations; yet the express inclusion of, or weight given to these factors in a fixed scoring rule can be difficult to quantify or be controversial. Thus, it is not surprising that in procurement auctions, the trend has been for an increasing amount of information going to the buyer even as there has been less information going to supplier(s); the exact selection criteria often remains unclear (Elmaghraby 2007).

The key research question is how the buyer's public commitment to a specific multi-attribute scoring rule impacts buyer surplus (or buyer utility) and overall wealth (for bidder and buyer combined) through bidding behavior, in contrast to when this scoring rule is concealed from bidders. To answer this question, we designed two laboratory experiments to study multi-dimensional (A+B) sealed-bid mechanisms. We compare two main scenarios: First, a sealed-bid *transparent scores* auction (TS); and second, a sealed-bid *multi-dimensional beauty-contest* auction (MDBC). In study 2, we consider a special case of the MDBC auction where the evaluation rule was revealed after bids were received, which we termed an *ex-post transparent* (EPT) auction. In the TS auction, suppliers know both their production cost and the exact evaluation rule used by the buyer to rank bids before they submit them (Che 1993, extended by Asker and Cantillon 2008). In contrast, the evaluation rule used by the buyer in a MDBC auction is concealed from the bidder. The MDBC is a "beauty-contest auction", because suppliers choose levels of quality and price to try to maximize their expected utility given their estimate of the buyer's preference for quality.

Our work is the first to empirically compare the use of transparent and concealed scoring rules in multi-dimensional procurement bidding decisions using sealed-bid first scores. We investigate the costs associated with a buyer concealing information on the assignment of multi-dimensional, quality-dependent contracts, as opposed to using a pre-revealed rule. Our results show a substantial decrease in buyer's surplus when the buyer uses a non-transparent rule as opposed to when the rule is publicly announced. On the other hand, sellers obtained higher profits in the non-transparent scenario, due to less competitive bidding.

2. Motivation, Approach and Contributions

Motivation. This research originated in response to a specific question from a public procurement agency in the nation of Chile. Created in 2003, *ChileCompra* is the central agency for governmental procurement (Gur et al. 2013, Corvalán 2013, Reyes 2015). ChileCompra is supervised by the President, has a publically-appointed Director and is part of the Chilean Ministry of Finance. It regulates all public sector auctions and framework agreements, and over 850 governmental offices use ChileCompra's electronic procurement platform to purchase and contract goods and services from over 100,000 pre-qualified suppliers and vendors. These transactions amount to approximately 3.5% of Chile's annual GDP. Specifically, ChileCompra wanted to know if the considerable efforts and expense they undertook to develop and publicize transparent scoring rules for procurement auctions improved outcomes, or if they would be better off implementing nonbinding or buyer-determined mechanisms.

One of the explicit goals for ChileCompra was to increase transparency in governmental acquisitions while generating savings and efficiency due to supplier competition. As in the US state departments of transportation (e.g., Lewis and Bajari 2011) and other examples, the communication of transparency and fairness in contract assignments has been one of ChileCompra's historical institutional concerns since its foundation. To accomplish these goals, ChileCompra uses multi-dimensional score auctions, incorporating price and one or more measures of quality, to procure goods and services for public agencies.

We had seven meetings with the chief manager and two analysts of the research division[1] at ChileCompra. The meetings were aimed at understanding how ChileCompra conducts their auctions, as well as discussing specific concerns they had about the way the agency operates. Specifically, ChileCompra had made substantial efforts to be transparent in their assignments, communicating criteria, revealing information, and avoiding, not only corruption, but the *perception* of corruption. Yet these efforts were costly, requiring substantial resources to establish valid, precise and transparent scoring rules prior to opening the bidding. The main question discussed with the studies division at ChileCompra was whether there in fact was economic and social value in doing all the pre-work required for transparent procurement rather than awarding contracts in a more discretional manner. In the words of the chief manager: "*We want to know what the price of transparency is. What is the cost associated to make scores and bids transparent in our*

[1] The seven meetings took place between December 21st 2011 and June 4th 2013. The names of the participants from ChileCompra will be included at the start of the acknowledgements section.

procurement processes? By concealing this information, does our buyer's surplus become worse or better?" In sum, their question was whether it was worthwhile to make costly efforts to conduct transparent auctions. Although there was substantial prior literature on procurement auctions, we are not aware of any research studies that would assist in addressing this specific question in the context of sealed-bids.

Note that though our research is inspired by ChileCompra, many governmental agencies, private universities, and companies use similar types of contract assignment mechanisms, especially for large and complex projects. For example, state departments of transportation in the U.S. frequently assign contracts using these types of mechanisms (Lewis and Bajari 2011, Gupta et al. 2015, Bajari et al. 2014). Government procurement auctions frequently use sealed bid systems, and government purchases are said to account for ten percent of gross domestic product (McAfee and McMillan 1987). Private firms also use sealed bid multidimensional auctions, for example, to select among insurance companies to provide a corporate health insurance package plan for their employees, where insurers bid on multi-dimensional provisions as well as money transfers (Zheng 2000). As such, our research answers questions relevant to a spectrum of private and public contexts.

Prior Literature and Our Research Approach. In simple terms, a procurement auction is a purchasing process where a buyer evaluates and ranks offers made by potential suppliers to source a good or service. Sometimes, these auctions award contracts based on price only, such as the sealed-bid first price setting described by Olivares et al. (2012), where the lowest priced bidder won the right to supply school meals in Chile. In contrast to price-only auctions, this research investigates multi-attribute auctions, where the offers include price and any other attributes that impact the buyer's utility; a *scoring rule* can be implemented to weight these attributes. Beyond the number of attributes, procurement auctions can also be broadly classified as either open progressing or sealed bid auctions. Though most auction models traditionally consist of price offers subject to fixed quality characteristics, scoring rules allow for the more general case where the buyer can consider and evaluate offers differing in many other quality dimensions beyond price. These may include specific supplier attributes related to the design, completion time or quality preferences. As such, each offer or bid can be written as a list or vector, comprised of price and quality levels, presented by each supplier to the buyer. These are multi-dimensional bids, distinct from cases where the only decision that a supplier needs to offer is price. As Asker and Cantillon (2008, Theorem 6) showed, fixing the quality characteristics at a level, and assigning according to prices only is strictly worse for the buyer than using a score auction where the score reveals the buyer's preferences.

It is important to differentiate sealed bid auctions (the focus of this paper) from open progressing auctions. Briefly, in open progressing auctions, sellers compete to supply a good or service at a progressively better offer for the buyer which improves during the duration of the auction. Several software platforms (such as Ariba or FreeMarkets) were primarily designed to implement open progressing auctions, where each bidder is aware of the best competing bid during the auction, and may choose to improve their bid over time. At the end of the auction, the seller who has offered the best bid for the buyer is awarded the contract, and is expected to supply under the specified price and terms. In contrast, in sealed bid auctions, each supplier submits a unique final offer. All offers are simultaneously revealed, and the contract is assigned to the supplier with the offer that generates the highest net benefit (utility minus paid price) for the buyer. In contrast to open progressing auctions, sellers do not know the offers made by any competitor until they are revealed, and are prevented from adjusting their bid in response to other bids.

Prior research based in analytical modeling has used market design and game theory tools to analyze different hiring/purchasing mechanisms, providing useful theoretical insights when bids are either fully-transparent (e.g., Che 1993, Asker and Cantillon 2008), open-progressing (e.g, Chen-Ritzo et al. 2005), or single-dimensional (e.g., Kostamis et al. 2009). However, for a large number of the auctions conducted by ChileCompra (sealed-bid, first score, multi-dimensional), there is no equilibrium-in-pure-strategies solution when the assignment rule is non-transparent and feedback is restricted to bid levels only instead of actual scores. Figure 1 displays the difference between the two scenarios.

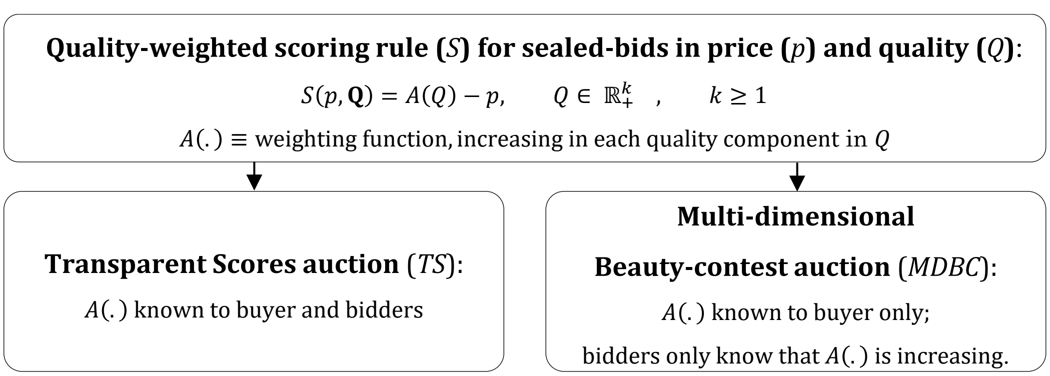

Figure 1: Sealed-bid mechanisms under comparison.

This research was developed to help answer the question of the cost of transparency faced by procurement agencies such as ChileCompra. Although a field experiment directly manipulating the information available would have been desirable from a scientific perspective, it is naturally problematic to conduct any empirical study that manipulates the trust or transparency environment: At ChileCompra, they maintain that a high commitment to the announced rules, information disclosure, and enforceability were all desirable market traits not to be sacrificed. Given these constraints, we decided to contrast both mechanisms through the use of an incentive-compatible controlled laboratory experiment. Using an environment that is analogous to the governmental procurement auctions used by agencies such as ChileCompra, we experimentally compare the bidding decisions of suppliers in sealed-bid auctions with multi-dimensional decisions with and without disclosure of the quality-weighted scoring rule as in Figure 1.

Our work is the first to analyze multi-dimensional procurement bidding decisions using sealed-bid first scores with independent private costs, contrasting transparent to concealed scoring rules in a laboratory setting. Existing experimental research using independent private costs has provided insights for related, but fundamentally different mechanisms and is summarized in Figure 2.

	Bidding mechanism			*Number of decisions per bidder*	
	Sealed-bid, 1st-score	**Sealed-bid, 2nd-score**	**Open progressing bids**	**Single dimensional bids (quality exogenously determined by Buyer)**	**Multi-dimensional quality + price bids (fully determined by Supplier)**
Chen-Ritzo et al. (2005)			√		√
Engelbrecht-Wiggans et al. (2007)	√	√		√	
Haruvy and Katok (2013)	√		√	√	
Fugger et al. (2015)			√	√	
Our work	√				√

Figure 2: Experimental research on Multi-dimensional procurement auctions

In contrast to sealed-bid first score auctions, Chen-Ritzo et al. (2005) studied multi-dimensional bidding in open score auctions with a scoring rule that could be inferred from the feedback generated by open-progressing bidding. In such an auction, bidders revise their offers based on their relative rank feedback as the bidding takes place – they can observe where their quality (speed of delivery) is too low, or their price is too high, compared to the current winning bid, and submit a revised counter-offer. In practice, the feedback provided gives a full display of the buyer's preference[2].

Fugger et al. (2015) studied multiattribute assignments in open-progressing procurement auctions with price-only bids; this transforms the quality dimension into a problem of bidding with an additive exogenous disturbance term. Engelbrecht-Wiggans et al. (2007) experimentally tested second-price vs. price-and-quality procurement auctions using sealed-bids. In their setting, the quality dimension was considered exogenous, random and observable for the bidder and the buyer (but not by any other competing bidder), making bid decisions effectively single-dimensional. Haruvy and Katok (2013) allowed for both sealed and open bidding and public concealment/disclosure of bids, but quality-levels remained exogenously determined by the buyer. A related empirical study by Moreno and Terwiesch (2014) showed that bidders who have access to externally-generated reputation ratings increase their prices as their reputations increase. However, the reputation scores are not controllable by the suppliers, who only bid on price in each auction. Therefore, in these articles, quality is not directly controlled by the bidder, and thus bidding decisions are price-only.

In a similar multidimensional bidding problem, Bichler (2000) studied an electronic brokerage system trading stocks and stock options with fully transparent assignment. Their system used a proportional scores mechanism, and their focus was on the bidding behavior in this brokerage context. More importantly, their study used different score valuations, making impossible to compare the effect of the different mechanisms for the buyer. An extension of this work, Bichler and Kalagnanam (2005), studied the related problem of multi-sourcing using multidimensional bidding.

A common characteristic that our work shares with Chen-Ritzo et al. (2005), Bichler (2000) and Bichler and Kalagnanam (2005), is that we allow for bidder-determined quality, as opposed to exogenous buyer-determined preferences for bidders. In our context, the ability to make non-price bidding decisions (such as acceleration in completion time, number of dedicated service employees, or any other quality dimension) is essential. These attributes are also contractible, visible to competing buyers, and to the public.

---

[2] They also compared their main multi-dimensional setting to the case where the quality attributes were fixed and bidding was on price-only, for benchmarking purposes. The multi-dimensional mechanism strongly outperformed the price-only case.

In order to isolate the effect of transparency loss, suppliers in our experiments were provided with perfect information about their own private cost structures – the *independent private values* paradigm. In practice, this paradigm precludes the possibility of suppliers falling prey to a "winner's curse", where they might win an auction for a lower value than the actual worth of the contract. Similarly, we assume perfect commitment: A bidder who wins an auction event will automatically supply the same price and quality levels offered, and the buyer will pay that amount for the good. In reality, risks of default and hold-up exist, and Bajari et al. (2014) provide ample evidence of those effects in governmental procurement. However, such complexities are not part of our research.

Understanding quality. In general, we refer to all the "B" components of an A+B bid as *quality*. For the purposes of this research, quality consists of any non-price attribute dimension present in a product offered by a supplier that affects the utility of the buyer. These attributes, in our research, are assumed verifiable, measurable and contractible/enforceable by the buyer, and can thus be used to rank bids and assign contracts. To understand the notion of quality that can be incorporated into an auction by a buyer, we refer to the taxonomy offered by Ketokivi and Schroeder (2004). They identify three possible types of measures that can be used to quantify quality attributes: operationally-defined measures, perceptual measures, and quasi-perceptual measures.

Operationally-defined measures are quality metrics which are objectively quantifiable; in procurement auctions, common examples include the acceleration time to fulfill a contract (compared to a maximal acceptable delivery time), the specific quantity of materials used, the hours of operation of an office, or the number of employees assigned to staff a point of service during peak time. These and other measures are objective, reliable and verifiable by both bidders and buyers, making their contractibility and enforceability possible. These are the least problematic for the buyer and suppliers to communicate and assess. However, buyers often consider other attributes that generate positive utility that include perceptual components in their measure.

Perceptual and *Quasi-perceptual measures* include any metrics that are subjectively quantified by the buyer. In general, the consideration of quality attributes with perceptual components in governmental procurement auctions is problematic due to the impression of impropriety when transparency is an explicit goal. While purely perceptual measures are based on impressions and experiences, quasi-perceptual measures also often have an embedded and underlying operational component. For example, governmental agencies sourcing via ChileCompra often would like to include the perceptual quality attribute called "bidder experience", where inclusion of such an attribute reduces scores for inexperienced or non-incumbent suppliers. These attributes are difficult to objectively measure: Usually when a bidder is evaluated, "experience" is expressed as a relative measure of the *perceived disutility* generated by sourcing from a non-incumbent bidder.

The presence of a perceptual or quasi-perceptual measure also impacts the modeling of optimal outcomes. Specifically, when a perceptual measure in a sealed-bid multi-dimensional bidding auction does not provide a universally understandable metric for the attribute it measures, the use of perceptual or quasi-perceptual measures is equivalent to the sealed-bid multi-dimensional beauty-contest auction described here. As an example, consider sealed bids (p, q) on price and quality, and set a scoring rule to rank bids: $S(p, q) = a\,q - p$, with $a > 0$ known only to the buyer and unknown to the bidders. All that a bidder knows in this case is that S is increasing in q and decreasing in p. If the scoring rule was instead $S(p, q) = a\,w(q) - p$, with $a > 0$ common knowledge to buyer and all suppliers, but with $w(.)$ unknown to the supplier except for $w'(.) > 0$, then both scoring rules are analogous in terms of their lack of transparency. The former case is what we term a sealed-bid multi-dimensional beauty-contest, whereas the latter case is an instance of a quality attribute quantified as a perceptual measure, even if the scoring weight for quality is known. As we elaborate in the next section, when we introduce our scenarios, in some instances is impossible for a buyer not to use perceptual measures, particularly when the type of product or service being procured is destined to satisfy a new need or necessity.

3. Environment

In this section, we formalize a unified framework in sealed-bid multi-dimensional score auction models that is both general and tractable. This framework is analogous to similar elements found in prior literature (e.g., Che 1993, Zheng 2000, Asker and Cantillon 2008, Lewis and Bajari 2011) and that we have tied to practice. Our notation is summarized in Figure 3.

The specifics of the modeling environment are crucial in establishing an appropriate benchmark for comparison with the empirical results. In the present section, we provide structure to the general bidding problem which makes it possible to draw specific conclusions. For example, the specification allows for as many quality dimensions as desired. We define a cost function that yields a log-linear supply for quality while allowing for many possible values for the elasticity of supply for quality; this specification is more general than Lewis and Bajari (2011). Therefore, our environment embeds many possible specifications of an admittedly complex decision environment. Any empirically observed deviations from optimality in the simplest instance of the model (see §4) will be amplified as the environment becomes more complex.

TS:	Transparent scores auction scenario
MDBC:	Multi-dimensional beauty contest auction scenario
EPT:	Multi-dimensional beauty contest auction scenario with ex-post transparency
N :	Number of bidders, indexed by j
Π_j, Π_B :	Utility (Profit) for bidder j and for the buyer, respectively
p :	Price paid per auctioned unit
$\mathbf{Q} \equiv (q_i)_{i=1}^{k}{}'$:	Quality stock for each non-monetary attribute i in auctioned unit
$\mathbf{A}$:	Non-negative weights given to each quality attribute in $\mathbf{Q}$. Common knowledge in the TS scenario, private information to the buyer in the case of the MDBC scenario.
$V(\mathbf{Q})$:	Valuation of a good with quality $\mathbf{Q}$ for the buyer, assumed as linear: $\mathbf{A'Q}$
$\mathbf{Z}_j \equiv \left(z_{ij}\right)_{i=1}^{k}{}'$:	Realized production efficiency (private signal) for bidder j for attribute i.
$1/\alpha$:	Elasticity of supply for quality (standardized at 1)
F:	Fixed production cost, only incurred by winning bidder
$C(\mathbf{Z}_j, \mathbf{Q})$:	Total production cost for winning bidder
$\mathbf{Q}^*(\mathbf{Z}_j)$:	Social first best level of quality achievable given $\mathbf{Z}_j$, $\mathrm{argmax}_{\mathbf{Q}}\{V(\mathbf{Q}) - C(\mathbf{Z}_j, \mathbf{Q})\}$
$S(p, \mathbf{Q})$:	Scoring auction rule, $s = \mathbf{A'Q} - p$. Also, since $V(\mathbf{Q}) = \mathbf{A'Q}$, the rule denotes the Buyer's Surplus when evaluated at the value of the winning auction.
β_i^L, β_i^U:	Bounds of the distribution of production efficiencies
$\mathbf{A'Q}$ - $C(\mathbf{Z}_j, \mathbf{Q})$:	Social welfare generated from quality $\mathbf{Q}$.
$v_j \equiv v(\mathbf{Z}_j)$**:**	Pseudo-type, the maximum social welfare a bidder of type $\mathbf{Z}_j$ can generate
$X_v(v_j)$:	Distribution of pseudo-types, with bounded support $[v_L, v_U]$. (TS only)
$b(v_j)$:	Expected profit maximizing bidding function (in scores) (TS only)

Figure 3: Notation Used in Model Formulation

Consider a buyer that seeks to procure an indivisible good (or service) for which there are N potential suppliers or bidders. Bidders, indexed by j, know the number of competitors they face, $N - 1$, before placing a bid. The good is characterized by its price, p, and quality attributes $\mathbf{Q} \equiv (q_i)_{i=1}^{k}{}'$ that can be contractually enforced. The buyer values the good according to quasi-linear preferences:

$$\Pi_B(p, \mathbf{Q}) = V(\mathbf{Q}) - p = \Sigma_{i=1}^{k} a_i q_i - p \tag{1}$$

where $a_i > 0$ are attribute-related constants. Hence, the buyer has increasing preferences in each quality dimension. Bidder *j*'s profit Π_j from selling good $(p, \mathbf{Q})$ consists in the quoted price p minus the production cost C if the contract is won:

$$\Pi_j(\mathbf{Z}_j, p, \mathbf{Q}) = p - C(\mathbf{Z}_j, \mathbf{Q}) \tag{2}$$

where $\mathbf{Z}_j = (z_{ij})_{i=1}^{k}{}'$ is a vector of random variables representing bidder *j*'s *production efficiency types*, with realizations known only to the bidder. We assume these efficiency types are defined on a strictly positive closed and bounded support defined by (3).

$$z_{ij} \in [\beta_i^L, \beta_i^U], \quad 0 < \beta_i^L < \beta_i^U \tag{3}$$

In order to model the cost function $C(\mathbf{Z}_j, \mathbf{Q})$, we follow an empirical strategy similar to Lewis and Bajari (2011) to model supplier's production costs in A+B highway construction auctions in California. The quality dimension considered in their study was the acceleration time to fulfill a contracted construction or repair work. Their single-dimensional quality analysis is generalizable to $k \geq 1$ quality dimensions. We structured supplier cost C as the sum of a quality-independent baseline cost and a quality-dependent cost structure:

$$C(\mathbf{Z}_j, \mathbf{Q}) = F + \sum_{i=1}^{k} \frac{q_i^{1+\alpha}}{(1+\alpha)z_{ij}} \tag{4}$$

We assumed the baseline cost parameter $F \geq 0$ to be a deterministic, non-sunk, fixed cost incurred by the bidder only if the auction is won and the good is sold; it is identical and known by all bidders. This baseline cost can be interpreted as the cost of completing a project at the minimal acceptable quality level. For example, the quality dimension in Lewis and Bajari (2011) is acceleration time (completion time savings); the baseline cost is interpreted as completing the project on the design engineer's original schedule. For the experiment, the baseline fixed cost is common to all bidders to avoid ex-ante asymmetry, as favoritism or skill in estimating production costs is not the focus of this study.

As for the quality-dependent portion of the cost structure, this specification yields a marginal cost (i.e., supply) function for quality dimension i given by $C_{q_i}(\mathbf{Z}_j, \mathbf{Q}) = q_i^{\alpha}/z_{ij}$, which implies a constant elasticity of supply for quality, $1/\alpha$. This flexible specification allows for a convex ($\alpha > 1$), concave ($\alpha < 1$) or linear ($\alpha = 1$) marginal cost for quality.

Social welfare, the sum of buyer and suppliers benefits, becomes $V(\mathbf{Q}) - C(\mathbf{Z}_j, \mathbf{Q})$. For simplicity, we have assumed costs to be independent across attributes, $V - C$ to be bounded, and $(V_{\mathbf{QQ}} - C_{\mathbf{QQ}})$ to be negative semi-definite. These mild regularity assumptions, based on existing literature (see Asker and Cantillon 2008), and satisfied by the experimental design presented in the next section, guarantee that social welfare $V - C$ is bounded and strictly concave in $\mathbf{Q}$, and, hence, $\mathbf{Q}^*(\mathbf{Z}_j) = \mathrm{argmax}_{\mathbf{Q}}\{V(\mathbf{Q}) - C(\mathbf{Z}_j, \mathbf{Q})\}$, the social first best level of quality attributes achievable by each j, is well defined and unique. Bidders' preferences are common knowledge, but bidders' types $\mathbf{Z}_j$ are private information: each bidder knows the realization of their own efficiency type, and knows the distributions from which each of their competitors' types are (independently) drawn. This is an important assumption: If bidders know their own cost realization, it is rationally impossible for any of them to bid below cost, precluding the possibility of the phenomenon known as *winner's curse* (e.g., Bajari and Hortaçsu 2003). We explicitly preclude that possibility with this independent private values (costs) structure, to minimize additional level complexity in the bidders' decisions.

In this context, we analyze two sealed-bid scenarios. The TS scenario is a *transparent scores* auction, where the assignment criterion is communicated to all suppliers before the start of the sealed-bid first score auction. The MDBC (*Multi-dimensional beauty contest*) scenario is identical to TS, except that suppliers do not know the weights given by the buyer to each attribute before submitting their bids. As the buyer weights are identical in both cases, we can compare performance between the two scenarios to evaluate the impact of transparency on performance and the specific benefits to the buyer.

Sealed-bid Transparent Scores Auction (TS). A score auction is a multiattribute auction where bids $(p, \mathbf{Q})$ are evaluated according to a pre-specified scoring rule, $S(p, \mathbf{Q})$, to which the buyer commits (and communicates) before the auction takes place. Che (1993) introduced this mechanism and first characterized its equilibrium solution. Asker and Cantillon (2008) offered a didactic and general characterization of its optimal bids and outcomes.

We assume the scoring rule in the auction to be a linear function, i.e.: $S(p, \mathbf{Q}) \equiv A(\mathbf{Q}) - p = \mathbf{A}'\mathbf{Q} - p$, which is consistent with many real-world implementations of multiattribute auctions (e.g., Lewis and Bajari 2011; Bajari et al. 2014; Snir and Gupta 2011). In principle, it is not clear whether linear rules are used as a translation of the buyer's preferences, because of their simplicity, or for other considerations. While some studies (e.g., Kersten 2014) have questioned the effectiveness of monotone scoring rules, we assume that the quality portion of the rule **is** a measure of the buyer's preferences for quality: For every possible value of $\mathbf{Q}$, $V(\mathbf{Q}) \equiv \mathbf{A}'\mathbf{Q}$. This is a reasonable assumption in practice; for example, the data analyzed by Lewis and Bajari (2011) comprised many A+B bidding auctions that considered price and completion time; the weights given to completion times were calculated by expert engineers on behalf of the California Department of Transportation, and measured the daily social cost to the public for the closure of the specific infrastructure (a bridge or street) in question. This preference structure also satisfies the assumptions on boundedness and strict concavity of $V - C$ discussed above. Thus, while ensuring participation of all bidders in equilibrium, $\mathbf{A}'\mathbf{Q} - C(\mathbf{Z}, \mathbf{Q})$ is effectively bounded, strictly concave in $\mathbf{Q}$ for all $\mathbf{Z}$, and such that $\max_{\mathbf{Q}}\{\mathbf{A}'\mathbf{Q} - C(\mathbf{Z}, \mathbf{Q})\} \geq 0$ for all $\mathbf{Z}$.

The *outcome* of the score auction is a probability of winning, x_j, and a score to fulfill when the bidder wins the contract, s_j. A winning bidder of type $\mathbf{Z}_j$ with score s_j will choose $(p, \mathbf{Q})$ to maximize profit $p - C(\mathbf{Z}_j, \mathbf{Q})$ subject to $s_j = \mathbf{A}'\mathbf{Q} - p$. Substituting the constraint in the objective function yields:

$$\max_{\mathbf{Q}}\left\{\mathbf{A}'\mathbf{Q} - C\left(\mathbf{Z}_j, \mathbf{Q}\right) - s_j\right\} \tag{5}$$

From (5), the optimal $\mathbf{Q}$ does **not** depend on s_j. Following Asker and Cantillon (2008) and others that have studied score auctions, we define a value function called *pseudo-type*, $v(\mathbf{Z}_j)$, the maximum social welfare a bidder of type $\mathbf{Z}_j$ can generate given the scoring rule. For instance, if supply elasticity $1/\alpha = 1$, we obtain a pseudo-type given by (6):

$$v\left(\mathbf{Z}_j\right) \equiv \max_{\mathbf{Q}}\left\{\mathbf{A}'\mathbf{Q} - C\left(\mathbf{Z}_j, \mathbf{Q}\right)\right\} = \max_{\mathbf{Q}}\left\{\sum_{i=1}^{k}\left(a_i q_i - \frac{q_i^2}{2z_{ij}}\right) - F\right\} \tag{6}$$

The quality choice that maximizes (5) is the same that maximizes (6). Solving for the optimal $\mathbf{Q}$ yields an optimal quality-dimension bid and a corresponding pseudo-type.

$$q_i^* = a_i z_{ij} \tag{7}$$

$$v(\mathbf{Z}_j) = \frac{1}{2}\left(\sum_{i=1}^{k} a_i^2 z_{ij}\right) - F \tag{8}$$

Because (8) is a function of the private efficiency types, the pseudo-type is an ex-ante random variable for each bidder. We call the distribution of this random variable $X_j \equiv X_v(v_j)$, with density function $x_v(v_j)$. Expected profit for bidder j is given by:

$$\mathrm{E}\Pi_j = \Pr(winning) \cdot \left(v(\mathbf{Z}_j) - s_j\right) \tag{9}$$

The above expression is mathematically equivalent to a sealed-bid first price auction with pseudo-types functioning as private valuations, and scores taking the role of prices. Therefore, bidders' preferences over outcomes (x_j, s_j) are fully captured by $v(\mathbf{Z}_j)$. As Asker and Cantillon (2008) show, pseudo-types are *sufficient statistics* for this auction mechanism: Once the distribution of pseudo-types is obtained, the result is a complete characterization of the score auction and of its outcome.

Given our assumption on the support of efficiency types in expression (3) and the optimal quality bidding described above, the distribution of pseudo-types is defined in the support given by expression (10):

$$v_j \equiv v(\mathbf{Z}_j) \sim X_v(v_j) \in \left[v_L \equiv \frac{1}{2}\left(\sum_{i=1}^{k} a_i^2 \beta_i^L\right) - F,\ \ v_U \equiv \frac{1}{2}\left(\sum_{i=1}^{k} a_i^2 \beta_i^U\right) - F\right] \tag{10}$$

Next, we solve for the optimal price portion of the bid. Let $b(.)$ be a monotone bidding function mapping pseudo-types $v_j = v(\mathbf{Z}_j)$ onto submitted scores $s_j = S(p_j, \mathbf{Q}_j)$ in equilibrium. Generalizing Riley and Samuelson (1981), the symmetric Bayesian-Nash equilibrium (SBNE) score bid $b(v_j)$ will satisfy differential equation (11) at optimality:

$$b'(v_j) = (v_j - b(v_j))(N-1)\frac{x_v(v_j)}{X_v(v_j)} \tag{11}$$

This differential equation completely characterizes the optimal score bid strategy function, $b(v_j)$. We note that this strategy will depend on the resulting distribution of pseudo-types induced by the bidders' efficiency types, as well as the number of bidders.

In the specific instance of uniformly distributed efficiencies, the distribution of pseudo-types is also uniformly distributed, with bounds given by the values in expression (10). That gives a simple closed form solution to expression (11):

$$b(v_j) = v_j \cdot \frac{(N-1)}{N} + v_L \cdot \frac{1}{N} = v_L + (v_j - v_L) \cdot \frac{(N-1)}{N} \tag{12}$$

We note that this allows consideration of the multi-dimensional decision setting that describes the procurement environment of interest to agencies such as ChileCompra.[3] We have specified a particular case of the sealed-bid transparent scores model that incorporates a cost structure for quality analogous to the total quality cost structure described by Lewis and Bajari (2011). This permits us to identify optimal bids not only at the pseudo-type and score level (given by (11)), but also by price and individual quality attributes (given by (7))

Naturally, solving for this differential equation is complex, even in the case where the closed-form analytical solution is relatively straightforward to compute (one such instance is experimentally tested as shown in §5), under the assumption that bidders know their own cost structure (i.e., under the private values paradigm). In particular, we expect bidding decisions to differ from the SBNE score bid defined above by chosen quality, price, and optimal scores. For taxonomy purposes: we refer to deviations consistent with more aggressive bids (smaller bidder profit margins as compared to the SBNE predictions) as *overbidding*, and deviations that translate into less aggressive bids as *underbidding*. While most existing experimental studies of sealed-bid, first price auctions have found systematic evidence of overbidding, we observed that in the experiments conducted by Bichler (2000), submitted scores in sealed-bid, first score auction tended to show, counterintuitively, evidence of *under*bidding.

Multi-dimensional Beauty-contest Auction (MDBC). For this scenario, consider the case where the scoring rule is hidden. In practice, we implement the same rule as the

[3] Specifically, this allows the supplier to bid on price and quality dimensions in sealed-bid procurement auctions. As in the literature described in §2, most prior experimental literature focused on either price-only, or price and pre-determined quality where the supplier's quality was fully exogenous (e.g., Engelbrecht-Wiggans et al. 2007).

preceding sealed-bid score auction scenario, where bidders submit sealed bids for prices and quality levels, and the winner is determined according to this rule. The only difference is that the attribute weights a_i used to implement the scoring rule were not disclosed to the bidders at the moment they submit their bids.

As discussed in §2, concealing the scoring rule is mathematically equivalent to not communicating the metric upon which the quality characteristics will be scaled or measured – effectively a (quasi-)perceptual measure of performance in the language of Ketokivi and Schroeder (2004). In this case, bidders still make bidding decisions on price as well as their controllable input to the buyer's perception of quality. However, even in cases when they are provided the direction given by the scoring rule, they are not informed about *how* the magnitude of increments in the latter dimension(s) affect their scores, and their subsequent likelihood of winning.

For example, consider a highway department requesting bids for a new bridge where contracts will be awarded based on a scoring rule that gives an explicit weight, say 15, to time acceleration in days with respect to a pre-determined deadline, plus a weight of 80 to an item called *local presence*, minus the bid price. Local presence usually refers to initiatives that generate substantial employment in a specific geographical zone. The bidder can offer a specific time acceleration component and a price, but lack a clear metric on how that "local presence" is quantified or evaluated. A bidder could commit to hiring a high percentage of local workers, start a local office, or source from local suppliers or subcontractors. Each is broadly consistent with local presence and the bidder knows that increasing the dimension should increase their score. However, the lack of a clear metric does not allow the bidder to control the impact of their offer on their bid. It becomes as nontransparent as announcing that the scoring rule will consider the number of local workers hired as a metric of local presence (a concrete measure) without announcing the exact weight of 80 on that metric.

It is important to note that in some instances, it is impossible for a buyer to transparently communicate either the exact weight or the metric for quality. In casual conversation, this is consistent with the colloquial expression "*I know quality when I see it*." Chen-Ritzo et al. (2005) wrote: "*An issue facing the real-world auction designer is that the buyer may not know his utility function. Instead, he must learn about what he wants by observing the different options available*." (p. 1762). Consider the following illustrative example of a supplier of adhesive bandages who bid to supply products via ChileCompra: A hospital wanted to procure a new type of surgical adhesive bandage. One key quality dimension considered was the degree of adhesiveness of the bandages. The metric involved a "just right" amount of adhesion, a value that was difficult to quantify in advance for both buyer and bidders. If the bandage was "not sticky enough", it would fall off and potentially fail to protect the wound, while if, on the contrary, it was "too sticky", it would be hard and painful to remove. Unfortunately, the evaluation of the "just right" was difficult, so bidders had to include samples of the type of bandage they offered in addition to the price charged. Only *after* the different samples were tested, the institutional buyer was able rank the submitted offers according to their adhesion quality, from higher to lower, and then accept the one with the best score (as a weighted combination of price and distance from the ex-post optimal adhesion quality). In this case, the rule could be revealed a posteriori, and made transparent after the fact, but the initial bidders had to submit offers "in the dark." Moreover, objectively, there are no guarantees that the evaluation of "just right" from this specific instance of this buyer's revealed preference would be repeated in the future.

We use the name *multi-dimensional beauty-contest* for this mechanism, because suppliers choose for their bids the levels of quality and price which maximize their expected utility given their personal estimate on the buyer's subjective preference for quality. If a buyer is a governmental agency whose interest is also to maximize benefits at a society level (making society's preferences their own), in a sense, suppliers are trying to estimate on those social benefit terms at large. The name "beauty-contest" (see Asker and Cantillon 2008) is inspired by a comment made by Lord Keynes (Keynes 1936) regarding beauty pageants: Any judge whose performance is evaluated on the quality of his/her own choice for the contest, should choose a winner not based on his or her own beauty preferences, but on what his or her own *perception* of what the general public considers the most desirable beauty features. Unlike beauty-contests in a game theoretical sense (i.e., a "majority rule", where decision makers try to guess as close as possible to x% of the average decision maker's guess; see Nagel 1995), for the mechanism described here, decisions of suppliers do not have power to influence the value of the buyer's preference for quality. This makes the belief formation process for a supplier impossible to determine without additional simplifying assumptions.

Starting with Chen-Ritzo et al. (2005), existing studies have derived bidding strategies for open-progressing MDBC-like mechanism environments, where the buyer's preference for quality is implicitly revealed via the relative ranking displayed through the auction's progress. The open-progressing environment is the key driver of their analytical results.

In contrast, for the sealed-bid case in question, the equilibrium bid is undetermined without assuming a *prior* (belief structure) on the buyer's concealed preferences, or without reducing the dimensionality of the problem. For instance, similar sealed-bid multiattribute mechanisms were considered by Kostamis et al. (2006, 2009). In their research, quality is not a decision variable for suppliers, but the buyer's private information about a supplier's desirability. That model is a tool of mechanism selection for pre-determined supplier discrimination; in our case each supplier can still affect their own desirability to the buyer through their chosen quality level, i.e., for our case, quality is endogenously determined.

In MDBC, bidders only learn whether they won or not at the end of the auction; though they are informed of the winning price and quality levels, they are not informed about the actual score associated with the winning bid. The difference in mechanism performance between MDBC and TS is our metric for the cost of transparency loss.

Hypothesizing on the preference of one scenario over the other. While the preceding auction theoretical intuition would appear to offer an advantage to TS over MDBC, there is no empirical evidence comparing the performance of the two, and the behavioral aspects of multi-dimensional bidding have not been extensively studied. In particular, the evidence of underbidding (i.e., suppliers bidding less aggressively than the SBNE suggests) found in the previous study of TS-types of auctions by Bichler (2000), combined with the lack of experimental studies devoted to MDBC, make it impossible to say whether MDBC or TS will perform better for the buyer or for the sellers. Would the sealed-bid, first score format make bidders shrink winning bids towards zero profit in the MDBC? For TS, would we find evidence of overbidding as in the traditional price-only world of sealed-bid, first price auctions, or would we instead see less aggressive bidding as found by Bichler (2000)? All these aspects make the answer to our research question of which of the two scenarios, between TS and MDBC, is preferable for the buyer, completely non-trivial.

4. Study 1: Comparing TS vs. MDBC

Experimental design. We designed study 1to compare both scenarios described above, TS vs. MDBC, in a way that gives the theoretical benchmark from the SBNE strategy under TS the best shot to work, by simplifying the participants' decision tasks while still capturing the complexity of the problem. We implemented two treatments, each representing one of the two auction scenarios. The streams of parameters and random production efficiency factors were the same for all treatments (although the efficiency factors were re-drawn after every period, every auction had the same N draws within a given round). Half of the participants were assigned to each treatment, and we used a between-subjects design. Subjects were matched in groups of $N = 4$ bidders for 40 rounds. Subjects did not know the identity of the bidder against whom they were competing and were rematched in every round to different competitors in the room. Subjects were informed at the beginning of the session of this periodic anonymous rematching. In each session there were 24 participants, all sitting in the same room as the rest, with each participant sitting in front of an individually isolated computer station for the length of the study.

Random rematching is a commonly used technique in experimental auctions research (e.g., Katok and Roth 2004, Ockenfels and Selten 2005, Wan et al. 2012, Chen et al. 2013). Briefly, though repeated social interaction would generate a need to study the bidding problem as a dynamically repeated game, the anonymous rematching in every round with other subjects in a relatively large room gives credence to the appropriateness of using a non-dynamic framework to analyze the auctions. Specifically, subjects had no way of knowing with whom they were matched, thus making it unlikely they could consider the behavior of specific competitors. Subjects were not allowed to participate in more than one session or more than one treatment.

Experimental parameterization. As noted, submitting a multi-dimensional bid is a complex task, even when the scoring rule is known. In our experiments we consider bi-dimensional bids, in price p and one quality dimension q, meaning that all i-indices for quality can be dropped. In sealed-bid environments, the simultaneity of the bidding precludes feedback on learning both the relative preference for alternative non-winning offers as well as the appropriateness of a bidders' decision, making the bidding process psychologically more complex than with open bids. Therefore, we restricted our design to a single non-monetary dimension.

We also consider the efficiency type z_j to be uniformly distributed between $\beta^L = 10$ and $\beta^U = 90$, the elasticity of supply for quality $1/\alpha = 1$,[4] and the baseline cost F to be

[4] While Lewis and Bajari (2011), in their own study of A+B auctions, estimated $1/\alpha$ to be around 0.3 for California highway contractors, different markets will have different elasticities for quality. In

equal to 500. We also fixed the parameter measuring preference for quality to $a = 10$. These parameters provide bidders a simple numerical setting that still captured real-world features of the problem.

For the TS auction, with this parameterization, the bid constituting a SBNE in pure strategies (in terms of quality q_j^* and price p_j^*) reduces to expression (13) below:

$$\begin{gathered} \boldsymbol{q_j^*} = az_j = \mathbf{10}\boldsymbol{z_j} \\ v_j = \frac{1}{2}a^2 z_j - F = 50z_j - 500 \sim U[v_L = 0, v_U = 4000] \\ \boldsymbol{s_j^*} = b(v_j) = v_j \frac{(N-1)}{N} + v_L \frac{1}{N} = \mathbf{37.5}\boldsymbol{z_j} - \mathbf{375} \\ \Rightarrow \ \boldsymbol{p_j^*} = aq_j^* - s_j^* = \mathbf{62.5}\boldsymbol{z_j} + \mathbf{375} \end{gathered} \tag{13}$$

Here, (13) is a relatively simple closed-form expression. In terms of scores, optimal bidding suggests that the SBNE scoring bid for the TS case with these parameters should be $(N-1)/N = 75\%$ of the bidder's pseudo-type v_j.

For the MDBC auction, an expression analogous to (13) is not well defined, because bidders are not informed of the value of the quality weight parameter a, which is private information of the buyer. The a parameter is both a measure of the true preference of the buyer with respect to quality and the weight used by the buyer in determining the winner of the auction, regardless of whether this value is communicated (TS auction) or not (MDBC auction). As such, the winning score bid is also the buyer's surplus in both treatments. For welfare comparisons, we maintained parallelism by keeping everything the same in both auctions. The only difference was revealing or concealing the numerical value of the quality weight parameter a.

Protocol. Each subject was compensated in proportion to their total earnings for all 40 rounds, plus a \$5 participation fee. $J_{TS} = 24$ subjects participated in the transparent scores (TS) treatment, and other $J_{MDBC} = 24$ participated in the multi-dimensional beauty-contest

---

the product quality literature, total cost has been usually assumed quadratic in quality (e.g., Karmarkar and Pitbladdo 1997, Karaer et al. 2015), and is also consistent with the multidimensional bidding model presented by Zheng (2000). This implies an elasticity of supply for quality of $1/\alpha = 1$, which also provides computational simplicity without sacrificing the real world characteristics of the problem.

(MDBC) treatment. The average pay was \$13.48, and each subject only participated in one treatment session. Payment took place in private at the end of the session; cash was the only incentive offered. All sessions were conducted at a major U.S. research university, using the subject pool associated with the business school. Subjects read the instructions (see Appendix A). After this, the experimenter read the instructions out loud and presented screen captures as examples. Approximately 10 minutes were used for reading instructions and answering questions, and 80 minutes were used for the decision task.

We programmed the experimental interface using the *z*Tree system (Fischbacher 2007). The bidding screen (see Appendix A) included the realization of efficiency type z_j as well as the parameters for fixed cost F = 500 and weight α = 10 (the weight was only displayed during the TS treatment). It also included a table that showed all relevant information from past periods (own offered price and quality, winning price and quality, profits, etc.). The cost function and scoring function equations were provided in the instructions and the bidding screen. Because the decision task is computationally challenging, we also provided subjects with a decision tool that automatically calculated costs and earnings using the efficiency type and their inputs of price and quality. Bidders were precluded from submitting bids that would generate a financial loss, as is common in independent private value auction experimental designs (e.g., Ockenfels and Selten 2005).

After all four bidders in an auction submitted their offers, each saw a results screen showing if they had won or not, their own offered price and quality, the auction's winning price and quality, their profits, and a table with their past history of all rounds. Bidders did not receive information on the other auctions in the room, avoiding to introduce noise in subsequent bids. The only information displayed every period was the actual bid placed by the winner in case the winner was not the same subject, which is analogue to one of the treatments (*NF*) in Ockenfels and Selten (2005) and in Katuščák et al. (2015). Subjects did not know the identity of their competitors in any round, and were informed they were rematched every round with different subjects.

5. Results for Study 1

In this section, we measure differences in both bidding decisions and mechanism performance. In §5.1, we provide an overview of *bidding behavior* and compare the outcomes of the TS and the MDBC auctions. In §5.2, we focus on *auction (winning)* *outcomes*, and compare the performance of the two auctions. §5.2 provides the key comparison of interest for the impact of transparency v/s non-transparency for the buyer.

5.1. Bidding decisions

Individual bidding behavior results for the TS and MDBC treatments are summarized in Table 1 and Table 2, looking at the performance of all bidders. Complete summary statistics for each subject are included in Tables B1 and B2 in Appendix B.

Table 1: Comparison of actual average bidder behavior and outcome (TS vs. MDBC)
(*Unit of analysis: Individual bidder, averaged over 40 auctions, $J_{TS} = J_{MDBC} = 24$ bidders*)

		Quality $\bar{q}_J$	**Price** $\bar{p}_J$	**Score** $\bar{s}_J$	**Profit** $\overline{\Pi}_J$
Actual TS	Mean	478.9	4416.3	372.9	173.4
	SD	133.5	988.2	841.2	114.2
Actual MDBC	Mean	509.1	4915.5	175.1	230.0
	SD	113.1	1151.2	751.9	96.5
Mann-Whitney tests	z_{M-W}	-0.753	**-2.062**	1.278	**-2.165**
	p-value	0.452	**0.039**	0.20	**0.030**

We note that the average quality levels ($z_{M-W} = -0.753$, $p = 45.17\%$) and equivalent scores ($z_{M-W} = 1.278$, $p = 20.11\%$) are not statistically different in the two treatments (Mann-Whitney unmatched samples test[5]). However, average prices ($z_{M-W} = -2.062$, $p = 3.92\%$) and profits ($z_{M-W} = -2.165$, $p = 3.04\%$) are statistically higher in the MDBC

[5] The nonparametric tests assume that observations are either individually independent (Mann-Whitney) or pairwise-independent (Wilcoxon). Though not strictly independent (subjects were grouped by session), subjects were randomly and anonymously rematched each round. Our use of average decisions and outcomes over all rounds as the unit of analysis accurately reflects the decisions of bidders. From a behavioral perspective, individuals were likely to make decisions in their best interest regardless of with whom they were matched. It was impossible for them to determine the identity of any specific competitor before submitting a bid. Several ways to address possible session effects have been offered in the existing experimental literature, but as Frechette (2012) remarks, all introduce additional problems that do not necessarily address the aim of the research question. Thus, our design and analysis is reasonable and consistent with prior literature (e.g., Ockenfels and Selten 2005, Schmidt et al. 2003, Katok and Roth 2004, Chen et al. 2013).

condition. In general, the higher average profits per bidder were driven by the higher priced bids observed in the MDBC condition.

Table 2: Comparison of bidder behavior and outcome (actual vs. predicted), TS treatment (*Unit of analysis: Individual bidder, averaged over 40 auctions, J_{TS} = 24 bidders*)

		Quality $\bar{q}_J$	Price $\bar{p}_J$	Score $\bar{s}_J$	Profit $\overline{\Pi}_J$
Actual TS	Mean	478.9	4416.3	372.9	173.4
	SD	133.5	988.2	841.2	114.2
Predicted TS	Mean	497.2	3482.6	1489.6	193.9
	SD	28.1	175.4	105.2	53.0
Wilcoxon tests	z_W	0.171	**3.629**	**-4.286**	-0.857
	p-value	0.864	**<.001**	**<.001**	0.391

In TS (Table 2), unsurprisingly we observe deviations from the predicted SBNE outcomes. Observed average bidder profit is lower than predicted, but not significantly so (Wilcoxon signed-rank test, z_W = 0.857, p = 39.14%). Underbidding is observed, and statistically significant in terms of lower average scores (z_W = -4.286, p < 0.01%), and higher average prices (z_W = 3.629, p = 0.03%). Lower-than-predicted average quality bids were observed as well, but they were statistically insignificant (z_W = 0.171, p = 86.39%). This is an interesting finding, which is in contrast to overbidding behavior commonly observed in sealed bid, first price auctions when bidding is single-dimensional. This provides further evidence of the underbidding behavior observed in previous research on multi-dimensional bidding (e.g., Bichler 2000, Chen-Ritzo et al. 2005).

These results indicate a lower level of competitiveness than the SBNE predicted for the TS scenario. This outcome is similar to the implicit collusion outcomes observed in price-and-quality open-progressing auctions. In Fugger et al. (2015), the buyer had a private random signal for quality for each supplier, unknown to all suppliers (even the bidder did not know his/her own quality), resulting in a steep increase in prices given the uncertainty in the quality component assigned by the buyer. However, in their case the result was observed through pegging behavior to the reserve price, while in our case (in absence of reserve prices or uncertainty in quality bids) the higher prices appear as driven by the complexity of the task.

Note that Tables 1 and 2 include a summary of the bidding behavior across all bidders, winners and losers. In order to analyze the effect for the buyer, we need to consider the results by auction, considering the behavior of winners only. That is the analysis we present in the following subsection.

5.2. Mechanism performance and welfare

Next, we consider our main comparison: the mechanism performance due to transparency by treatment. Since we kept the same realizations of types fixed by auction (round), we average the outcomes of each of our six observations per auction for each treatment. This makes every auction round comparable to each other round conducted in parallel. Those averages are reported in Table 3.

Table 3: Comparison of auction performance per treatment (winning bids only).

	Buyer surplus (Winning score)	*Bidder outcome (Winner profit)*
Predicted Theory (*TS*)	2326.48	775.49
TS	1842.65	693.74
MDBC	1614.65	919.90
Wilcoxon (*Theory-TS*)	**5.309**	**2.379**
	(< .0001)	**(.0174)**
Wilcoxon (*TS-MDBC*)	**3.844**	**-4.059**
	(.0001)	**(< .0001)**

Note: *Unit of analysis: T = 40 rounds, each averaged over all parallel auctions per treatment. Values reported are averages per auction of each performance metric. Test p-values in parentheses.*

Using each auction round as unit of analysis, we observed 58.75% of efficient assignments in the TS treatment and 60.83% in the MDBC treatment. This means that in approximately 2 out of 5 auctions, we observed winning bidders who should not have won that instance.

Buyer's surplus in the TS auctions was significantly larger (Wilcoxon signed-rank test $z_W = 3.844, p = 0.01\%$) than in the MDBC auctions. Winners, on the other hand, made a substantially higher profit in the MDBC auction compared to the TS case (Wilcoxon signed-rank test $z_W = -4.059, p < 0.01\%$).

Combining the results, we can arrive at some important conclusions. We observe that for the buyer, not disclosing the auction's scoring rule was highly detrimental to their surplus. Concealing the scoring rule in a multi-dimensional sealed-bid first score procurement auction significantly lowers buyer's surplus. Transparency loss, creates a negative impact for the buyer. Our first study provided evidence that the non-transparent auction was a strictly worse mechanism in the multi-dimensional sealed-bid first-offer context under study here. On the other hand, winning suppliers do obtain a higher profit in MDBC than both theoretically and empirically TS would suggest.

6. Study 2: Ex-Post Transparency (EPT)

As mentioned when the MDBC mechanism was introduced, in some instances, it is impossible for a buyer to transparently communicate either the exact weight or the metric for quality. The idea of "*I'll know quality when I see it*" is the motivation for our next study. We introduced a new experimental treatment, which we have called a *multi-dimensional beauty-contest with ex-post transparency* (EPT). In this case, the weight for the quality dimension is not known to bidders before submitting their bids, but it is revealed to all after the bidding took place.

Our main goal with this new study is to give an additional layer of realism to our multi-dimensional beauty-contest. In practice, it is important to know what happens if learning the weight for quality is too costly for the buyer before they open a request for quotes, but is effectively learned after the offers are received. This type of ex-post determination of weights in practice may open the door for corruption such as ex-post favoritism for one bidder, yet a controlled laboratory experiment precludes the occurance of such corruption. In absence of corruption, EPT offers a middle ground between TS and MDBC: The buyer's actions do not appear as fully transparent as in TS, but offers some feedback to bidders about the buyer's preferences after the contracts are awarded.

In order to make this a non-trivial manipulation (recall that for study 1, quality weight remained constant at 10), it was necessary to change the quality weight each round. For study 2, we drew a series of 40 uniformly random weights between 10.0 and 13.0, one for each round. Neither the weights nor the range were communicated to the participants in the beginning of the experiment. We replicated the two treatments for study 1 (TS and MDBC) under these different weights, to make the results from the new EPT treatment comparable. Aside from changing the weights from round to round as opposed to being fixed at 10, all other details of the design, parameterization and protocol for study 1 remained the same for study 2. We recruited 24, 24 and 16 participants for our new EPT, TS and MDBC treatments respectively. None of the subjects recruited had participated in the previous study, they were recruited from the same subject pool, and played in the same facility under the supervision of the same experimenter. We used the same instructions for treatments TS and MDBC as in study 1, and made minimal modifications for EPT. The complete instructions are available in Appendix A.1.

Just as the only difference between TS and MDBC in study 1 was the revelation or not of the quality weights, in study 2 the difference between the three cases was that, in TS, subjects were informed of the quality weight *before* submitting their bid, while in EPT they were informed of the quality weight *after* submitting their bid, and in MDBC they were not informed of the quality weight at all. As such, MDBC and TS are identical to the earlier study except for the varying weights.

7. Results for Study 2

Similarly to §5, in this section we measure differences in both bidding decisions and mechanism performance with regards to the three treatments in study 2. In §7.1, we provide an overview of bidding behavior and compare the outcomes of the TS, MDBC and EPT treatments. In §7.2, we compare the performance of the three treatments relative to the theoretical predictions in terms of outcome for the buyer and for the winning bidder.

7.1. Bidding decisions

Individual bidding behavior results for the TS and MDBC treatments are summarized in Table 4 and Table 5. Complete summary statistics for each subject are included in Tables B3, B4 and B5 in Appendix B.

Table 4: Comparison of actual average bidder behavior and outcome (TS vs. MDBC)
(*Unit of analysis: Individual bidder, averaged over 40 auctions, $J_{TS} = J_{EPT} =$ 24 bidders, $J_{MDBC} =$ 16 bidders*)

		Quality $\bar{q}_J$	Price $\bar{p}_J$	Score $\bar{s}_J$	Profit $\overline{\Pi}_J$
Actual TS	Mean	545.0	5786.5	541.7	216.7
	SD	94.8	998.4	1328.5	121.3
Actual EPT	Mean	517.0	5203.9	795.1	226.9
	SD	93.1	1268.4	1420.7	86.5
Actual MDBC	Mean	483.3	4830.9	769.6	243.9
	SD	87.8	1565.8	1574.9	172.9
Mann-Whitney tests	z_{M-W}	**2.181**	**2.733**	-1.132	0.055
(TS vs MDBC)	p-value	**(.0292)**	**(.0063)**	(.2577)	(.9560)
Mann-Whitney tests	z_{M-W}	1.000	1.856	-0.845	-0.165
(TS vs EPT)	p-value	(.3173)	(.0635)	(.3979)	(.8690)
Mann-Whitney tests	z_{M-W}	1.353	1.104	-0.469	-0.690
(EPT vs MDBC)	p-value	(.1761)	(.2695)	(.6388)	(.4901)

Table 5: Comparison of bidder behavior and outcome (actual vs. predicted), TS treatment
(*Unit of analysis: Individual bidder, averaged over 40 auctions, $J_{TS} =$ 24 bidders*)

		Quality $\bar{q}_J$	Price $\bar{p}_J$	Score $\bar{s}_J$	Profit $\overline{\Pi}_J$
Actual TS	Mean	545.0	5786.5	541.7	216.7
	SD	94.8	998.4	1328.5	121.3
Predicted TS	Mean	576.1	4526.1	2184.6	261.7
	SD	37.2	271.0	162.6	94.0
Wilcoxon tests	z_W	-1.629	**3.771**	**-4.286**	-1.171
	p-value	(.1034)	**(.0002)**	**(<.0001)**	(.2414)

Table 5 largely replicates the results from Table 2. In reality, the difference in observed bidding behavior regarding quality, and subsequently, profits, is not statistically significantly different to what theory predicts. Table 4, on the other part, shows that when the EPT is introduced, it effectively pegs in between of the two "purer" mechanisms. We observe that bidding behavior is not statistically different between bidding knowing the score before bidding versus after bidding takes place. Bidding behavior in EPT is also not significantly different from MDBC, and the only difference that became significant was between qualities and prices when comparing MDBC and TS.

Our design for study 2 also allows us to assess some learning effects. Subjects in TS for study 2 were not able to bid based on a fixed weight as in study 1, and many bidders might have focused on constructing a bid that was a *number* instead of a *function.* Because the weights moved each round in study 2 (and were either unavailable in MDBC or available after bidding in EPT), it was harder to refine a bidding strategy for all participants. Since weights are representations of a buyer's preference for quality, study 2 effectively mimics the open progressing bidding environment that Chen-Ritzo et al. (2005) had designed but places it in the context of a multi-attribute sealed bid auction. By concealing the explicit preference for non-price attributes, Chen-Ritzo et al. (2005) simplified the bidding environment faced by subjects. When bidders cannot get used to a static, constant, preference for quality as in study 1, and need to figure out how to bid for each instance of quality preference, the differences among the bidding behaviors are reduced. In consequence, when weights were variable, bidders in all three treatments of this study bid more similarly, whereas in study 1, it was easier for bidders to learn competitively to bid when one of the moving parts (buyer's preference for quality, or weight) was constant.

Just as in study 1 and Tables 1-2, since Tables 4 and 5 include a summary of the bidding behavior across all bidders (including winners and losers), to understand the effect of the different mechanisms for the buyer, we next present the results by auction, including the behavior of winners only, and this is presented in the following subsection.

7.2. Mechanism performance and welfare

Next, we consider our main comparison for Study 2: the mechanism performance due to transparency by treatment. Since we kept the same realizations of types fixed by auction (round), we average the outcomes of all parallel auctions taking place in every period, for each treatment. This makes every auction round fully comparable to the same round conducted in each treatment. Those averages are reported in Table 6.

Table 6: Comparison of auction performance per treatment (winning bids only).

	Buyer surplus (Winning score)	*Bidder outcome (Winner profit)*
Predicted Theory (*TS*)	3313.10	1046.95
TS	2778.27	866.71
EPT	2778.47	907.64
MDBC	2690.88	975.46
Wilcoxon (*Theory-TS*)	**4.167**	**3.737**
	(< .0001)	**(.0002)**
Wilcoxon (*TS-EPT*)	0.86	-0.941
	(.3897)	(.3468)
Wilcoxon (*TS-MDBC*)	**1.626**	**-2.097**
	(.1039)	**(.0360)**
Wilcoxon (*EPT-MDBC*)	1.331	1.116
	(.1833)	(.2646)

Note: *Unit of analysis: T = 40 rounds, each averaged over all parallel auctions per treatment. Values reported are averages per auction of each performance metric. Test p-values in parentheses.*

Very importantly, we observe that even when looking at the auction round as unit of analysis, the EPT treatment's results sit neatly in between the "pure" versions of MDBC and TS. In practice, since now we have three levels of transparency (TS is full transparency, MDBC is full non-transparency, and EPT is a partial transparency), our results suggest that when communicating preferences for quality is possible, the outcome is better than concealing the amount of preference for quality completely (with a p-value of about 10%), even in the more complex and noisy setting of study 2. In consequence, the results appear as directionally robust, particularly if we consider the noise that the variation in weights over time introduces in terms of learning how to bid over time. In terms of profits for the winning supplier, MDBC still yields the highest profits, TS is the lowest profits, and EPT sits in between the two.

Using each auction round as unit of analysis, we also observed 57.92% of efficient assignments in the TS treatment, 60.00% in the EPT treatment, and 60.63% in the MDBC treatment. Similar to study 1, in approximately 2 out of 5 auctions in study 2 we observed winning bidders who should not have won that instance. This has an important implication to explain why empirical results at the auction round level differ from the theoretical predictions for SBNE (TS) in both of our studies, as in about 40% of the instances, the suppliers with the highest efficiency, and thus who are in better position to provide the buyer with the highest surplus possible, are not winning the auction, irrespective of the transparency level.

8. Conclusion

We compared the performance of two similar multi-dimensional sealed-bid procurement auction mechanisms, one considering transparent rules, and another concealing said rules. Based on our results, even in the case of simple uncertainty of the size of the rule coefficient, the effect of transparency loss is statistically significant and substantial in magnitude. This negative impact is mitigated when the rule is truthfully revealed ex-post.

Our results are informative for sealed-bid purchase decisions where price and non-price components matter, and the latter have a monotonic effect in the buyer's utility. The policy implication of this result is that succumbing to social or political pressures towards relaxing the transparency requirements by procurement agencies such as ChileCompra can lead to the destruction of the institutional credibility of the transparent process. Therefore, we recommend in favor of transparency rules and against non-transparent assignment criteria, when possible.

For practitioners, we offer two specific recommendations for whenever a scoring rule is used: (a) the weights for each attribute should be clearly communicated to bidders, and (b) whenever possible, the measurement scale for attributes should be clearly explained, specified and measurable. This is important, because suggesting a scoring rule without clearly providing appropriate metrics for the quality attributes is analogous to concealing the specific weights of the scoring rule. Such lack of transparency is likely to lead to poorer outcomes. Moreover, in cases where the assessment of quality is costly, communicating the assignment rule after the auction reduces this negative impact for the buyer.

We acknowledge the appeal for a buyer in concealing the weights on certain attributes. Some attributes might be unpopular or controversial (e.g., buyers may have a stated preference to award business to historically under-represented firms, or to favor a firm that commits to employ a minimum number of local workers, etc.). There also might be genuine uncertainty of the worth of certain characteristics for the buyer, making it costly and difficult to determine the true value of the weight to place on such characteristics *a priori*. Yet, where it is possible to determine the value that each attribute provides to the buyer before request for bids is issued, our results provide the first clear evidence of the detrimental impact of using non-transparent assignment rules in sealed-bid procurement auctions for the buyer. We also observe that this lower profit is driven fundamentally by less competitive bidding in presence of non-transparency, translating into higher profits for the winning bidders.

An important (though non-trivial) extension to this research would occur when changing the production costs structural framework from independent private values to a *common values* case. In the common values framework, bidders only receive a signal that is correlated with their (common and unknown) cost structure realization, and do not know the cost realization until completion. That opens the door to the presence of the winner's curse through lack of information about her/his own true cost structure until the bidder wins a contract. Although our results under independent private values demonstrate that lack of transparency reduces buyer's surplus, the presence of the winner's curse would be expected to reduce the winner's profits. Additionally, the buyer's surplus would be adjusted downwards to consider supplier default under negative profits. Therefore, the overall effect of introducing common values would be expected to accentuate the gap between the two scenarios under comparison.

9. References

Asker, J., E. Cantillon. 2008. Properties of scoring auctions. *RAND Journal of Economics*, 39(1): 69-85.

Bajari, P., Hortacsu, A. 2003. The winner's curse, reserve prices, and endogenous entry: Empirical insights from eBay auctions. *RAND Journal of Economics*, 34(2), 329-355.

Bajari, P., S. Houghton, S. Tadelis. 2014. Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs. *American Economic Review*, 104(4): 1288-1319.

Bichler, M. 2000. An experimental analysis of multiattribute auctions. *Decision Support Systems*, 29 (3): 249-268.

Bichler, M., J. Kalagnanam. 2005. Configurable offers and winner determination in multiattribute auctions. *European Journal of Operational Research*, 160: 380-394.

Branco, F. 1997. The design of multidimensional auctions. *RAND Journal of Economics*, 28(1), 63-81.

Che, Y.K. 1993. Design Competition through Multidimensional Auctions. *RAND Journal of Economics*, 24(4): 668-680.

Chen-Ritzo, C.-H., T.P. Harrison, A.M. Kwasnica, D.J. Thomas. 2005. Better, Faster, Cheaper: An Experimental Analysis of a Multiattribute Reverse Auction Mechanism with Restricted Information Feedback. *Management Sci.*, 51(12): 1753-1762.

Chen, Y., P. Katuščák, E. Ozdenoren. 2013. Why Can't a Woman Bid More Like a Man? *Games and Economic Behavior*, 77(1): 181-213.

Corvalán, M. 2013. Jun 16. Las tareas pendientes de ChileCompra a 10 años de su creación ("Pending tasks in ChileCompra 10 years from its inception", in Spanish). *La Tercera: Negocios* (Print edition): 12-13.

Elmaghraby, W. 2007. Auctions Within E-Sourcing Events. *Production and Operations Management*, 16(4), 409-422.

Engelbrecht-Wiggans, R., E. Haruvy, E. Katok. 2007. A Comparison of Buyer-Determined and Price-Based Multiattribute Mechanisms. *Marketing Sci.*, 26(5): 629-641.

Fischbacher, U. 2007. z-Tree: Zurich Toolbox for Ready-made Economic Experiments. *Experimental Economics*, 10(2): 171-178.

Frechette, G.R. 2012. Session-effects in the laboratory. *Experimental Economics*, 15: 485-498.

Fugger, N., E. Katok, A. Wambach. 2015. Collusion in Dynamic Buyer-Determined Reverse Auctions. *Management Sci.*, 62(2): 518-533.

Goeree, J. K., C. A. Holt, T. R. Palfrey. 2002. Quantal Response Equilibrium and Overbidding in Private-Value Auctions. *Journal of Economic Theory*, 104(1): 247-272.

Gupta, D., E. Snir, Y. Chen. 2015. Contractors' and Agency Decisions and Policy Implications in A+B Bidding. *Production and Operations Management*, 24(1): 159-177.

Gur, Y., L. Lu, G.Y. Weintraub. 2013. Framework Agreements in Procurement: An Auction Model and Design Recommendations. Working paper, Columbia Business School.

Haruvy, E., E. Katok. 2013. Information and Revenue in Procurement Auctions. *Production and Operations Management*, 22(1): 19-35.

Joglekar, N. R., J. Davies, E.G. Anderson. 2016. The Role of Industry Studies and Public Policies in Production and Operations Management. *Production and Operations Management*, 25(12): 1977–2001.

Karaer, O., T. Kraft, J. Khawam. 2015. Buyer and Nonprofit Levers to Improve Suppliers' Environmental Performance. Darden Business School Working Paper No. 2575878.

Karmarkar, U.S., R.C. Pitbladdo. 1997. Quality, Class, and Competition. *Management Sci.*, 43(1): 27-39.

Katok, E., A.E. Roth. 2004. Auctions of Homogeneous Goods with Increasing Returns. *Management Sci.*, 50(8): 1044-1063.

Katuščák, P., F. Michelucci, M. Zajiček. 2015. Does Feedback Really Matter in One-Shot First-Price Auctions? *Journal of Economic Behavior & Organization*, 119: 139-152.

Kersten, G.E. 2014. Multiattribute Procurement Auctions: Efficiency and Social Welfare in Theory and Practice. *Decision Analysis*, 11(4): 215-232.

Ketokivi, M.A., R.G. Schroeder. 2004. Perceptual Measures of Performance: Fact or Fiction? *Journal of Operations Management*, 22(3): 247-264.

Keynes, J.M. 1936. The State of Long-Term Expectation. *The General Theory of Employment, Interest and Money*, Chapter 12. New York: Harcourt Brace and Co.

Kostamis, D., D.R. Beil, I. Duenyas. 2006. Total-Cost Procurement Auctions: Should the Buyer Reveal Suppliers' Quality? Ross School of Business Working Paper Series No. 1058.

Kostamis, D., D.R. Beil, I. Duenyas. 2009. Total-Cost Procurement Auctions: Impact of Suppliers' Cost Adjustments on Auction Format Choice. *Management Sci.*, 55(12): 1985-1999.

Lewis, G., P. Bajari. 2011. Procurement contracting with time incentives: Theory and evidence. *The Quarterly Journal of Economics*, 126(3): 1173-1211.

Lucas, R.E. 1976. Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy,* 1(1): 19-46.

McAfee, R.P., J. McMillan. 1987. Auctions and Bidding. *Journal of Economic Literature*, 25(2): 699-738.

Moreno, A., C. Terwiesch. 2014. Doing Business with Strangers: Reputation in Online Service Marketplaces. *Information Systems Research*, 25(4): 865-886.

Nagel, R. 1995. Unraveling in Guessing Games: An Experimental Study. *American Economic Review*, 85(5): 1313-1326.

Ockenfels, A., R. Selten. 2005. Impulse Balance Equilibrium and Feedback in First Price Auctions. *Games and Economic Behavior*, 51(1): 155-170.

Olivares, M., G.Y. Weintraub, R. Epstein, D. Yung. 2012. Combinatorial Auctions for Procurement: An Empirical Study of the Chilean School Meals Auction. *Management Sci.*, 58(8), 1458-1481.

Reyes, C. 2015. Feb 7. En 11 años el Fisco ahorró US$ 4 mil millones por ChileCompra ("In 11 years the Government saved US$ 4 billion through ChileCompra", in Spanish). *La Tercera: Negocios* (Print edition): 46 (upper half).

Riley, J.G., W.F. Samuelson. 1981. Optimal Auctions. *American Economic Review*, 71(3): 381-392.

Schmidt, D., R. Shupp, J.M. Walker, E. Ostrom. 2003. Playing Safe in Coordination Games: The Roles of Risk Dominance, Payoff Dominance, and History of Play. *Games and Economic Behavior*, 42(2): 281-299.

Snir, E., D. Gupta. 2011. A+B Bidding: Contractor and Agency Perspectives. Working paper, University of Minnesota Department of Industrial and Systems Engineering.

Wan, Z., D.R. Beil, E. Katok. 2012. When does it pay to delay supplier qualification? Theory and experiments. *Management Sci.*, 58(11): 2057-2075.

Zheng, C.Z. 2000. Optimal Auction in a Multidimensional World. Working paper, Northwestern University Department of Economics.

Appendix A: Experimental Materials

Included below are the participants' instructions as delivered at the start of every session, and an example screenshot of the experimental computer interface

A.1. Instructions

(Instructions for both conditions (TS vs. MDBC auction) were identical, except for one paragraph regarding the definition of "Buyer Quality Weighting" which varied from one treatment to the other as indicated below.)

Instructions

You are about to participate in an experiment in the economics of decision-making. If you follow these instructions carefully and make good decisions, you will earn money that will be paid to you in cash at the end of the session. If you have a question at any time, please raise your hand and the experimenter will come to your station and answer it. We ask that you not talk with one another for the duration of the experiment.

Government agencies or companies often ask suppliers to participate in a Procurement Auction by submitting a bid to supply equipment or construct a building, a bridge or a highway. These bids are often at a combination of price, quality and possibly other factors.

In this experiment, you will be in the role of a **Supplier (Bidder)** who participates in 40 Procurement Auctions, competing against **three (3)** human competitors (there are 4 suppliers in total). Your competitors are determined through a random match with other individuals in the room. You will be matched with different people in each round.

The session consists of **40** separate auction rounds, and you will submit a bid in each round. The unit of exchange in all the transactions is called experimental currency units (ECUs). At the end of the session, your earnings in ECUs will be converted to US dollars at a pre-specified rate, and paid to you in cash.

<u>How you earn money</u>

Each bid you make consists of (A) a money price you request in ECUs ("Your Price"), and (B) a quality level ("Your Chosen Quality"). The buyer has established a scoring rule to compare bids submitted by each supplier. Supplier Score **increases** with Chosen Quality, and decreases with Price:

Supplier Score = (Buyer Quality Weighting * Your Chosen Quality) – Your Price

"Buyer Quality Weighting" represents the evaluation the buyer makes about the quality level offered by any supplier. It is a positive constant, **known and common** to all bidders. (displayed for TS only)

"Buyer Quality Weighting" represents the evaluation the buyer makes about the quality level offered by any supplier. It is a positive constant, **unknown and common** to all bidders. Only the Buyer knows their Quality Weighting and the Supplier Score. (displayed for MDBC only)

"Buyer Quality Weighting" represents the evaluation the buyer makes about the quality level offered by any supplier. It is a positive constant, **common** to all bidders. The Buyer only learns its Quality Weighting **after** all offers are received, and uses it to compute all Suppliers' Scores. (displayed for EPT only)

The highest score among all four competitors **wins** the contract (Any ties in supplier scores are broken at random). You earn money each time **you win** a contract at a good price and appropriate quality. If you **win** the contract:

$$\text{Your Earnings} = \text{Your Price} - \text{Your Cost}$$

If you **don't win** the contract, your profit, your earnings and your cost will be 0 regardless of your bid.

Only if you win the contract, you incur a production cost ("Your Cost") to produce the good at Your Chosen Quality. This cost varies based on Your Efficiency and Fixed Costs:

$$\text{Your Cost} = \left(0.5 \times \frac{(\text{Your Chosen Quality})^2}{\text{Your Efficiency}}\right) + \text{Fixed Costs}$$

You will see Your Efficiency before deciding how much to bid, but not the efficiency of any other competitor. For each bidder, the higher the efficiency, the cheaper it is to offer the chosen quality level. Your Efficiency will change from round to round; it is independent from your competitors' efficiencies, and is unrelated to your efficiency in any other auctions. Your Efficiency is a random number within the range **[10, 90]**, with all values in that range being equally likely. Your competitors' efficiencies are also random numbers drawn from the same range, independent from yours and from each other's. Fixed Costs are the same and are known for all bidders before each auction, but are only incurred by the winner.

For each auction, enter a number for Your Price and another for Your Chosen Quality in the boxes on your computer screen, and then click the "Make Offer" button. For price decisions you can chose any number between 0 and above. For quality levels, you can choose any number from 0 and above. Your decisions can have up to two decimals.

With the above information, you will be able to calculate Your Cost and Your Earnings in each round in case you win. Our software has an embedded simulator you can use for this. Each bidder knows only their **own** cost, but not the cost of any other bidder.

Note: If a supplier bids his/her Price below his/her Cost, and wins the auction, he/she loses money. Therefore, carefully choose your price and quality. Hint: Use the Simulator tool for this purpose.

<u>Information you will see at the end of each auction</u>

At the end of each auction you will see the following information:

- Your own efficiency and bid (price and quality) in this auction.
- Your Cost (which will only be different from 0 if you win).
- The winning bid (price and quality).
- Your earnings from the auction.

You will also have access to this information for all past auctions.

<u>How you will be paid</u>

At the end of the session you will see the final screen summarizing your earnings for the session. This screen will calculate your net profits from the 40 auctions, convert them to US dollars at the rate of 1,100 ECU per $1, and add them to your $5 participation fee.

Please use this information to fill out your check-out form and wait quietly until the experimenter calls you to come to the front of the room and be paid your earnings in private and in cash. After you have been paid, you may leave the laboratory.

A.2. Computer interface screenshots

This example is for the MDBC treatment; the TS case was virtually identical (the term "UNKNOWN" was replaced by the corresponding number).

Bidding Entry Screen:

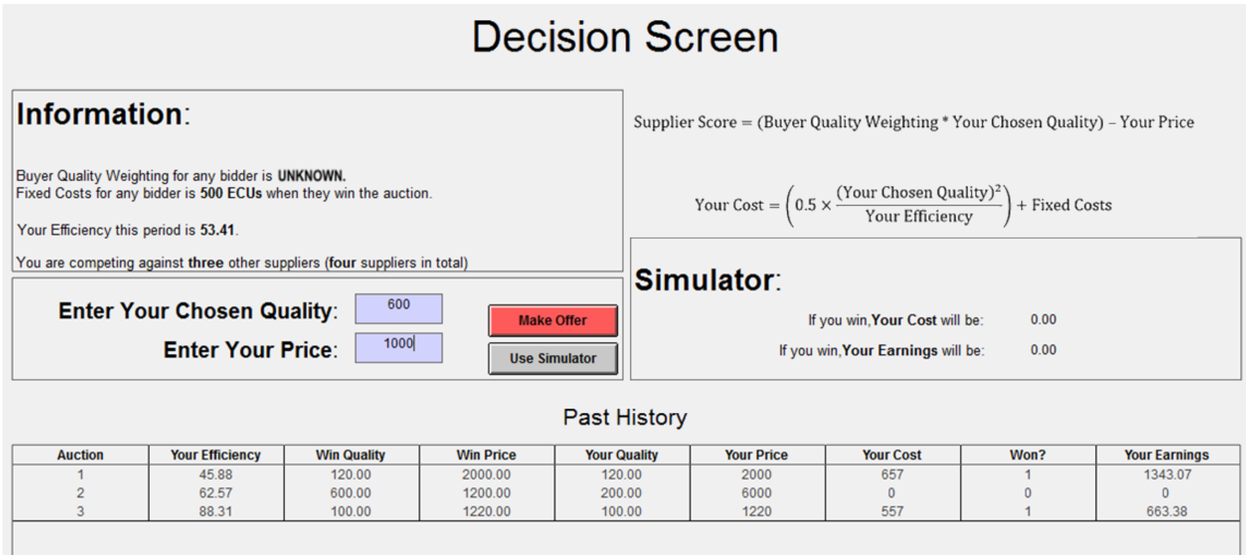

Auction	Your Efficiency	Win Quality	Win Price	Your Quality	Your Price	Your Cost	Won?	Your Earnings
1	45.88	120.00	2000.00	120.00	2000	657	1	1343.07
2	62.57	600.00	1200.00	200.00	6000	0	0	0
3	88.31	100.00	1220.00	100.00	1220	557	1	663.38

Results Screen:

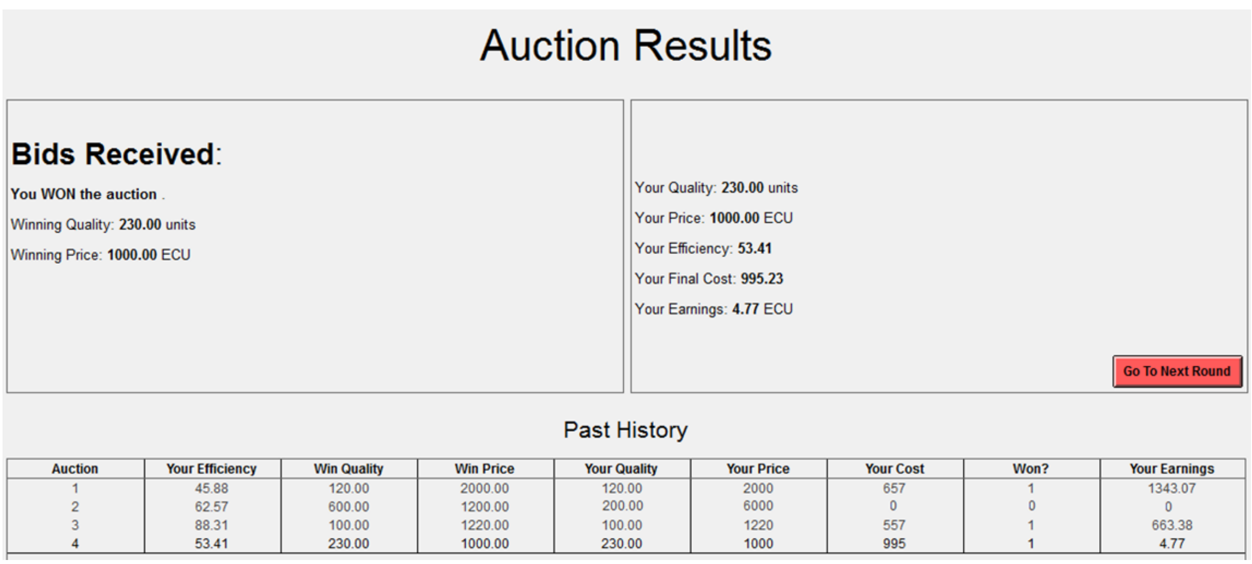

Auction	Your Efficiency	Win Quality	Win Price	Your Quality	Your Price	Your Cost	Won?	Your Earnings
1	45.88	120.00	2000.00	120.00	2000	657	1	1343.07
2	62.57	600.00	1200.00	200.00	6000	0	0	0
3	88.31	100.00	1220.00	100.00	1220	557	1	663.38
4	53.41	230.00	1000.00	230.00	1000	995	1	4.77

Appendix B: Average individual bidding behavior and outcome

Average individual bidding behavior is summarized in Tables B1-B4. Note that both sets of predicted behaviors in each treatment does not perfectly coincide, because we controlled that every round and every auction had the same four random realizations (key for subsequent welfare analyses), and subjects were randomly rematched every period. For every subject, the stream of realizations depended on which "role" they had in the auction (i.e., what order statistic in terms of efficiency types they had been assigned to in every period); that assignment, undisclosed to participants, was randomized during the session.

Table B1: Study 1 - Average bidder behavior and outcome (actual vs. predicted)
Sealed-bid Transparent Scores. (TS)
(*Unit of analysis:* $J_{TS} = 24$ *subjects, each averaged over 40 auctions*)

	Actual subjects behavior					*Predicted behavior for TS*					
j	$\#w_j$	$\overline{q_j}$	$\overline{p_j}$	$\overline{s_j}$	$\overline{\Pi_j}$	$\#w_j^*$	$\overline{q_j^*}$	$\overline{p_j^*}$	$\overline{s_j^*}$	$\overline{\Pi_j^*}$	$\overline{z_j}$
11001	1	357.8	4243.5	-665.8	19.2	13	517.4	3608.5	1565.1	251.1	51.7
11002	11	615.5	5774.2	381.1	237.6	10	494.1	3463.0	1477.8	197.5	49.4
11003	15	548.0	4215.9	1264.1	148.4	12	492.1	3450.6	1470.3	223.2	49.2
11004	4	491.3	5263.7	-351.2	88.4	13	508.4	3552.6	1531.6	252.1	50.8
11005	9	584.0	4701.0	1139.0	156.2	13	532.5	3703.3	1622.0	234.5	53.3
11006	17	646.8	5557.4	910.9	84.3	9	491.9	3449.7	1469.8	178.0	49.2
11007	15	574.5	4439.7	1305.4	139.7	6	522.9	3643.0	1585.8	121.6	52.3
11008	7	487.6	3830.2	1045.9	175.7	7	444.9	3155.8	1293.5	142.3	44.5
11009	15	468.9	3487.6	1200.9	102.9	8	471.8	3323.9	1394.3	157.2	47.2
11010	10	539.6	4691.3	705.0	194.4	8	449.6	3185.2	1311.1	146.6	45.0
11011	8	506.9	4371.9	696.9	152.0	12	540.3	3751.6	1651.0	240.6	54.0
11012	8	425.2	3416.2	835.9	162.1	9	500.7	3504.1	1502.5	181.9	50.1
12013	7	561.3	5343.3	269.5	283.8	13	532.2	3701.1	1620.7	258.2	53.2
12014	2	118.4	3401.7	-2217.4	3.9	9	475.0	3344.0	1406.4	176.8	47.5
12015	19	508.4	3962.5	1121.8	532.7	16	524.3	3651.9	1591.2	312.0	52.4
12016	8	165.2	1257.5	394.3	48.1	11	494.6	3466.3	1479.8	201.4	49.5
12017	12	466.7	4107.0	559.8	197.8	8	518.8	3617.3	1570.4	159.9	51.9
12018	14	607.3	5726.4	346.8	289.8	11	490.5	3440.4	1464.2	210.9	49.0
12019	12	500.7	4659.0	348.3	288.7	13	531.8	3698.6	1619.1	258.9	53.2
12020	13	466.4	4247.4	416.3	312.4	10	493.2	3457.2	1474.3	198.1	49.3
12021	8	476.0	4644.1	115.9	140.5	5	478.6	3366.2	1419.7	87.8	47.9
12022	14	642.2	5868.4	553.4	164.5	9	514.1	3587.9	1552.7	174.9	51.4
12023	3	293.7	4276.9	-1339.9	50.8	8	452.7	3204.4	1322.7	161.8	45.3
12024	8	441.9	4504.7	-86.2	188.4	7	461.0	3256.1	1353.6	125.8	46.1
Mean		**478.9**	**4416.3**	**372.9**	**173.4**		**497.2**	**3482.6**	**1489.6**	**193.9**	**49.7**
SD		**133.5**	**988.2**	**841.2**	**114.2**		**28.1**	**175.4**	**105.2**	**53.0**	**2.8**

Note: Variables presented are number of wins ($\#w$), qualities (q), prices (p), scores (s), profits (Π) and private types (z); with upper bars denoting averages over all 40 periods and asterisks indicating predicted behavior, measured as the decisions that a decision-maker would have made if pre-programmed to make SBNE decisions knowing its own efficiency type z_j and the scoring rule.

Table B2: Study 1 - Average bidder behavior and outcome (actual vs. TS predicted)
Multi-dimensional Beauty-contest (MDBC).
(*Unit of analysis:* $J_{MDBC} = 24$ *subjects, each averaged over 40 auctions*)

	Actual subjects behavior					*Predicted behavior for TS*					
j	$\#w_j$	$\overline{q_j}$	$\overline{p_j}$	$\overline{s_j}$	$\overline{\Pi_j}$	$\#w_j^*$	$\overline{q_j^*}$	$\overline{p_j^*}$	$\overline{s_j^*}$	$\overline{\Pi_j^*}$	$\overline{z_j}$
21001	15	282.2	2026.6	795.2	141.6	11	513.3	3582.9	1549.8	214.7	51.3
21002	6	453.6	5122.1	-585.8	123.5	8	514.1	3588.4	1553.0	158.4	51.4
21003	6	514.0	5705.0	-565.0	211.8	7	417.9	2987.2	1192.3	134.4	41.8
21004	11	544.3	4920.0	522.5	248.8	11	559.2	3869.8	1721.9	209.2	55.9
21005	5	415.5	5147.5	-992.5	243.5	13	507.6	3547.2	1528.3	261.3	50.8
21006	17	529.1	4389.8	901.5	333.6	10	483.4	3396.4	1437.8	203.5	48.3
21007	12	706.5	6547.6	517.4	236.6	10	533.0	3706.5	1623.9	187.2	53.3
21008	5	417.7	6038.1	-1860.9	134.0	11	540.1	3750.8	1650.5	212.8	54.0
21009	7	588.4	6055.1	-171.1	214.5	7	473.5	3334.3	1400.6	140.9	47.3
21010	16	622.4	6007.3	216.2	402.9	10	448.6	3178.5	1307.1	183.8	44.9
21011	8	675.9	7060.0	-301.3	203.0	8	464.9	3280.9	1368.6	149.9	46.5
21012	12	494.8	4352.5	595.0	403.6	14	510.9	3568.4	1541.1	270.5	51.1
22013	4	180.0	2402.5	-602.5	36.8	13	514.7	3591.6	1555.0	252.1	51.5
22014	9	562.1	5331.1	289.4	350.5	12	517.5	3609.5	1565.7	221.7	51.8
22015	10	528.1	4490.2	791.0	248.6	12	469.2	3307.4	1384.4	224.4	46.9
22016	12	437.5	3941.8	433.2	149.4	9	513.6	3585.3	1551.2	187.1	51.4
22017	9	515.4	4925.0	228.5	233.1	8	455.5	3222.1	1333.2	165.2	45.6
22018	22	528.5	3526.0	1759.0	187.5	10	525.4	3658.5	1595.1	190.5	52.5
22019	14	599.7	5242.8	754.0	360.4	13	507.2	3544.9	1526.9	239.8	50.7
22020	3	468.9	4905.5	-216.5	57.2	6	429.2	3057.4	1234.5	113.5	42.9
22021	8	554.6	4926.3	620.0	240.0	10	536.6	3728.7	1637.2	204.4	53.7
22022	7	479.8	4847.5	-50.0	223.1	12	516.1	3600.3	1560.2	223.6	51.6
22023	13	554.8	4785.5	762.0	306.2	8	469.2	3307.2	1384.3	161.7	46.9
22024	9	563.9	5275.8	363.5	229.3	7	512.6	3578.4	1547.1	142.5	51.3
Mean		**509.1**	**4915.5**	**175.1**	**230.0**		**497.2**	**3482.6**	**1489.6**	**193.9**	**49.7**
SD		**113.1**	**1151.2**	**751.9**	**96.5**		**36.2**	**226.5**	**135.9**	**41.9**	**3.6**

Notes:

i. Variables presented are number of wins (#w), qualities (q), prices (p), scores (s), profits (Π) and private types (z), with upper bars denoting averages over all 40 periods and asterisks indicating predicted behavior, measured as the decisions that a decision-maker would have made if pre-programmed to make SBNE decisions knowing its own efficiency type z_j and the scoring rule.

ii. Predicted averages do not correspond to the MDBC mechanism, but to the TS case, and are included here only as a reference point, given that participants would have needed the actual scoring rule to be able to determine those predicted decisions were optimal.

Table B3: Study 2 - Average bidder behavior and outcome (actual vs. predicted)
Sealed-bid Transparent Scores. (TS)
(*Unit of analysis:* J_{TS} *= 24 subjects, each averaged over 40 auctions*)

	Actual subjects behavior					*Predicted behavior for TS*					
j	$\#w_j$	$\overline{q_J}$	$\overline{p_J}$	$\overline{s_J}$	$\overline{\Pi_J}$	$\#w_j^*$	$\overline{q_J^*}$	$\overline{p_J^*}$	$\overline{s_J^*}$	$\overline{\Pi_J^*}$	$\overline{z_J}$
31001	10	506.9	5485.0	406.9	261.3	9	612.3	4801.3	2349.7	243.1	52.7
31002	11	704.3	6675.0	1471.8	371.3	14	625.2	4874.9	2393.9	347.9	54.1
31003	14	776.3	7950.6	1047.9	23.6	8	537.3	4240.1	2013.0	178.8	46.4
31004	9	594.8	6690.0	234.5	263.8	9	564.6	4468.6	2150.0	244.4	48.4
31005	7	547.0	5380.6	985.2	167.2	4	532.5	4215.3	1998.1	96.8	45.9
31006	15	538.8	4190.0	2042.7	245.8	10	580.3	4546.1	2196.6	241.4	50.2
31007	1	388.1	7570.0	-3077.3	29.9	15	586.2	4593.0	2224.7	407.4	50.6
31008	11	549.0	5508.8	893.2	394.6	13	643.9	5016.3	2478.7	347.8	55.6
31009	10	661.3	6203.8	1512.7	271.8	7	550.9	4351.8	2080.0	206.3	47.4
31010	7	440.1	5785.0	-688.6	203.9	9	538.8	4251.9	2020.0	236.0	46.5
31011	12	575.3	5900.0	725.5	397.6	14	609.3	4753.2	2320.8	386.5	52.8
31012	13	595.0	5323.8	1594.8	322.1	8	531.7	4201.2	1989.6	204.4	45.9
32013	14	568.8	4988.4	1630.8	440.5	13	620.1	4851.7	2379.9	336.3	53.5
32014	6	475.3	5454.5	64.3	249.9	14	606.0	4745.1	2315.9	391.3	52.3
32015	16	571.8	5005.2	1601.0	174.9	13	626.0	4897.2	2407.2	332.4	53.9
32016	14	644.6	6523.0	1014.2	131.1	12	552.2	4387.9	2101.7	314.8	47.2
32017	14	574.9	5592.5	1069.9	65.2	4	523.0	4149.5	1958.6	126.3	45.0
32018	8	483.5	4344.3	1271.1	169.3	15	599.5	4684.1	2279.3	387.6	51.9
32019	5	521.3	5882.9	152.1	133.9	4	533.8	4195.1	1986.0	96.5	46.3
32020	1	380.0	6259.5	-1830.7	2.6	7	533.3	4222.6	2002.5	205.7	45.9
32021	8	474.1	3917.5	1642.7	216.9	10	609.5	4803.5	2351.0	283.1	52.2
32022	8	506.4	6321.5	-423.7	116.5	7	556.9	4367.1	2089.1	158.3	48.3
32023	17	581.0	4942.6	1777.5	303.7	12	589.8	4603.6	2231.1	296.6	51.2
32024	9	421.8	6982.5	-2117.1	242.9	9	562.8	4406.4	2112.7	211.8	48.9
Mean		**545.0**	**5786.5**	**541.7**	**216.7**		**576.1**	**4526.1**	**2184.6**	**261.7**	**49.7**
SD		**94.8**	**998.4**	**1328.5**	**121.3**		**37.2**	**271.0**	**162.6**	**94.0**	**3.2**

Note: Variables presented are number of wins (#*w*), qualities (*q*), prices (*p*), scores (*s*), profits (Π) and private types (*z*); with upper bars denoting averages over all 40 periods and asterisks indicating predicted behavior, measured as the decisions that a decision-maker would have made if pre-programmed to make SBNE decisions knowing its own efficiency type z_j and the scoring rule.

Table B4: Study 2 - Average bidder behavior and outcome (actual vs. TS predicted)
Multi-dimensional Beauty-contest with Ex-Post Transparency (EPT).
(*Unit of analysis:* J_{EPT} *= 24 subjects, each averaged over 40 auctions*)

	Actual subjects behavior					*Predicted behavior for TS*					
j	$\#w_j$	$\overline{q_j}$	$\overline{p_j}$	$\overline{s_j}$	$\overline{\Pi_j}$	$\#w_j^*$	$\overline{q_j^*}$	$\overline{p_j^*}$	$\overline{s_j^*}$	$\overline{\Pi_j^*}$	$\overline{z_j}$
51001	10	536.0	5327.5	883.0	207.3	10	573.1	4490.8	2163.4	264.0	49.6
51002	16	507.8	4226.0	1667.8	360.5	12	643.1	5004.1	2471.4	300.4	55.6
51003	1	322.4	8559.0	-4872.4	41.8	12	646.5	5035.0	2489.9	322.8	55.9
51004	3	422.2	4031.5	917.0	129.3	5	508.6	4051.4	1899.7	130.6	43.7
51005	12	603.5	5500.0	1488.6	253.3	8	552.1	4341.6	2073.8	196.4	47.8
51006	15	707.3	7106.0	1085.1	419.7	10	576.9	4535.3	2190.1	286.9	49.7
51007	7	557.4	5068.5	1405.1	252.1	9	596.9	4670.9	2271.4	212.0	51.6
51008	17	686.4	6372.5	1604.9	231.6	13	567.1	4468.8	2150.2	345.6	48.9
51009	11	374.0	2977.5	1357.7	328.4	10	578.6	4539.2	2192.4	280.9	50.0
51010	7	568.8	5799.3	799.1	231.1	10	548.9	4331.1	2067.6	242.7	47.3
51011	13	585.0	5016.5	1831.0	314.4	8	536.7	4275.9	2034.5	237.3	45.9
51012	8	504.7	5202.5	648.0	266.4	13	584.5	4569.5	2210.6	321.3	50.7
52013	4	423.3	5705.0	-825.6	161.8	11	576.4	4472.5	2152.4	287.1	50.4
52014	10	527.4	6133.3	47.4	216.7	7	522.0	4161.0	1965.5	184.3	44.7
52015	10	638.5	6977.5	451.5	229.4	8	563.8	4453.3	2140.9	207.3	48.5
52016	5	525.4	5932.9	120.8	199.8	6	566.7	4472.6	2152.5	139.5	48.7
52017	14	591.5	5259.3	1615.9	147.9	9	603.9	4747.9	2317.6	230.4	51.9
52018	15	496.5	3822.5	1922.8	222.0	9	584.8	4601.2	2229.6	244.5	50.3
52019	13	502.0	4265.0	1544.6	352.4	16	607.1	4767.7	2329.5	435.2	52.2
52020	5	447.9	4575.8	600.5	99.2	11	571.6	4459.0	2144.3	280.2	49.7
52021	10	486.5	4049.4	1589.3	152.6	11	546.7	4295.2	2046.0	275.3	47.4
52022	16	468.9	3520.0	1912.9	179.9	10	554.1	4362.9	2086.7	258.7	47.8
52023	13	517.2	4300.0	1713.1	246.4	11	620.8	4835.4	2370.2	293.1	53.8
52024	5	406.6	5167.5	-425.0	201.8	11	595.0	4685.0	2279.9	305.3	51.1
Mean		**517.0**	**5203.9**	**795.1**	**226.9**	**10.0**	**576.1**	**4526.1**	**2184.6**	**261.7**	**49.7**
SD		**93.1**	**1268.4**	**1420.7**	**86.5**	**2.4**	**33.7**	**241.2**	**144.7**	**66.3**	**3.0**

Notes:

i. Variables presented are number of wins (#*w*), qualities (*q*), prices (*p*), scores (*s*), profits (Π) and private types (*z*), with upper bars denoting averages over all 40 periods and asterisks indicating predicted behavior, measured as the decisions that a decision-maker would have made if pre-programmed to make SBNE decisions knowing its own efficiency type z_j and the scoring rule.

ii. Predicted averages do not correspond to the beauty contest mechanism, but to the TS case, and are included here only as a reference point, given that participants would have needed the actual scoring rule to be able to determine those predicted decisions were optimal.

Table B5: Study 2 - Average bidder behavior and outcome (actual vs. TS predicted) Multi-dimensional Beauty-contest (MDBC).
(*Unit of analysis:* J_{MDBC} *= 16 subjects, each averaged over 40 auctions*)

	Actual subjects behavior					*Predicted behavior for TS*					
j	$\#w_j$	$\overline{q_j}$	$\overline{p_j}$	$\overline{s_j}$	$\overline{\Pi_j}$	$\#w_j^*$	$\overline{q_j^*}$	$\overline{p_j^*}$	$\overline{s_j^*}$	$\overline{\Pi_j^*}$	$\overline{z_j}$
41001	7	407.5	5336.1	-619.7	374.6	10	572.5	4520.9	2181.5	269.3	49.2
41002	13	502.5	4381.0	1418.0	429.8	8	570.9	4469.6	2150.6	194.2	49.5
41003	20	732.0	6881.1	1553.2	493.0	12	594.4	4651.4	2259.8	305.3	51.5
41004	2	451.0	9235.8	-3986.0	70.2	8	552.1	4354.1	2081.4	208.8	47.6
41005	8	440.8	4743.0	362.3	76.0	9	527.4	4156.3	1962.7	247.8	45.7
41006	7	365.0	4676.5	-467.4	354.6	6	544.7	4285.0	2039.9	149.5	47.2
41007	4	422.5	5252.5	-372.5	174.8	15	594.2	4687.6	2281.5	421.0	50.9
41008	19	581.6	5023.1	1734.3	659.3	12	652.3	5084.2	2519.4	297.8	56.2
42009	13	511.7	4448.7	1470.2	137.6	16	604.3	4723.0	2302.7	389.1	52.3
42010	8	424.4	3155.3	1771.2	124.4	8	507.3	4026.0	1884.5	211.5	43.8
42011	7	523.8	5840.2	248.9	127.2	6	544.2	4284.9	2039.9	150.6	47.1
42012	11	454.3	3199.5	2096.4	197.6	7	604.8	4756.9	2323.0	228.5	51.9
42013	10	396.6	2851.8	1751.5	117.5	7	542.4	4291.5	2043.8	180.0	46.7
42014	9	512.5	4413.8	1526.2	157.2	11	609.2	4768.4	2330.0	289.9	52.6
42015	12	532.3	4161.5	1997.6	301.1	12	632.0	4915.4	2418.1	301.4	54.8
42016	10	474.5	3694.5	1829.1	106.8	13	564.6	4443.0	2134.7	342.9	48.7
Mean		**483.3**	**4830.9**	**769.6**	**243.9**		**576.1**	**4526.1**	**2184.6**	**261.7**	**49.7**
SD		**87.8**	**1565.8**	**1574.9**	**172.9**		**39.4**	**289.1**	**173.4**	**80.3**	**3.4**

Notes:

i. Variables presented are number of wins ($\#w$), qualities (q), prices (p), scores (s), profits (Π) and private types (z), with upper bars denoting averages over all 40 periods and asterisks indicating predicted behavior, measured as the decisions that a decision-maker would have made if pre-programmed to make SBNE decisions knowing its own efficiency type z_j and the scoring rule.

ii. Predicted averages do not correspond to the MDBC mechanism, but to the TS case, and are included here only as a reference point, given that participants would have needed the actual scoring rule to be able to determine those predicted decisions were optimal.