Essays on Procurement Auctions and Incentives

Inauguraldissertation
zur
Erlangung des Doktorgrades
der
Wirtschafts- und Sozialwissenschaftlichen Fakultät
der
Universität zu Köln

2025

vorgelegt
von

Jeffrey Virgil Starck

aus
Kleve

Referent:	Prof. Dr. Nicolas Fugger
Koreferent:	Jun.-Prof. Dr. Henrik Franke
Tag der Promotion:	11. Dezember 2025

To my parents.

Acknowledgements

None of this would have been possible without the support of my family and loved ones. Thank you for always being there for me, for encouraging me on the hard days, and for teaching me persistence. I love you more than words can say.

My heartfelt gratitude goes out to my supervisor, Nicolas Fugger, to whom I am deeply indebted. Thank you for providing an exceptional framework and the necessary support for me conduct this research, for always making time to speak with me, and for continually giving me guidance, excellent advice, and food for thought. It was a pleasure having you accompany me on this journey.

I am also indebted to my other co-authors, with whom it was only ever a pleasure to work with. I need to thank Sandra Hartmann for the time and effort devoted to our paper, Philippe Gillen for always providing interesting angles to consider, and Ulrich Laitenberger, not only for the hours spent working on our paper, but also for inviting me to spend some time at Tilburg University. This invitation had far-reaching and unexpected consequences.

A special thanks goes to Henrik Franke for refereeing this thesis and to Andreas Fügener for chairing the defense board.

Finally, I am grateful to everyone in the Supply Chain Management Area, especially Stephanie Rauscher for helping me navigate the treacherous administrative waters, Thomas Vogt for bringing me on board, and Moritz Wegener, David Könen and Dylan Gellert for being a joy to work with.

Contents

List of Figures V

List of Tables VI

1 Introduction 1

2 The Role of Commitment in First-Price Auctions 4

2.1 Introduction 4

2.2 Literature 7

2.3 Theory 10

2.3.1 Model 10

2.3.2 Analysis 11

2.4 Experiment 16

2.4.1 Design 16

2.4.2 Organisation 17

2.4.3 Hypotheses 18

2.5 Results 19

2.5.1 Decision support 24

2.5.2 Simulations 31

2.5.3 Artificial intelligence 32

2.6 Conclusion 38

2.A Experiment: screenshots 40

2.A.1 Pre-treatment 40

2.A.2 Commitment 43

2.A.3 No-Commitment 47

2.A.4 No-Commitment DS 56

2.A.5 Final screen 65

2.B Commitment AI treatment 66

2.C No-Commitment AI treatment 71

3 Commitment in Auctions: Extensions 84

3.1 Introduction 84

3.2 Literature 86

3.3 First-price auctions 88

3.3.1 Bidder sophistication 89

3.3.2 Continuous types with discrete bid and counteroffer grid . 93

3.4 Second-price auctions 98

3.5 Conclusion 106

4 Procurement Auctions with Demand Uncertainty 107

4.1 Introduction 107

4.2 Related Literature 111

4.3 Theory 114

4.3.1 Setting 114

4.3.2 Analysis 116

4.4 Experiment 121

4.4.1 Design 121

4.4.2 Organisation 123

4.4.3 Hypotheses 124

4.4.4 Results 126

4.5 Discussion of the model 140

4.6 Conclusion 142

4.A Experiment: screenshots . . . 145
4.A.1 Pre-treatment . . . 145
4.A.2 Uniform price . . . 148
4.A.3 Price premium . . . 152
4.A.4 Negotiation . . . 156
4.A.5 Survey . . . 160
4.A.6 Final screen . . . 163

5 Improving Incentive Schemes for Procurement Managers 164
5.1 Introduction and motivation . . . 165
5.2 Related Literature . . . 168
5.3 Model . . . 171
5.3.1 General Framework . . . 171
5.4 Analysis . . . 174
5.4.1 First-best . . . 174
5.4.2 Delegation to a procurement manager . . . 179
5.4.3 No-commitment . . . 182
5.4.4 Commitment . . . 187
5.5 Discussion . . . 192
5.5.1 Investment costs of the procurement manager . . . 192
5.5.2 Level of inefficiency . . . 193
5.5.3 Firm as sole beneficiary of a resolved inefficiency . . . 194
5.5.4 Assigning decisions to different managers . . . 195
5.5.5 Incentivizing action instead of outcome . . . 195
5.6 Conclusion . . . 196
5.A Appendix . . . 198
5.A.1 Appendix: First-best . . . 198
5.A.2 Appendix: No-commitment . . . 199
5.A.3 Appendix: Commitment . . . 203

References **206**

List of Figures

2.1 Overview of all treatments ran. 20
2.2 Comparison of bids, prices, counteroffer margins, and buyer expenses across the no-commitment treatments. 29
2.3 Breakdown of auction prices into average costs, margins, and buyer expenses across the no-commitment treatments. 30

4.1 Overview of all treatments ran. 124

5.1 Timeline of information, decisions, and outcomes 172
5.2 Investment decision of high-cost firm depending on i and n 176
5.3 Qualification and investment decision dependent on q and $\bar{i}$ 178

List of Tables

2.1 Average bid, auction price, buyer expenses, and efficiency across treatments. 21
2.2 Random effects model for buyer margin in the no-commitment treatment. 23
2.3 Average bid, auction price, buyer expenses, and efficiency across treatments. 26
2.4 Random effects model for buyer counteroffers in the no-commitment DS treatment. 28
2.5 Average bid, auction price, buyer expenses, and efficiency across AI treatments. 34

4.1 Average bids, prices, buyer profit, efficiency, and rejection rate across treatments. 128
4.2 Summary statistics of survey results. 132
4.3 Random effects panel regression of the sellers' bids for the fixed demand across treatments and including personality traits. 134
4.4 Random effects panel regression of the sellers' requests for the additional demand including personality traits. 136
4.5 Random effects panel regression of the buyers' maximum willingness-to-pay for the additional demand including personality traits. 138
4.6 Random effects panel regression of the price for additional quantity on the winner's costs and round. 139

5.1 First-period prices and inefficiency probabilities for $n = 2$ and $n = 3$ suppliers. 175

5.2 First-best qualification thresholds and total cost depending on $\bar{i}$. 177
5.3 Qualification incentives for different $\bar{i}$ and q 185

Chapter 1

Introduction

This dissertation consists of four distinct chapters. The general theme of these chapters is commitment, demand uncertainty, and incentive schemes, and the influence that these have on (auction) outcomes in a procurement context.

Chapter 2 carries the title *The Role of Commitment in First-Price Auctions*, and is joint work with Nicolas Fugger and Philippe Gillen.[1] In this chapter, we theoretically and experimentally study what the implications are of a buyer who lacks the commitment not to negotiate with the winning seller upon conclusion of the procurement auction. We show theoretically that without commitment, competition breaks down, leading to higher buyer expenses. Experimentally, we find the opposite: buyer expenses are actually significantly *lower* in the no-commitment setting. This result proves robust not only to buyers and sellers who are equipped with decision support, but also holds when letting state-of-the-art artificial intelligence agents take the role of buyer and seller. As a result, at least at this moment in time, a lack of commitment may actually be beneficial to the buyer.

Chapter 3, titled *Commitment in Auctions: Extensions*, is a single-author paper. It expands the theoretical findings of the previous chapter by studying how robust the finding is that a lack of commitment leads to a break down of competition. To this end, we consider different auction formats and seller costs. We demonstrate that our theoretical findings are not a relic of first-price auctions with continuous cost types and a continuous bid grid. Rather, our findings on the importance of commitment carry over more generally. To be precise, in the context of first-price auctions, commitment matters even when sellers are bidding on a discrete grid. However, introducing bidders who differ in their level

[1]Financial support of the *Deutsche Forschungsgemeinschaft* (DFG) is gratefully acknowledged.

of sophistication, we demonstrate that there are cases in which it is beneficial for the buyer not to commit. Next, we consider sealed-bid second-price and English auctions. In both cases—for both continuous and discrete seller costs—we demonstrate that a lack of commitment can lead to a break down of competition. This is an interesting result, given that English auctions are historically considered immune to the commitment issues an auctioneer in a sealed-bid second-price auction is faced with.

Chapter 4, with the title *Procurement Auctions with Demand Uncertainty*, is joint work with Nicolas Fugger and Ulrich Laitenberger.[2] In this chapter, we consider a buyer facing demand uncertainty. The question we seek to answer is: how should the buyer account for this demand uncertainty when running their procurement auction? We consider three distinct possibilities—(1) having the auction price serve as a framework contract that is valid irrespective of the actual demand, (2) defining contracts contingent on the actual demand, and (3) negotiating with the contract supplier only if and when additional demand occurs. We demonstrate theoretically that all three options yield the same expected profits to both buyers and sellers. However, this equivalency does not hold experimentally. As a result of aggressive bidding, sellers do best with contingent contracts. Buyers, on the other hand, do best in the simplest of the three settings we consider: the framework contract. However, we demonstrate that this result is driven by aggressive negotiations on the part of buyers, leading to a breakdown of negotiations in about a third of rounds.

Chapter 5 carries the title *Improving Incentive Schemes for Procurement Managers*, and is joint work with Sandra Hartmann. In this chapter, drawing inspiration from procurement practice, we consider a procurement manager who is rewarded for both securing low initial contract prices and for negotiating

[2]Financial support of the *Deutsche Forschungsgemeinschaft* (DFG), as well as of Tilburg University, is gratefully acknowledged.

price reductions over the duration of the contract. There is anecdotal evidence that such dual incentives create a trade-off between receiving a bonus now and receiving a bonus later. To formally capture this dynamic, we introduce a game-theoretic model of a procurement manager with savings targets in a dynamic game of incomplete information, where the firm delegates costly qualification and investment decisions to the manager. Qualification, which increases competition by adding more bidders, lowers initial prices, but also reduces the future value of investment—the effort devoted to identifying and resolving inefficiencies with the supplier—thereby making it harder to secure a second-period bonus. We show that even under optimal incentivisation, the procurement manager's trade-off leads to inefficient qualification and/or investment choices, which directly translate into higher procurement expenses for the firm. Our analysis further demonstrates that these inefficiencies can be minimised if the firm commits to an incentive scheme across both periods from the outset.

Chapter 2

The Role of Commitment in First-Price Auctions

Consider a first-price procurement auction in which the buyer lacks the commitment not to renegotiate with the winning seller upon conclusion of the auction. The theoretical prediction for such a setting is stark: only pooling equilibria exist in which bids are both high and uninformative to the buyer. The buyer responds by selecting a winning seller, and making them a take-it-or-leave-it offer. This leads to higher buyer expenses than in a setting with commitment. Of course, this equilibrium requires a high degree of sophistication on the part of the sellers. We therefore take this setting to the laboratory. While we find that sellers bid significantly higher in the no-commitment setting vis-à-vis the setting with commitment, buyer expenses are actually significantly *lower* in the no-commitment setting. This finding holds even when providing buyer and sellers with decision support. It also holds when letting artificial intelligence partake in these auctions. These findings moderate the theoretical prediction, suggesting that, at least at this moment in time, a lack of commitment may actually be beneficial to the buyer.

2.1 Introduction

Commitment is widely regarded by both practitioners and scholars as a critical determinant of success in negotiations and procurement processes. In recent reports of procurement professionals (see e.g. PwC, 2025; Inverto, 2025*a*,*b*), management and stakeholder commitment to the negotiation process—whether bilateral or competitive— is overwhelmingly identified as one of the most important drivers of project outcomes. Organizations that secure such commitment are consistently more likely to achieve favourable results, while the absence of commitment is frequently cited as a primary reason for project underperformance. Despite this broad consensus among practitioners, systematic research on the causal effects of commitment remains limited.

In spite of near-universal agreement about its importance, commitment remains notoriously difficult to test and quantify in organizational research. Commitment, by its nature, is both a process variable and an outcome, shaped by context, incentives, and discretion. Existing studies in negotiation and procurement management often invoke commitment as an assumed input, but few offer credible empirical strategies to isolate its causal impact or measure its effects in a controlled environment. This disconnect between the perceived centrality of commitment in practice and the limitations of empirical research design leaves a significant gap in both theory and practice.

This study addresses this challenge by leveraging a unique research design that allows commitment to be rigorously tested within a process that is both theoretically tractable and widely used in practice: the first-price procurement auction. First-price auctions represent a cornerstone of industrial procurement, with applications spanning public and private sectors and annual volumes measured in the billions (Bajari et al., 2009). Crucially, while economic theory typically presumes that buyers are bound by auction rules and cannot renegotiate after bids are submitted, in reality, many procurement processes permit or even expect some degree of post-auction negotiation (Jap, 2002). This discrepancy between theoretical models and practical realities is an inspiration to study a setting that allows us to examine the true value of commitment.

By exploiting the structure of first-price auctions—where the buyer's ability to commit can be directly manipulated—this study provides a direct experimental test of commitment in a negotiation process that closely mirrors real-world procurement practice. The theoretical framework delivers sharp predictions regarding equilibrium behaviour and outcomes with and without commitment. Specifically, we will compare a procurement auction in which the buyer commits to the rules of the auction to a procurement auction in which, upon seeing the bids, the buyer may make the winning seller a counteroffer.We show that this seemingly trivial change—the *option*, not the *obligation*, to make a counteroffer on the part of the buyer—destroys competition. This leads to higher expected buyer expenses vis-à-vis an auction with commitment. Our results hold independently of both the distribution of sellers' costs, and the numbers of sellers in the auction.

We then take this setting to the laboratory. The experimental approach, using both human subjects and artificial intelligence agents, tests our predictions in controlled yet realistic settings. While we find that sellers bid significantly higher when the buyer lacks commitment, buyer expenses are actually significantly *lower*. This finding holds even when providing buyer and sellers with decision support. Moreover, it also holds when letting artificial intelligence partake in these auctions. Taken together, these findings imply that, at least in the short run, a lack of commitment in first-price procurement auctions may actually be beneficial to the buyer.

In sum, this research makes three main contributions. First, it bridges the gap between theory and practice by empirically examining commitment in a procurement setting of direct managerial importance. Second, it demonstrates that commitment can be isolated and tested using an industry-standard process, enabling a more nuanced understanding of its effects on negotiation outcomes. Third, it provides actionable insights for practitioners on the design and governance of procurement processes where commitment cannot be taken for granted.

The rest of the paper is structured as follows: section 2.2 connects this research with the existing literature; section 2.3 introduces the model, and derives the theoretical predictions; section 2.4 introduces the experiment, while section 2.5 discusses the experimental results. Finally, section 2.6 concludes.

2.2 Literature

The importance of having commitment has long been recognised in the literature on auction theory. However, its presence is often presumed in this body of literature. Consider, for example, the seminal work of Vickrey (1961). While Vickrey assumes the auctioneer is bound to outcomes of the auction, he recognises that the auctioneer may indeed face commitment issues, and discusses this issue explicitly in the case of sealed-bid second-price auctions.

This leaves us with a notable paradox. If commitment is so important and so widely acknowledged, why is it so frequently absent in practice, and why has it received relatively limited systematic attention in empirical research? One explanation lies in the inherent challenges associated with defining, measuring, and manipulating commitment as a variable. Commitment is fundamentally relational and context-dependent: it is shaped by organizational culture, incentive structures, and interpersonal dynamics (Ring and Van de Ven, 1994; Schelling, 1960). Its presence or absence is often revealed only ex-post, and it resists straightforward quantification in surveys or archival datasets (Gulati, 1995). As a result, most empirical research in negotiation and procurement has tended to focus on more readily observable factors—such as contract terms, bidding mechanisms, or price outcomes—rather than on the softer, yet foundational, process variable of commitment (Macaulay, 2018; Poppo and Zenger, 2002).

Furthermore, commitment often operates in the background, as an assumed norm or cultural expectation, rather than as an explicit managerial lever. This may lead to a "taken-for-granted" status in both organizations and research designs, where its effects are assumed rather than directly tested. The limited experimental and field research that does exist has focused primarily on the consequences of contract enforcement or renegotiation, rather than on the broader organizational process of securing and sustaining commitment (Malhotra and Lumineau, 2011).

As a result, while practitioners almost universally recognize the importance of commitment, the difficulties inherent in defining, observing, and experimentally varying this construct have contributed to a relative paucity of direct empirical research—a gap that this study seeks to address. That being said, there is a small body of literature that explicitly studies the theoretical implications of (a lack of) commitment in auctions. However, in this literature, what exactly the word *commitment* entails differs strongly.

In Fugger et al. (2016), the lack of commitment comes in the form of a buyer-determined auction, that is, an auction in which the buyer need not necessarily procure from the lowest bidder. In the procurement setting they consider, sellers have uncertainty about the quality score that the buyer associates with each of the sellers. The authors show that in open-outcry auctions (for example, an English auction), even a small amount of uncertainty can lead to non-competitive prices. The reason is that sellers cannot be sure that a reduction in price leads to a higher probability of winning. However, Fugger et al. find that the existence of the pooling equilibrium depends on both the distribution of the quality components, and the value of the reserve price.

Liu et al. (2019) study the implications of a buyer who cannot commit to the reserve price. They demonstrate that an effective reserve price requires commitment. To be precise, they consider a setting in which a buyer runs a procurement auction that is profit-maximising in the one-shot case, i.e. a standard auction with an optimal reserve price. However, if nobody bids in the auction, the buyer may increase the reserve price. They find that in this case, sellers may like to wait for an auction with a higher reserve price, and thus, an auction without a reserve price is optimal.

Fugger et al. (2019) consider a multi-period procurement setting. In their model, the buyer lacks the commitment not to renegotiate the terms of a two-period contract after the first period. In a first-price auction, they find that whenever the second-period contract is sufficiently high, suppliers will pool on a high bid to conceal their private information, since this information would be used against them in the renegotiation. However, the pooling equilibrium they find is not unique.

How does our paper relate to the papers listed above? While we also make the point that lacking commitment can be detrimental to the auctioneer, we do so in a simpler model. To be precise, we show that a lack of commitment is not only problematic when the auctioneer's preferences are unknown, in the face of a reserve price, or when the procurement process is sequential. Commitment matters *generally*, even in plain vanilla one-shot auctions. As soon as the sellers expect the buyer to make a counteroffer, there can only exist pooling equilibria.

This brings us to the paper with the setting most similar to ours, the one of Shachat and Tan (2015). They consider an English procurement auction in which the buyer reserves the right to bargain further concessions from the winner. Unlike in our paper—or the papers mentioned previously—they show that this lack of commitment is *not* detrimental to the buyer. Sellers still have a weakly dominant strategy of exiting the auction at their costs. We show that this result hinges crucially on the auction format. In particular, it does not carry over to first-price auctions.

2.3 Theory

In this section, we will introduce the theoretical model. Thereafter, we will derive the equilibrium predictions for the commitment and no-commitment settings, respectively.

2.3.1 Model

Consider a first-price procurement auction. A buyer is looking to buy a good from one of n potential sellers, with $n \geq 2$. We assume that all involved parties are risk-neutral profit maximisers.

Each seller's cost c is independently drawn from $C = [\underline{c}, \overline{c}]$, with $\underline{c} \geq 0$, according to the cumulative density function F with a strictly positive probability density function f on its support. A seller's cost is their private information; however, the cost distribution is commonly known. The buyer's commonly-known value for the good is $v = \overline{c}$. The buyer implements a reserve price $r = \overline{c}$.

Each seller submits a bid $b \in [0, r]$. At this point, the two settings we consider differ.

Commitment In the *commitment setting*, the seller with the lowest bid b^* wins, and is paid what they bid, yielding a payoff of $b^* - c$. The losing sellers' payoffs are all 0. The buyer's payoff is given by the difference between the value for the good and the price paid for it, $v - b^*$.

No-Commitment In the *no-commitment setting*, after seeing the bids, the buyer either accepts the lowest bid—yielding the payoffs described above for the commitment setting—or makes the winning seller a counteroffer $o \in [0, b^*]$, which the winning seller may either accept or reject.

If the counteroffer is accepted, o is the price at which trade takes place. The winning seller's payoff is $o - c$; the losing sellers' payoffs are all 0. Again, the buyer's payoff is given by the difference between the value for the good and the price paid for it, $v - o$.

If the counteroffer is rejected, the buyer falls back on the lowest-bidding losing seller and pays them their bid b'. Therefore, this losing seller's payoff is $b' - c$, all other losing sellers' payoffs are 0. The winning seller's payoff is also 0. The buyer's payoff is given by $v - b'$.

2.3.2 Analysis

In this section, we will derive the theoretical predictions of the commitment and no-commitment settings.

Commitment setting

Proposition 1. *The unique equilibrium in the commitment setting is separating: the sellers bid according to the bidding function*

$$\beta(c) = -\frac{1}{G(c)} \int_c^{\bar{c}} x \cdot g(x)\ dx = c + \frac{1}{G(c)} \int_c^{\bar{c}} G(x)\ dx$$

where $G(c) = (1 - F(c))^{n-1}$ *denotes the probability of a seller with costs of* c *winning the auction in a symmetric equilibrium.*

Proof. The equilibrium bidding function is derived analogously to the one of a forward auction (see e.g. Krishna, 2010). □

No-commitment setting

Proposition 2. *The unique equilibrium in the no-commitment setting is a pooling equilibrium. This equilibrium can be characterised as follows:*

(a) All sellers pool on the reserve price r.[1]

(b) The buyer responds to these bids by selecting a winning seller, and offering them a counteroffer of

$$\arg\max_{o}\ (v - o)F(o).$$

When observing a bid $\neq \bar{c}$, the buyer makes a counteroffer of $\underline{c}$ to (one of) the winning seller(s).

(c) The winning seller accepts the counteroffer whenever it is weakly profitable, and rejects the counteroffer otherwise.

Proof. First, note that this indeed constitutes an equilibrium. Sellers pool on the reserve price $r = \bar{c}$. The winning seller is made a counteroffer as defined above. Evidently, it is optimal to accept this counteroffer whenever it is at least as high as their cost, and to reject it otherwise. Note that sellers have no incentive to bid lower, since a deviating bid will be met with a counteroffer of $\underline{c}$. For all possible seller costs, this guarantees a payoff of 0. Similarly, the buyer has no incentive to deviate either. The bids they receive are uninformative of sellers' costs. Thus, they cannot do better than selecting one of the sellers as the winner and making the optimal counteroffer as defined above.

To establish the uniqueness of this equilibrium, we will show that there exists neither a fully separating equilibrium, nor a partial pooling equilibrium. In a fully separating equilibrium, the bidding function is bijective: different seller cost types submit different bids. That is, for any two costs $c_1, c_2 \in C$, the bidding function $\beta(c_1) \neq \beta(c_2)\ \forall c_1 \neq c_2$. Therefore, in such an equilibrium, the buyer is able to infer a seller's cost from their bid. By inverting the bidding function

[1] If $r > \bar{c}$, then a continuum of pooling equilibria exist: sellers mix over a set of bids such that the lowest bid in this set is $\geq \bar{c}$.

and offering the winning seller a counteroffer of $\beta^{-1}(b^*)$, the buyer is able to extract full surplus. The winning seller would then simply receive their cost as a counteroffer, and can profitably deviate by misrepresenting as a type with higher cost. This rules out the existence of a fully separating equilibrium. We adopt the Intuitive Criterion for out-of-equilibrium beliefs. Under this refinement, any deviation from a pooling set is attributed to the highest cost type in that set, which then no longer finds the deviation profitable. Hence, partial pooling and mixed equilibria are eliminated, leaving the unique full-pooling equilibrium.

To see this, consider first a partial pooling equilibrium, i.e. an equilibrium in which multiple types submit the same bid. We will show that such an equilibrium cannot exist. The argument is by contradiction. Let $C_p \subset C$ be the set of types that submit the same bid b_p in a partial pooling equilibrium, and let $\tilde{c}$ denote the type with the highest cost in C_p. Evidently, b_p must be smaller than $\overline{c}$, since otherwise, the sellers in C_p could increase their probability of winning by excluding type $\overline{c}$, which can be achieved through a deviation to $b_p - \epsilon$. If, upon conclusion of the auction, b_p is the lowest bid, the buyer will make a counteroffer at most $\tilde{c}$ to one of the sellers in C_p. However, this means that type $\tilde{c}$ is guaranteed a payoff of 0, and can thus profitably deviate by misrepresenting as a type with higher cost. This contradicts the assumption that type $\tilde{c}$ bids b_p in equilibrium. There therefore exists no partial pooling equilibrium. Note that a similar argument rules out the existence of a mixed strategy equilibrium. For such an equilibrium to exist, a seller would have to be indifferent between all bids they mix over. But now note that, just like in the proof above, the highest cost type that places a certain bid is guaranteed zero profits with that bid, which contradicts the assumption that the seller be indifferent between this bid and all others they mix over.

Finally, let us turn to the buyer's counteroffer when observing a bid lower than $\overline{c}$. Assume to the contrary that the buyer makes a counteroffer greater than $\underline{c}$

to a deviating seller. Some types will then find a deviation away from the pooling equilibrium profitable. However, just like in the argument for the partial pooling equilibrium, observe that the highest cost type in this set of deviating types is guaranteed a payoff of 0; a contradiction to the assumption that this type wants to deviate away from the pooling equilibrium. For any counteroffer in the case of deviation that is smaller than $\underline{c}$, no type wants to deviate from the pooling equilibrium. Now recall that counteroffers are $\in C$. Taken the points made above together, we are left with the unique equilibrium as described in the proposition. □

At this point, we want to stress that proposition 2 holds for any shape of the sellers' cost distribution, provided it is continuous with strictly positive density over its support, and is independent of (i) the particular functional form of that distribution, (ii) the number of sellers in the auction, and (iii) the presence of a fallback option. The first two points follow immediately since the proof does not depend on these features. For the last point, note that in equilibrium, the buyer makes no profit in case the counteroffer is rejected by the winning seller: they fall back on the bid $r = v$. Therefore, the proposition also holds if there is no fallback option, that is, if no trade takes place in case the counteroffer is rejected.

In what follows, we will consider the following metrics: bids, auction prices, buyer expenses, and efficiency. In both treatments, the auction price is given by the winning bid. In the commitment setting, buyer expenses are given by the auction price. By contrast, in the no-commitment setting, buyer expenses are not equal to the auction price whenever the buyer decides to make a counteroffer: if the counteroffer gets accepted, buyer expenses are given by the counteroffer; if the counteroffer gets rejected, buyer expenses are given by the lowest losing bid. Lastly, we will call an outcome efficient only if the seller with the lowest costs supplies the buyer. Observe that in the no-commitment setting, the outcome can also be efficient if the auction was won by an inefficient seller, namely in the case where this seller rejects the buyer's counteroffer and the efficient supplier was the runner-up in the auction. There are multiple theoretical predictions we can derive from proposition 2.

Proposition 3. *In equilibrium, both the expected bid and the expected auction price are higher in the no-commitment setting than in the commitment setting.*

Proof. This follows directly from the separating equilibrium of the commitment setting, in which only the type $\bar{c}$ submits a bid of $\bar{c}$. In the no-commitment setting, *all* types submit a bid of $\bar{c}$. The claim in the proposition follows. □

Proposition 4. *In equilibrium, expected buyer expenses are higher in the no-commitment setting than in the commitment setting.*

Proof. Note that, because $v = \bar{c}$, if the counteroffer is rejected, the buyer makes a profit of zero, since the lowest-bidding losing seller's bid is given by $\bar{c}$. In the no-commitment setting, then, the situation the buyer is in is mathematically equivalent to one in which they are facing only one seller. By contrast, in the commitment setting, the buyer is running an auction with $n \geq 2$ sellers. This being the case, we may leverage the result derived by Bulow and Klemperer (1996): from the buyer's perspective, additional competition is more valuable than running an auction with an optimal reserve price. That is, having an additional seller in an auction is preferable in terms of the expected buyer expenses to *any* mechanism one could run with one less seller. This completes the proof. □

Proposition 5. *In equilibrium, expected efficiency is lower in the no-commitment setting than in the commitment setting.*

Proof. Note that in the commitment setting, the outcome is always efficient, since the seller with the lowest cost always wins. In contrast, inefficient outcomes occur with positive probability in the no-commitment setting, e.g. if the inefficient seller is selected as the winner, and made a counteroffer which they accept. □

The theoretical predictions for the commitment setting and the no-commitment setting differ widely. To determine whether this stark difference holds in practice—even in the face of decision support and artificial intelligence—we take these settings to the laboratory.

2.4 Experiment

In this section, we will begin by describing the design and organisation of the experiment, before deriving concrete hypotheses for both our settings.

2.4.1 Design

In the experiment, one buyer is matched with $n = 2$ sellers. Sellers' private costs are independently and uniformly drawn from the set $C = [0, 100]$. The buyer's value for the good is set to $v = 100$. Bids and counteroffers are both in C. Ties are broken randomly.

We have two treatments: *commitment* and *no-commitment*.

In the commitment treatment, the lowest bidder always wins and is paid their bid. In this treatment, the buyer has a passive role, since they cannot influence their payoff.[2] However, given the active role of the buyer in the no-commitment treatment, the human buyer was kept in order to rule out that any observed differences are driven by sellers facing a computerised buyer.

[2]In order to provide buyers with a task, at the end of every auction, we asked them to submit their best guess of the winning seller's costs. This guess did not influence buyers' payoffs.

In the no-commitment treatment, upon seeing the bids, the buyer either accepts the lowest bid b^*, or makes the winning seller a counteroffer $o < b^*$. The counteroffer is automatically accepted by the winning seller if it is at least as high as their costs, and rejected if it is below their costs.

In all treatments, at the end of each round, all subjects are informed of the two bids that were placed, what the buyer's decision was in the case of the no-commitment treatments—including, if applicable, the counteroffer and whether or not the winning seller accepted it—and what their resulting payoff is.

2.4.2 Organisation

The experiment was conducted on-site in a computer laboratory using the software oTree (Chen et al., 2016). Subjects were recruited from the subject pool of a large western European university, with cash being the only incentive offered. The average payment was 23.01 EUR, which corresponds to about 25.31 USD at the time of the experiment. Subjects were mostly undergraduates and came from a variety of majors.

In order to ensure an understanding of the rules, all participants had to correctly answer the control questions before the experiment could begin. Thereafter, in order to familiarise themselves with the setting, all subjects played 10 rounds of a commitment auction in the role of a seller. They were competing against a computerised seller who bids uniformly between 50 and 100. The computerised buyer bought the good from the lowest seller, and paid them their bid. In the results section below, these ten initial rounds will be called the *pre-treatment*.

Thereafter, the respective treatment was played. Subjects were randomly assigned a role—either buyer or seller—and this role remained constant throughout the 40 rounds played. Subjects were divided into matching groups of six, with each group consisting of four sellers and two buyers. In each round, a buyer was randomly paired with two sellers. For each treatment, 60 subjects were recruited. In total, this gives us 10 independent observations per treatment.

2.4.3 Hypotheses

In the commitment setting, applying proposition 1, risk-neutral sellers are predicted to bid according to the bidding function

$$\beta(c) = \frac{100 + c}{2}. \tag{2.1}$$

In the no-commitment setting, applying proposition 2, sellers are predicted to pool on the highest possible bid of 100. The buyer responds by selecting a winning seller and making them an optimal take-it-or-leave-it offer of 50.

We can now translate propositions 1–2 into concrete hypotheses. The expected bid in the commitment setting is 75, and the expected winning bid (i.e. the auction price) is 66.7. By contrast, both the expected bid and the auction price bid are 100 in the no-commitment setting. This gives us the first hypothesis.

Hypothesis 1. *The bids are* higher *in the no-commitment setting than in the commitment setting.*

Hypothesis 2. *The auction prices are* higher *in the no-commitment setting than in the commitment setting.*

Buyer expenses in the commitment setting are given by the auction price. In the no-commitment setting, however, we must distinguish between two cases: namely the case in which the winning seller accepts the counteroffer, and the case in which they do not. Since the optimal counteroffer is given by 50, each of these two cases occurs with a probability of $\frac{1}{2}$. Recall that in the case of rejection, buyer expenses are given by the losing seller's bid. The expected buyer expenses in the no-commitment setting are therefore $\frac{1}{2} \cdot (50 + 100) = 75$. This gives us our next hypothesis.

Hypothesis 3. *Buyer expenses are higher in the no-commitment setting than in the commitment setting.*

Next, let us turn to efficiency. We will call an outcome efficient if the seller with the lowest cost supplies the buyer. In the commitment setting, the winning seller is always the one with the lowest cost, yielding an efficiency of 100%. By contrast, in the no-commitment setting, half the time, the winning seller is the one with the lowest cost, who accepts the counteroffer with a probability of 75%. Equally half the time, the winning seller is the one with the higher cost, who rejects the counteroffer with a probability of 75%. This yields an expected efficiency of 75% in the no-commitment setting.[3]

Hypothesis 4. *Expected efficiency is lower in the no-commitment setting than in the commitment setting.*

2.5 Results

In this section, we will discuss the experimental results. We obtain independent observations by aggregating on the matching group level. Unless stated otherwise, all comparisons in this section are made using Wilcoxon rank-sum tests on independent observations.

Figure 2.1 below provides a summary of the experiments we ran. To start, each subject played ten rounds of the pre-treatment. Thereafter, subjects moved

[3] These probabilities were calculated using the cumulative density functions of the lowest- and second-lowest of $n = 2$ cost draws from a uniform distribution on $[0, 1]$, respectively, which we will denote by $G_1(c)$ and $G_2(c)$. $G_1(c) = 2c - c^2$ and $G_2(c) = c^2$. Calculating the probability that the seller with lower costs accepts the buyer's counteroffer, as well as the probability that the seller with the higher costs rejects it, we get $G_1(0.5) = 1 - G_2(0.5) = 0.75$.

on to the main treatment. In each of the main treatments, a matching group consisting of six subjects played forty rounds of the experiment. The figure also shows three greyed-out main treatments. These will be described in sections 2.5.1 and 2.5.3, respectively.

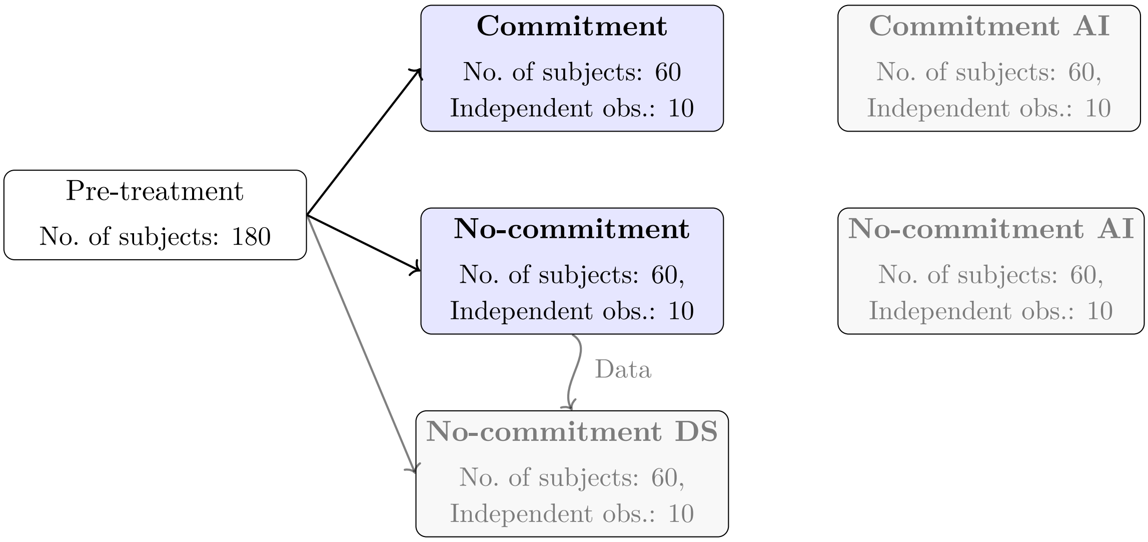

Figure 2.1: Overview of all treatments ran.

Comparing the bids of the initial ten rounds of the pre-treatment across the two main treatments allows us to check for differences in the subject pools. A two-sided Wilcoxon rank-sum test yields no significant difference ($p = 0.39$). This gives us an indication that our randomisation of subjects across treatments was successful.

Table 4.1 below reports the results of the experiment. For each treatment, we report the averages, as well as the theoretical predictions based on the realised cost draws of the experiment, including non-parametric tests. In addition to the averages and the results of two-sided tests against the theoretical predictions, we also report the results of one-sided tests that compare the two treatments across the four metrics. This allows us to directly use table 4.1 to test the hypotheses from above.

Let us begin by comparing bids.

	Observed		Theoretical	
Metric	**Commitment**	**No-commitment**	**Commitment**	**No-commitment**
Bid	69.7 (2.1)	72.0* (3.3)	74.4††† —	100.0††† —
Price	60.6 (3.5)	61.0 (4.6)	64.8††† —	100.0††† —
Expenses	60.6 (3.5)	57.5** (3.4)	64.8††† —	75.0††† —
Efficiency [%]	85.0 (4.5)	79.0** (6.6)	100.0††† —	75.0 —

Table 2.1: Average bid, auction price, buyer expenses, and efficiency across treatments. Standard errors in parentheses. H_0: Identical to commitment treatment, H_1: greater or smaller than commitment treatment; $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.
H_0: Observed = Theoretical, H_1: Observed $\neq$ Theoretical; $^{\dagger}p < 0.1$, $^{\dagger\dagger}p < 0.05$, $^{\dagger\dagger\dagger}p < 0.01$.

Result 1. *In line with hypothesis 1, bids are significantly higher in the no-commitment treatment (*$p = 0.10$*).*

While bids are higher under no-commitment, auction prices are not.

Result 2. *In aggregate, auction prices are not significantly different in the no-commitment treatment (*$p = 0.91$*). Therefore, we find no support for hypothesis 2.*

For the central metric of our analysis—buyer expenses—counter to the theoretical prediction, we observe *lower* buyer expenses under no-commitment.

Result 3. *In contrast to hypothesis 3, buyer expenses are significantly* lower *in the no-commitment treatment (*$p = 0.03$*).*

Let us turn next to our last metric: efficiency.

Result 4. *In line with hypothesis 4, we find that efficiency is significantly lower in the no-commitment treatment relative to the commitment treatment.*

Looking holistically at our results, the higher bids observed under no-commitment suggest that sellers recognise the buyer's lack of commitment should lead them to bid less aggressively, in anticipation of a subsequent counteroffer. Interestingly, auction prices do not differ significantly between the commitment and no-commitment settings. If *all* sellers anticipated that buyers were likely to make a counteroffer to the winning seller—which occurs in more than three-quarters of cases—we would expect bids to be adjusted upward, resulting in significantly higher auction prices. The absence of such a difference implies that winning sellers are disproportionately those who are less sensitive to the distinction between commitment and no-commitment, and therefore fail to adjust their bids adequately to account for the possibility of a counteroffer. While auction prices remain similar across treatments, buyers' counteroffers are accepted by the winning seller in roughly two-thirds of the cases. This enables buyers to achieve substantially lower total expenses under the no-commitment condition.

A median split of buyers on their payoff in the main treatment allows us to demonstrate this. Buyers in the top half of the payoff distribution (henceforth *successful buyers*) use their option to make counteroffers more frequently—78% vs 75% of the time—however, this difference is not statistically significant ($p = 0.36$). Nevertheless, the difference between the auction price and the counteroffer—henceforth, the *margin*—is significantly higher for buyers in the top half of the payoff distribution: 15.3 versus vs 11.9 ($p < 0.01$). In other words, buyers who make more frequent and aggressive counteroffers attain higher payoffs in the experiment.

Looking more closely at buyers' behaviour when making counteroffers: the margin should be positively associated with the auction price and negatively associated with the distance between winning and losing bid. For the former, this is because the higher the auction price, the more the buyer stands to gain by making a counteroffer. For the latter, this is because the further the winning and losing bid are apart, the worse a rejected counteroffer is for the buyer, which should lead to lower margins. Table 2.2 provides the results of a random-effects panel regression of the buyer's margin on the auction price and the distance between winning and losing bid. As predicted, looking at model (1) we find a positive coefficient for the auction price and a negative coefficient for the distance. Table 2.2 also includes models with a dummy variable for successful buyers.[4] As demonstrated by models (2) and (3), successful buyers make more aggressive counteroffers and are less influenced by the auction price, but more influenced by the distance between winning and losing bid.

	Dependent variable:		
		Margin	
	(1)	(2)	(3)
Constant	15.774***	13.226***	8.448*
	(3.411)	(3.831)	(5.090)
AuctionPrice	0.059	0.060*	0.126**
	(0.036)	(0.036)	(0.054)
DistanceWinningLosingBid	−0.071*	−0.072*	−0.046
	(0.041)	(0.040)	(0.065)
SuccessfulBuyer		4.996	13.775**
		(3.379)	(6.826)
AuctionPrice × SuccessfulBuyer			−0.123*
			(0.072)
SuccessfulBuyer × DistanceWinningLosingBid			−0.054
			(0.083)
Observations	613	613	613
R^2	0.053	0.057	0.064
Adjusted R^2	0.050	0.052	0.056
F Statistic	32.318***	34.664***	39.277***
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 2.2: Random effects model for buyer margin in the no-commitment treatment.

Taken together, the results above demonstrate that while seller bidding behaviour and buyer counteroffer behaviour is directionally in line with the theoretical predictions, sellers may benefit from a better understanding of what counteroffers to expect for a given bid, while buyers may benefit from information on the which seller costs are associated with a given winning bid. As a result, there may well be room for improvement for both. We therefore introduce another treatment

[4]Note that the correlation between a buyer's counteroffer margins and the dummy for whether or not they were successful is −0.09. As such, we can rule out that these two variables are simultaneously determined by a buyer's payoff.

that includes decision support.

2.5.1 Decision support

In this section, we introduce a new treatment—no-commitment with decision support, henceforth *no-commitment DS*—to study the effects of decision support in the no-commitment setting. The no-commitment DS treatment is identical to *no-commitment*, except that subjects are additionally given decision support. This treatment was run after all no-commitment sessions had been concluded. We use the empirical data collected from the no-commitment treatment to give sellers and buyers decision support. To be precise, in each round, sellers are informed which bid would have maximised their expected profit in the no-commitment treatment without DS if (1) the buyer accepts the lowest bid, and if (2) the buyer makes the counteroffer which would have maximised their expected profit in the no-commitment treatment without DS. Additionally, sellers were informed of the range of counteroffers the buyer will see if they win with their currently entered bid.[5] Similarly, buyers are informed what the profit-maximising counteroffer would have been against the two observed bids in the no-commitment treatment without DS.[6]

Recall that we observed no difference in auction prices between commitment and no-commitment. Decision support was intended to make salient to sellers that the buyer is likely to make them a counteroffer. It also provides information on the counteroffer recommendation the buyer will see, which allows sellers to better account for the counteroffer when submitting their bids. We therefore hypothesise

[5] For example, assume the seller's currently entered bid is 75. Winning with this bid means that the other seller bids $\in [75, 100]$. For each of these possible bids, we compute the buyer's profit-maximising counteroffer. We then report the minimum and maximum of these possible counteroffers to the seller. In this concrete example, the minimum counteroffer is 25, and the maximum counteroffer is 66.

[6] If the buyer sees a winning bid that was not observed in the no-commitment treatment, the suggestion is based on the next-highest observed winning bid.

that providing decision support will move us closer to the theoretical predictions. That is, we predict that DS leads to higher bids and auction prices. The effect on buyer expenses is less clear. On the one hand, higher auction prices may drive buyer expenses up. On the other hand, buyer's decision support may result in more aggressive counteroffers, driving buyer expenses down. In aggregate, the effect of decision support on buyer expenses—if any—is unclear.

This gives us three hypotheses.

Hypothesis 5. *The bids are higher in the no-commitment DS setting than in the no-commitment setting.*

Hypothesis 6. *The auction prices are higher in the no-commitment DS setting than in the no-commitment setting.*

Hypothesis 7. *Counteroffers are more aggressive in the no-commitment DS setting than in the no-commitment setting.*

The experimental setup is just as described above for the other treatments: subjects in the no-commitment DS play ten rounds of the pre-treatment, before playing 40 rounds of the main experiment in a matching group consisting of six subjects. Just like in the other two treatments, we collect ten independent observations.

Comparing the bids of the initial ten rounds of the pre-treatment across all three main treatments allows us to check for differences in the subject pools. A Kruskal-Wallis rank sum test yields no significant difference between the three groups ($p = 0.58$). Similarly, when making pair-wise comparisons, two-sided Wilcoxon rank-sum tests also yield no significant differences ($p = 0.39$ for the comparison between commitment and no-commitment, $p = 0.52$ for the comparison between commitment and no-commitment DS, and $p = 0.57$ for the comparison between no-commitment and no-commitment DS).

Table 2.3 below reports the results of the no-commitment DS experiment. For ease of comparison, we have included the results of the no-commitment treatment. Moreover, in italics, we report the results of *no-commitment SIM* and *SIM2*—that is, a simulation of the no-commitment setting assuming both the sellers and the buyer or only the buyer had adhered to the DS recommendation, respectively.

Just like before, in addition to the averages and the results of two-sided tests against the theoretical predictions, we also report the results of one-sided tests that compare *no-commitment DS* with *no-commitment* across the four metrics.

	Observed				*Theoretical*	
	No-co DS	**No-co**	***No-co SIM1***	***No-co SIM2***	**No-co DS**	**No-co**
Bid	75.0^{**} (3.3)	72.0 (3.3)	*86.2**** —	*75.0*** *(3.3)*	$100.0^{\dagger\dagger\dagger}$ —	$100.0^{\dagger\dagger\dagger}$ —
Price	64.2^{*} (3.9)	61.0 (4.6)	*81.9**** —	*64.2** *(3.9)*	$100.0^{\dagger\dagger\dagger}$ —	$100.0^{\dagger\dagger\dagger}$ —
Expenses	57.0 (4.2)	57.5 (3.4)	*68.2**** —	*53.9*** *(3.5)*	$75.0^{\dagger\dagger\dagger}$ —	$75.0^{\dagger\dagger\dagger}$ —
Efficiency [%]	79.0 (3.2)	79.0 (6.6)	*70.0**** —	*78.5* *(6.5)*	$75.0^{\dagger\dagger}$ —	75.0 —

Table 2.3: Average bid, auction price, buyer expenses, and efficiency across treatments. Standard errors in parentheses. Italics indicate a simulation of the no-commitment setting assuming both the sellers and the buyer (*No-co SIM1*) or only the buyer (*No-co SIM2*) had adhered to the DS recommendation.
H_0: Identical to no-commitment treatment, H_1: greater or smaller than no-commitment treatment; $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.
H_0: Observed = Theoretical, H_1: Observed $\neq$ Theoretical; $^{\dagger}p < 0.1$, $^{\dagger\dagger}p < 0.05$, $^{\dagger\dagger\dagger}p < 0.01$

Result 5. *In line with hypothesis 5, bids are higher in the presence of DS ($p = 0.04$).*

Result 6. *In line with hypothesis 6, auction prices are higher in the presence of DS ($p = 0.05$).*

As hypothesised, in the presence of DS, we find significantly higher bids and auction prices compared to the no-commitment setting. However, a two-sided test yields no significant difference in buyer expenses ($p = 0.74$).

Let us look at this result in more detail. First, DS makes the possibility of receiving a counteroffer more salient to sellers, who, as predicted, respond by bidding higher. This leads to significantly higher auction prices. On the buyer's side, we find that DS causes them to make counteroffers significantly more often ($p = 0.09$): counteroffers were made 77% and 83% of the time in the no-commitment and no-commitment DS settings, respectively. Looking at the acceptance rate of counteroffers, while DS leads to buyers making counteroffers that winning sellers accept more often—72% of the time compared with 67% of the time—a two-sided test yields no significant difference ($p = 0.35$).

Table 2.4 reports the results of a random effects regression of the buyer's counteroffer on the auction price, the distance between winning and losing bid, the recommended counteroffer of the DS, as well as whether or not a buyer was successful, that is, fell in the top half of the payoff distribution.[7] As seen in model (1), buyers are in line with theoretical predictions: counteroffers are less aggressive for higher auction prices and larger distances between winning and losing bid. Moreover, counteroffers are positively associated with the DS recommendation. Once more, just like in the no-commitment setting, looking at model (2), successful buyers are those who make more aggressive counteroffers. Importantly, model (3) demonstrates that successful buyers are those who put more weight on the DS recommendation.

Why do we not see a significant difference in buyer expenses? The answer can be found by comparing buyers' average difference between auction price and counteroffer under no-commitment with and without DS. Under no-commitment DS, this counteroffer margin is 16.6 on average; without DS, on average 13.6.

Result 7. *In line with hypothesis 7, buyer counteroffer margins are higher in the presence of DS* ($p = 0.099$).

In other words, as exemplified in figure 2.2, while bids and auction prices are higher with DS, more frequent and aggressive counteroffers—which also get accepted more often—drive the buyer expenses back down. In aggregate, these

[7]Note that the correlation between a buyer's counteroffers and the dummy for whether or not they were successful is 0.11. As such, we can rule out that these two variables are simultaneously determined by a buyer's payoff.

	Dependent variable:		
	Counteroffer		
	(1)	(2)	(3)
Constant	4.494	8.080**	15.370***
	(3.805)	(3.914)	(3.993)
AuctionPrice	0.336***	0.331***	0.307***
	(0.047)	(0.048)	(0.046)
DistanceWinningLosingBid	0.160***	0.158***	0.147***
	(0.048)	(0.048)	(0.046)
CounterofferDS	0.426***	0.425***	0.305***
	(0.027)	(0.028)	(0.032)
SuccessfulBuyer		−6.422***	−18.250***
		(2.420)	(3.067)
CounterofferDS × SuccessfulBuyer			0.272***
			(0.042)
Observations	660	660	660
R^2	0.565	0.567	0.594
Adjusted R^2	0.563	0.564	0.591
F Statistic	835.882***	834.213***	932.972***
Note:			*p<0.1; **p<0.05; ***p<0.01

Table 2.4: Random effects model for buyer counteroffers in the no-commitment DS treatment.

two opposing effects cancel each other out. Another way to see this is to compare payoffs with and without DS. Using a two-sided test, we find no statistically significant difference, neither for sellers, nor for buyers ($p = 0.74$ and $p = 0.85$, respectively). Therefore, it seems that both sellers and buyers profit from DS to a similar degree.

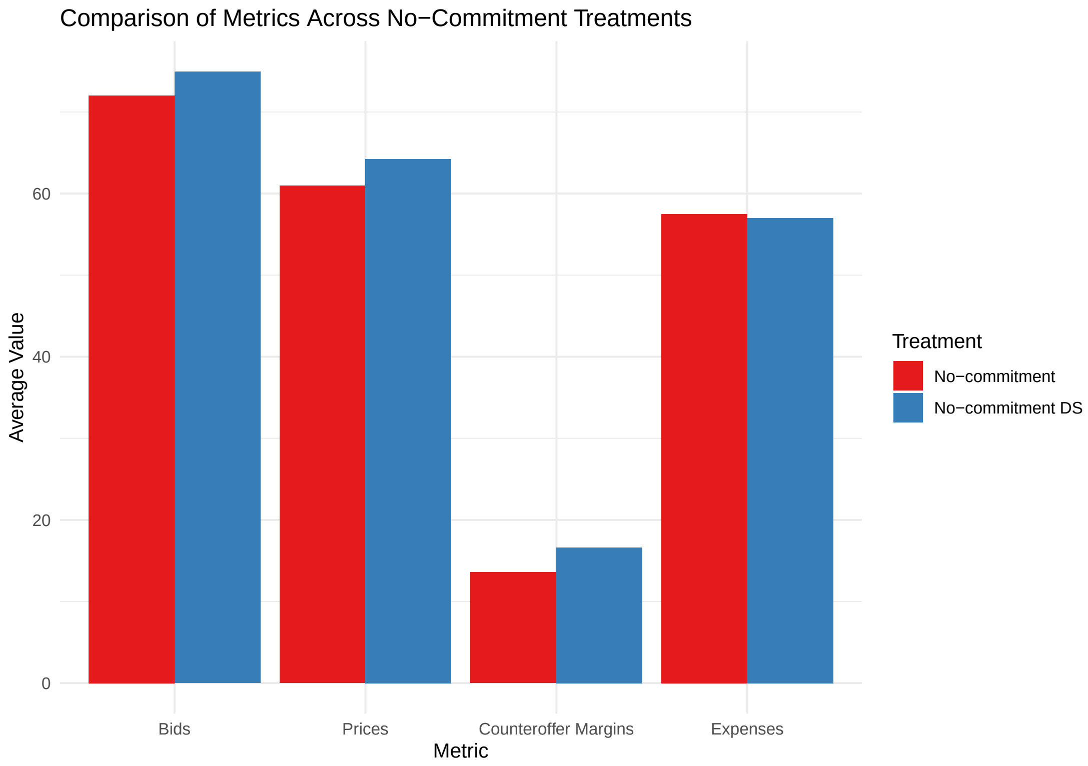

Figure 2.2: Comparison of bids, prices, counteroffer margins, and buyer expenses across the no-commitment treatments.

Figure 2.3 breaks the auction prices down into the winning seller's costs and their margin, as well as into buyer expenses and their counteroffer margin. We display cases in which the buyer's counteroffer was accepted separately. Comparing the counteroffer margin to the winning seller's cost margin, the figure demonstrates that there is room for improvement on the buyer's side when making counteroffers.

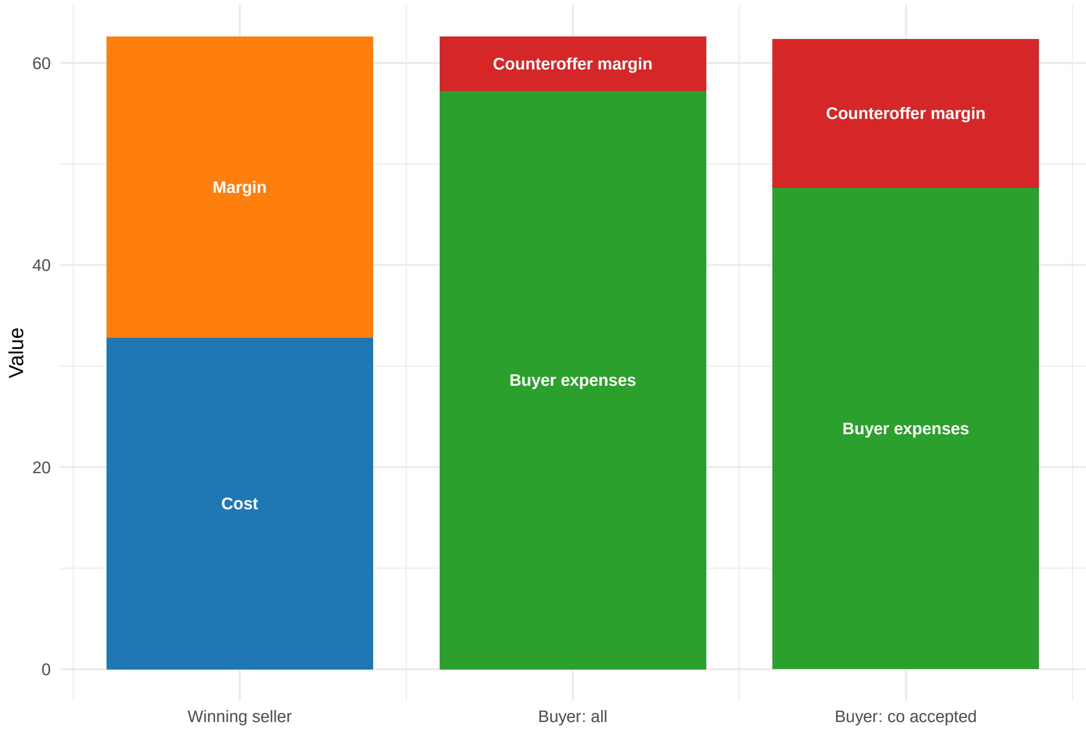

Figure 2.3: Breakdown of auction prices into average costs, margins, and buyer expenses across the no-commitment treatments. The height of a bar corresponds to the respective auction price.

In the following section, we will use the data we collected to run simulations that allow us to study what would have happened had subjects stuck to DS recommendations.

2.5.2 Simulations

In the previous section, we saw that DS leads to higher bids and auction prices. However, the central metric, buyer expenses, remains unchanged. The data allows us to run simulations to reveal what would have happened had (1) buyers and sellers, and (2) only buyers followed DS advice.

What if sellers and buyers had both followed DS advice? To be precise, imagine sellers had always placed the bid that would have maximised their expected profit in the no-commitment treatment, assuming the buyer will respond with a counteroffer that maximises their expected profit. What would the outcomes be?

Referencing *No-co SIM1* in table 2.3 above, we find that—relative to the no-commitment with decision support treatment—bids, auction prices, and buyer expenses are all higher, while efficiency is lower. Note that these differences are all statistically significant (all $p < 0.01$). In short: had both sellers and buyers followed the DS recommendations, sellers would be better off, buyers would be worse off, and from an efficiency standpoint, the outcomes are less desirable too.

What if buyers had followed DS advice? In 82.5% of rounds, buyers made a counteroffer in the no-commitment with DS treatment. However, in only approximately a fifth of these rounds did they follow the advice of the decision support system regarding the counteroffer.[8] This begs the question: would buyers have been better off following the advice given by the decision support system. A

[8]In 60% of cases, they departed from the recommendation by more than 5 ECU.

numerical simulation of the counterfactual allows us to answer this question.

In the no-commitment with DS treatment, buyer expenses were, on average, 57.0 ECU. As demonstrated by *No-co SIM2*, had buyers followed the advice regarding counteroffers, buyer expenses would have averaged 53.9 ECU. This is significantly lower ($p = 0.095$). Put simply, buyers would have attained significantly higher profits had they followed the DS advice.

Summing up the results thus far, we find a lack of commitment to be beneficial to buyers when faced with human subjects. This finding holds even when subjects are equipped with decision support. However, we also find that buyers' lack of commitment could be detrimental if decision support is followed more closely. Therefore, as a buyer, an important question to ask is how relevant commitment is when faced not with human subjects, but with AI. To answer this question, we run another experiment.

2.5.3 Artificial intelligence

With the rising importance of artificial intelligence (AI) in the procurement context (see e.g. Cui et al., 2022), AI-supported—or even AI-only auctions—are no longer a thing of the future. A natural question this raises is how important commitment is in the face of AI subjects. We expect commitment to matter more if we have AI instead of humans subjects in the auction. This is because we expect the difference between an auction with and without commitment to be clearer to AI than to human subjects. Therefore, just like in the case of DS, we expect the AI treatments to be closer to the theoretical predictions. Focusing on our central metric of buyer expenses, this gives us two hypotheses.

Hypothesis 8. *Buyer expenses in the commitment AI treatment are higher vis-à-vis the treatment with human subjects.*

Hypothesis 9. *Buyer expenses in the no-commitment AI treatment are higher vis-à-vis the treatment with human subjects.*

Moreover, we expect that the hypotheses 1 to 3 also apply to AI subjects. We therefore expect higher bids, auction prices, and buyer expenses in the no-commitment setting vis-à-vis the commitment setting.

Setup

To test the hypotheses above, we let OpenAI's GPT-4o play the commitment and no-commitment treatments.[9] We went through exactly the same steps as in the experiment with human subjects. To mirror the matching group structure of the experiment with human subjects, a group consisted of six separate AI threads: two buyers, and four sellers. In each round, one buyer was randomly matched with two sellers. Each AI thread was given the instructions of the respective treatment, and was informed that its goal is to maximise its total profit over all 40 rounds of the experiment. It was asked to devise a strategy to do just that. Moreover, each AI thread was provided with the control questions and their answers before the experiment began. Thereafter, each AI thread was informed of their role, and then the experiment proceeded precisely as described above for human subjects.

To be precise, in each round, sellers were informed of their private costs before submitting a bid. Then—in the no-commitment treatment—the buyer was informed of the bids before deciding whether or not to make a counteroffer. Thereafter, everyone was informed of all the decisions that were made in that round—bids, and if applicable, whether the buyer accepted the lowest bid, what counteroffer the buyer made to the winning seller if they did not accept the lowest bid, as well as whether or not this counteroffer was accepted by the winning seller. Lastly, each thread was informed what their resulting profit is.

[9] We used OpenAI's GPT-4o engine with the temperature set to 0.5.

Note that in each round, the respective AI thread includes the entire history of the game thus far, i.e. it includes the instructions, control questions, their strategy to maximise profits, as well as their decisions in and the results of previous rounds. For exemplary purposes, sections 2.B and 2.C contain the full thread of an AI subject for the first three rounds in the commitment and no-commitment setting, respectively.

Just like in the case of human subjects, we let AI independently play each treatment ten times. This gives us the same number of independent observations as in the experiment with human subjects.

Results

Let us turn to the results of the experiment. In the commitment treatment, 79.8% of the time, AI sellers bid according to the the risk-neutral profit-maximising bidding function as described in eq. (2.1) above. A less aggressive bid is submitted only 3.2% of the time, which means that 17.1% of the time, AI sellers are bidding more aggressively. As reported in table 2.5, this yields, on average, a bid of 71.2, buyer expenses of 59.5, and an efficiency of 94.4%.

	Observed				*Theoretical*			
	Co AI	**No-co AI**	**Co**	**No-co**	**Co AI**	**No-co AI**	**Co**	**No-co**
Bid	71.2 (4.0)	67.7^{**} (4.4)	69.7 (2.1)	$72.0^{‡‡}$ (3.3)	$74.4^{†††}$ —	$100.0^{†††}$ —	$74.4^{†††}$ —	$100.0^{†††}$ —
Price	59.5 (6.2)	54.6^{**} (6.4)	60.6 (3.5)	$61.0^{‡‡}$ (4.6)	$64.8^{††}$ —	$100.0^{†††}$ —	$64.8^{†††}$ —	$100.0^{†††}$ —
Expenses	59.5 (6.2)	53.3^{**} (5.0)	60.6 (3.5)	$57.5^{‡‡}$ (3.4)	$64.8^{††}$ —	$75.0^{†††}$ —	$64.8^{†††}$ —	$75.0^{†††}$ —
Efficiency [%]	94.4 (6.2)	85.4^{**} (10.1)	$85.0^{‡‡‡}$ (4.5)	79.0 (6.6)	$100.0^{†††}$ —	$75.0^{†††}$ —	$100.0^{†††}$ —	75.0 —

Table 2.5: Average bid, auction price, buyer expenses, and efficiency across AI treatments. Standard errors in parentheses. Human treatments listed for comparison.
H_0: Identical to no-commitment treatment, H_1: greater or smaller than no-commitment treatment; $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.
H_0: Observed = Theoretical, H_1: Observed $\neq$ Theoretical; $^{†}p < 0.1$, $^{††}p < 0.05$, $^{†††}p < 0.01$.
H_0: Identical to human treatment, H_1: greater or smaller than human treatment; $^{‡}p < 0.1$, $^{‡‡}p < 0.05$, $^{‡‡‡}p < 0.01$.

Let us start with the commitment treatment. The majority of the results are statistically indistinguishable between human and AI subjects. Two-sided Wilcoxon rank-sum tests yield no significant difference for bids ($p = 0.43$) or buyer expenses ($p = 0.79$). However, we do observe higher efficiency in the commitment setting with AI subjects ($p < 0.01$).

Result 8. *In contrast to hypothesis 8, we do not observe higher buyer expenses in the commitment AI treatment.*

Human subjects are known to bid more aggressively than the risk-neutral profit-maximising bidding function in first-price auctions with commitment.[10] It seems that the same is true for a share of AI subjects. Taken together, while the vast majority of AI subjects bid according to the risk-neutral profit-maximising strategy, the presence of the small share of aggressive sellers means that, in aggregate, we observe lower buyer expenses than theoretically predicted ($p = 0.04$).

In the no-commitment treatment, bidding behaviour is more heterogeneous. A large share of AI sellers are impervious to the difference between an auction with and an auction without commitment: 27% of all bids are derived using the risk-neutral profit-maximising bidding function of the commitment setting. More surprisingly, only 18% of bids are above this benchmark. This means that the remaining 55% of the time, AI sellers are bidding *more aggressively* than the risk-neutral profit-maximising benchmark of the commitment setting. In aggregate, this means that bids, auction prices, and buyer expenses are significantly lower with AI subjects than with human subjects ($p = 0.01$, $p = 0.01$, and $p = 0.02$, respectively).

Result 9. *In contrast to hypothesis 9, we find* lower *buyer expenses in the no-commitment treatment with AI subjects.*

When it comes to the comparison between the two treatments with AI subjects, we expected AI sellers in the no-commitment setting to adjust their bids upwards

[10] See e.g the discussion on overbidding in first-price auctions started by Cox et al. (1982).

from eq. (2.1) in order to account for the buyer's possibility to make counteroffers. However, we find the opposite.

Result 10. *Compared to the commitment setting, bids and auction prices are significantly lower in the no-commitment setting (both* $p = 0.03$*).*

The same is true for buyer expenses.

Result 11. *Compared to the commitment setting, buyer expenses are significantly lower in the no-commitment setting (*$p = 0.01$*).*

Experiments with AI subjects give us the opportunity to directly ask subjects what the reasoning behind their decisions is. We did just that: we asked AI subjects to describe their bidding strategy before the actual experiment began. The analysis shows that the large share of aggressive bids is driven by flawed reasoning on the part of the AI sellers. The assumption of some AI sellers is that the buyer will either make conservative counteroffers that get accepted by the winning seller, or simply accept the winning bid, and thus, it is beneficial to bid aggressively to win the auction more often. While this is true, bidding *more* aggressively than in the commitment setting—which subjects did more than half the time—is not optimal.

To see a concrete example of this, consider the AI seller in section 2.C. We asked subjects in the no-commitment treatment to first describe their bidding function in the commitment setting, and then to modify this bidding function to account for the buyer's lack of commitment. The AI subject in question correctly determined the risk-neutral profit-maximising function for the commitment setting. However, when asked to determine the bidding function for the no-commitment setting, they set their bidding function to $\beta(c) = c + 5$, in order to ensure that the markup on costs is "small enough to remain competitive but large enough to allow room for a profitable counteroffer". Even if the buyer always accepts the outcome of the auction, the bidding function above cannot be optimal, since the profit is bounded by this small markup on costs.

Note that this strategy of essentially bidding costs is not unique to this example. Of the 40 sellers in the no-commitment treatment, across all rounds, seven sellers added on average less than 5 ECU to their costs when submitting their bids. The presence of these bidders translates directly into lower buyer expenses.

Note that we can rule that subjects' flawed reasoning comes from a lack of understanding of the no-commitment setting. This is because the strategy to maximise total profits was derived at a point in time in which the control questions and their answers had already been given to the AI. Therefore, the setting—especially the fact that it is possible to make a profit as the losing seller, namely in the case where the winning seller rejects the buyer's counteroffer—is known to the AI.

What about the AI buyers? Their counteroffer strategies are similar: for low winning bids, simply accept; otherwise, derive the counteroffer by subtracting e.g. 10 from the winning bid. It is interesting to note that this strategy was predicted by the small share of sellers who, in anticipation of a counteroffer, adjusted their bids upward from the risk-neutral profit-maximising bidding function of the commitment setting by a constant.

To conclude then, for our central metric—buyer expenses—it does not matter whether humans or AI are participating in an auction with commitment. However, in the no-commitment setting, the buyer obtains more favourable results if they are facing AI sellers. As a result—at least at this moment in time—buyer expenses are lower in auctions without commitment. We have shown that this result is robust to both subjects equipped with decision support as well as AI agents.

2.6 Conclusion

In this paper, we study the role of commitment in first-price auctions. The question we seek to answer is: how important is it for a buyer to commit not to renegotiate with the winning seller? From a theoretical standpoint, we show that commitment is indispensable. Without it, competition breaks down, leading to higher buyer expenses than in a setting with commitment. However, the equilibrium in the no-commitment setting requires a high level of sophistication on the part of sellers. Therefore, we take this setting to the laboratory.

We show experimentally that while sellers bid higher in an auction without commitment, auction prices remain the same. This finding is in line with winning sellers who are not sensitive to the difference between an auction with and without commitment. That is, winning sellers essentially bid as if they were participating in an auction with commitment. As a result of their ability to make counteroffers, buyers obtain significantly lower expenses under no-commitment.

We extend our experimental analysis by introducing decision support for both sellers and buyers that aims at making the implications of a lack of commitment clearer. To be precise, sellers are informed what counteroffers to expect for a given bid, while buyers are recommended a counteroffer given the auction outcomes. While decision support leads to higher auction prices, it also leads to more aggressive counteroffers on the buyer's side. As a result, in aggregate, total buyer expenses remain unchanged. However, we show using simulations that the full potential of the decision support was not exploited: both sellers and buyers could have done better by adhering more closely to the decision support recommendations.

Due to the rising importance of artificial intelligence (AI) in the procurement context, we extend our experimental analysis even further by allowing current state-of-the-art artificial intelligence agents to take the role of seller and buyer. We find no difference between AI and human subjects in the commitment setting. However, under no-commitment, flawed reasoning on the seller's part leads to significantly lower buyer expenses than with human subjects.

Our experimental results demonstrate that buyer commitment is not as important as theoretically predicted: sellers—even those equipped with decision support, or state-of-the-art AI agents—simply do not adjust their bids upwards enough to account for the buyer's lack of commitment. As a result, buyers' ability to make counteroffers allows them to obtain significantly lower expenses under no-commitment.

The main managerial implication for real-world procurement auctions is that, as a buyer, running an auction without commitment is beneficial. Surprisingly, this finding proves robust both to sellers with decision support as well as AI agents, with the buyer actually preferring to face AI instead of human sellers. Whether this result is one that the advance of artificial intelligence may overturn in the future is an interesting question to ponder for future research.

Appendix of Chapter 2

2.A Experiment: screenshots

2.A.1 Pre-treatment

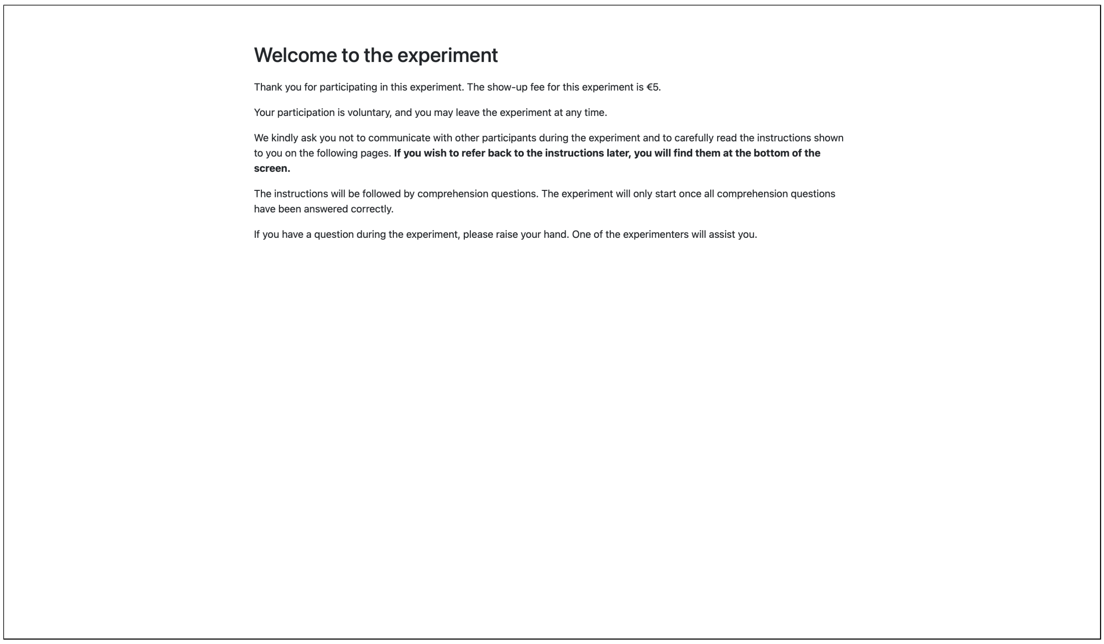

Welcome to the experiment

Thank you for participating in this experiment. The show-up fee for this experiment is €5.

Your participation is voluntary, and you may leave the experiment at any time.

We kindly ask you not to communicate with other participants during the experiment and to carefully read the instructions shown to you on the following pages. **If you wish to refer back to the instructions later, you will find them at the bottom of the screen.**

The instructions will be followed by comprehension questions. The experiment will only start once all comprehension questions have been answered correctly.

If you have a question during the experiment, please raise your hand. One of the experimenters will assist you.

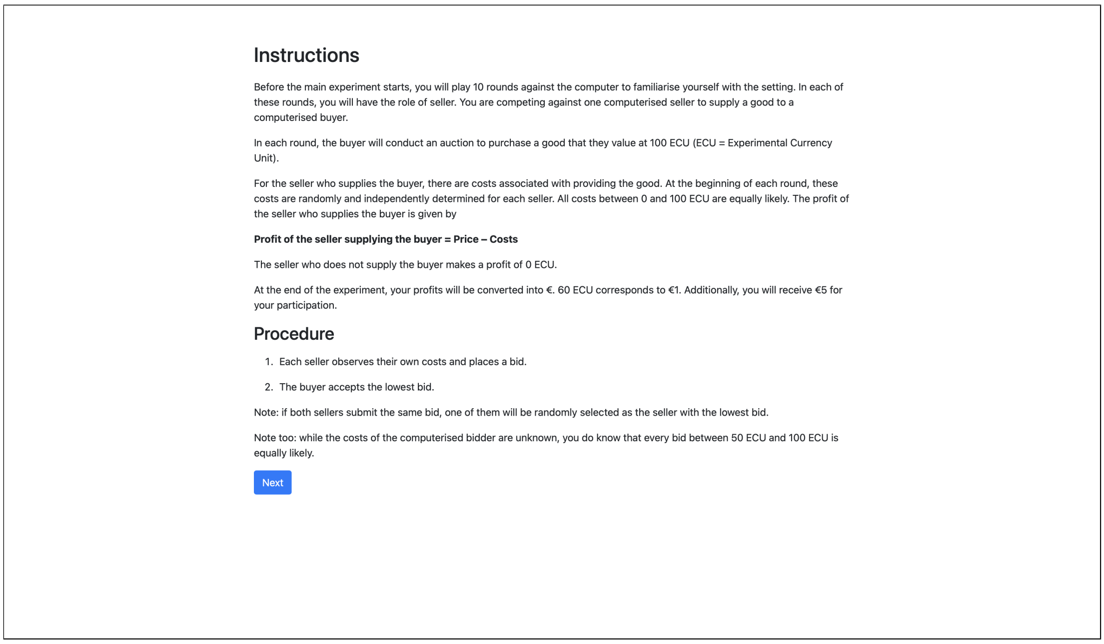

Instructions

Before the main experiment starts, you will play 10 rounds against the computer to familiarise yourself with the setting. In each of these rounds, you will have the role of seller. You are competing against one computerised seller to supply a good to a computerised buyer.

In each round, the buyer will conduct an auction to purchase a good that they value at 100 ECU (ECU = Experimental Currency Unit).

For the seller who supplies the buyer, there are costs associated with providing the good. At the beginning of each round, these costs are randomly and independently determined for each seller. All costs between 0 and 100 ECU are equally likely. The profit of the seller who supplies the buyer is given by

Profit of the seller supplying the buyer = Price – Costs

The seller who does not supply the buyer makes a profit of 0 ECU.

At the end of the experiment, your profits will be converted into €. 60 ECU corresponds to €1. Additionally, you will receive €5 for your participation.

Procedure

1. Each seller observes their own costs and places a bid.
2. The buyer accepts the lowest bid.

Note: if both sellers submit the same bid, one of them will be randomly selected as the seller with the lowest bid.

Note too: while the costs of the computerised bidder are unknown, you do know that every bid between 50 ECU and 100 ECU is equally likely.

Next

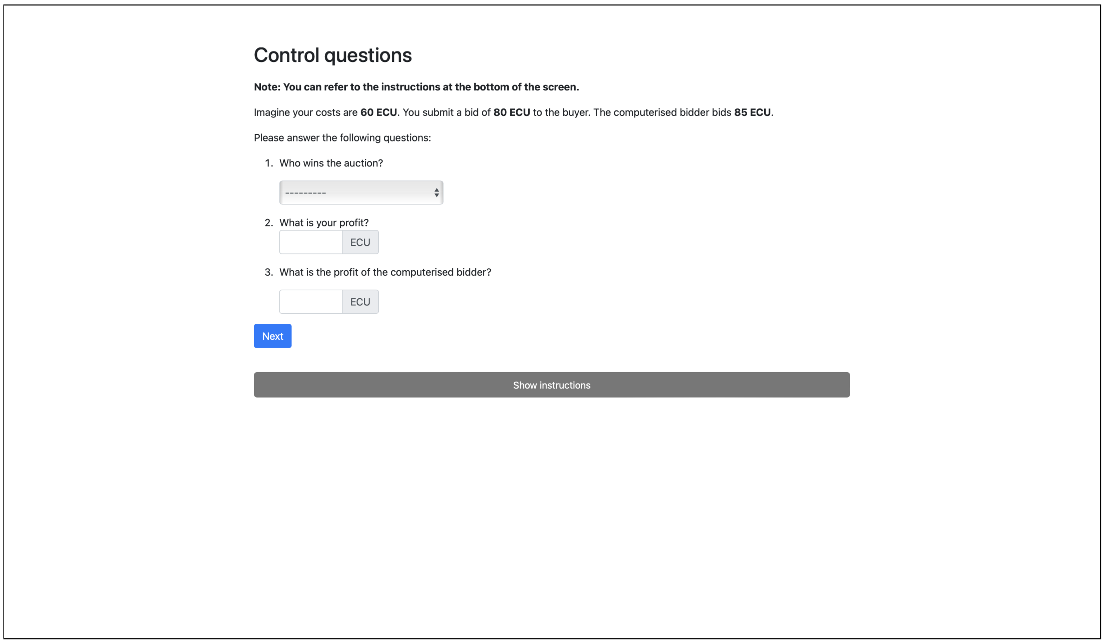
Control questions
Note: You can refer to the instructions at the bottom of the screen.
Imagine your costs are 60 ECU. You submit a bid of 80 ECU to the buyer. The computerised bidder bids 85 ECU.
Please answer the following questions:
1. Who wins the auction?
----
2. What is your profit?
ECU
3. What is the profit of the computerised bidder?
ECU
Next
Show instructions

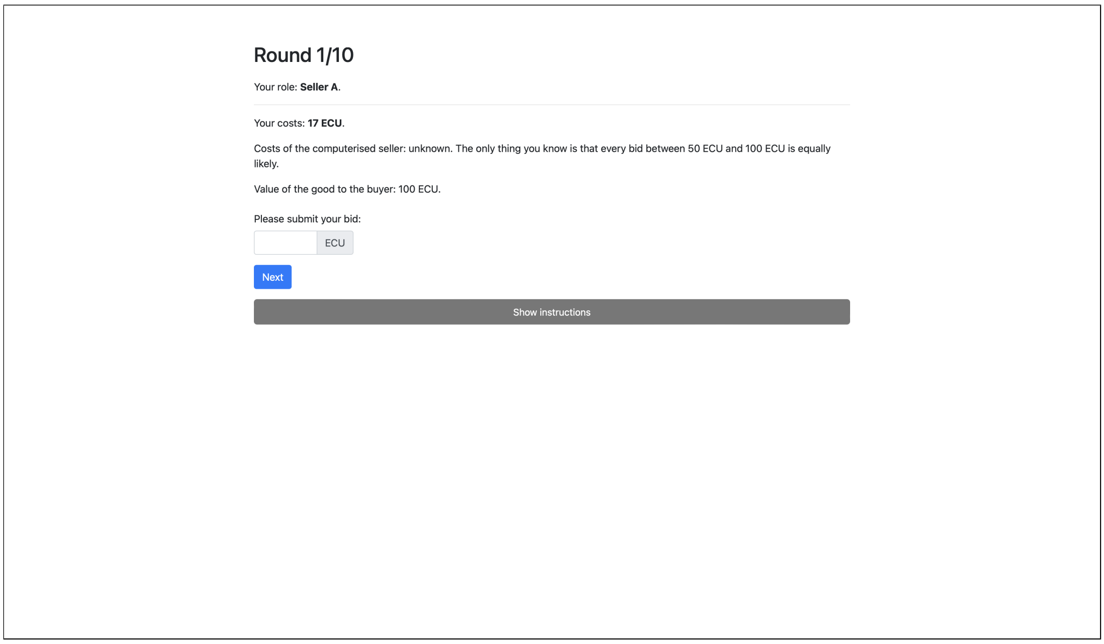
Round 1/10
Your role: Seller A.
Your costs: 17 ECU.
Costs of the computerised seller: unknown. The only thing you know is that every bid between 50 ECU and 100 ECU is equally likely.
Value of the good to the buyer: 100 ECU.
Please submit your bid:
ECU
Next
Show instructions

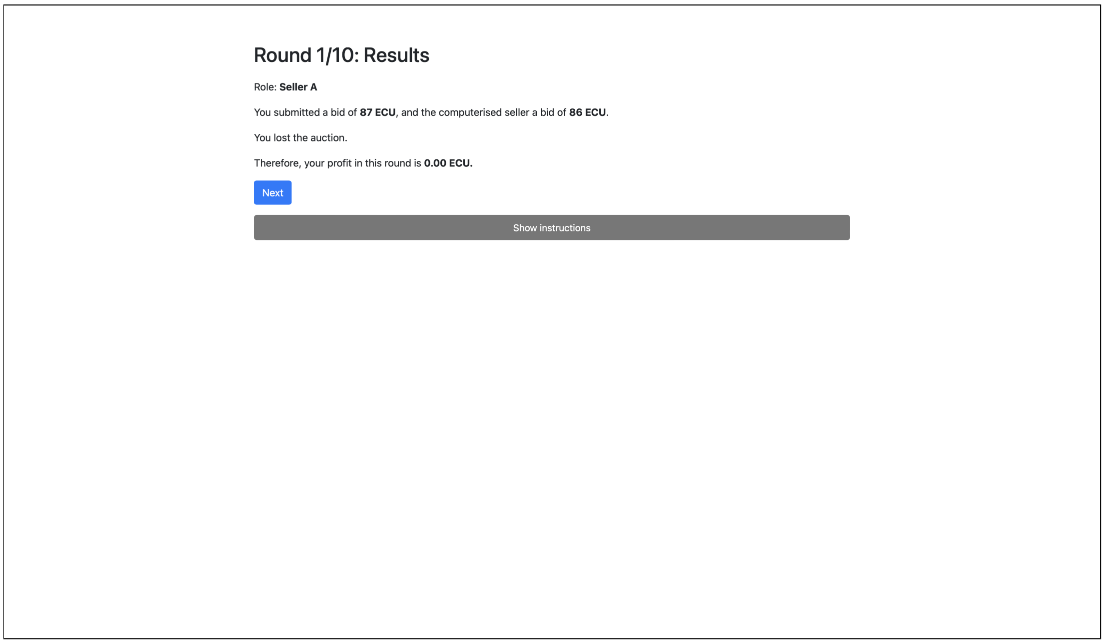
Round 1/10: Results
Role: **Seller A**
You submitted a bid of **87 ECU**, and the computerised seller a bid of **86 ECU**.
You lost the auction.
Therefore, your profit in this round is **0.00 ECU.**
Next
Show instructions

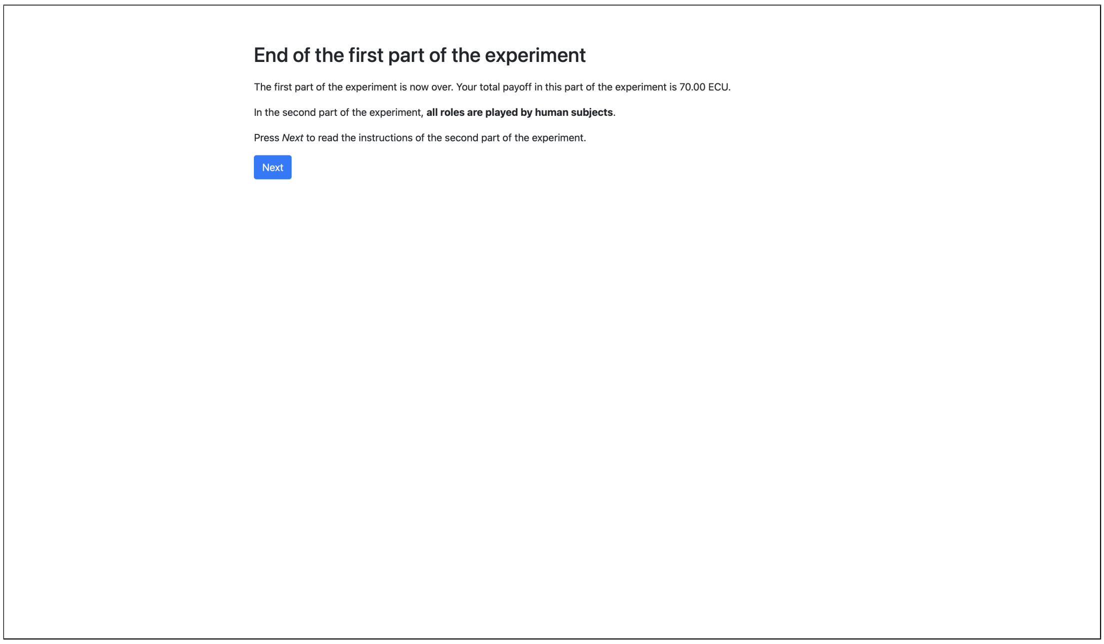
End of the first part of the experiment
The first part of the experiment is now over. Your total payoff in this part of the experiment is 70.00 ECU.
In the second part of the experiment, **all roles are played by human subjects**.
Press *Next* to read the instructions of the second part of the experiment.
Next

2.A.2 Commitment

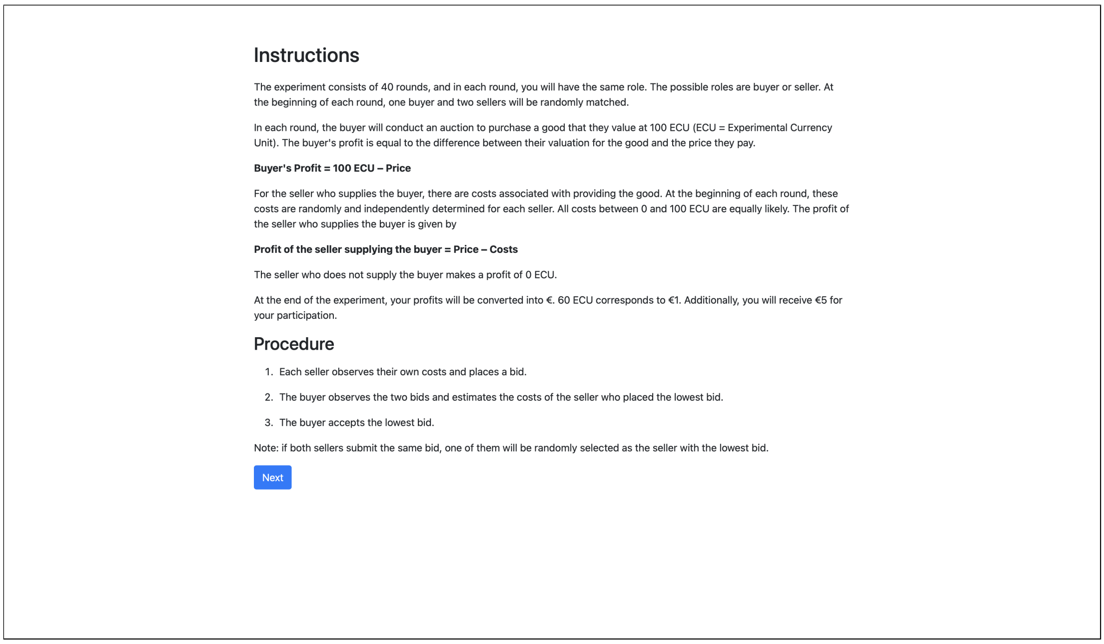

Instructions

The experiment consists of 40 rounds, and in each round, you will have the same role. The possible roles are buyer or seller. At the beginning of each round, one buyer and two sellers will be randomly matched.

In each round, the buyer will conduct an auction to purchase a good that they value at 100 ECU (ECU = Experimental Currency Unit). The buyer's profit is equal to the difference between their valuation for the good and the price they pay.

Buyer's Profit = 100 ECU – Price

For the seller who supplies the buyer, there are costs associated with providing the good. At the beginning of each round, these costs are randomly and independently determined for each seller. All costs between 0 and 100 ECU are equally likely. The profit of the seller who supplies the buyer is given by

Profit of the seller supplying the buyer = Price – Costs

The seller who does not supply the buyer makes a profit of 0 ECU.

At the end of the experiment, your profits will be converted into €. 60 ECU corresponds to €1. Additionally, you will receive €5 for your participation.

Procedure

1. Each seller observes their own costs and places a bid.
2. The buyer observes the two bids and estimates the costs of the seller who placed the lowest bid.
3. The buyer accepts the lowest bid.

Note: if both sellers submit the same bid, one of them will be randomly selected as the seller with the lowest bid.

Next

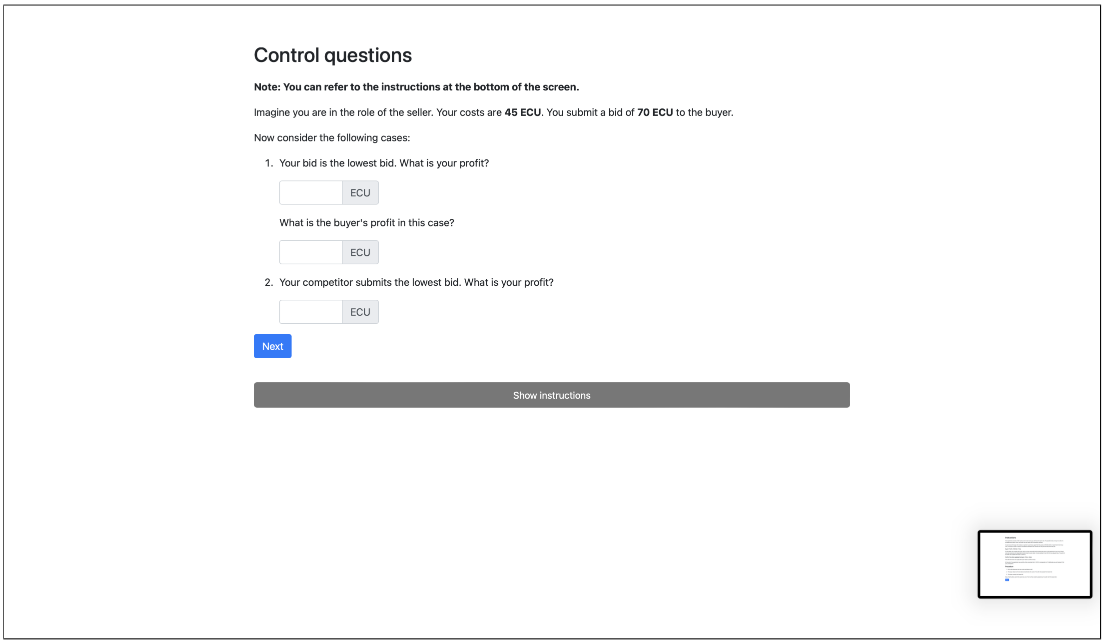

Control questions

Note: You can refer to the instructions at the bottom of the screen.

Imagine you are in the role of the seller. Your costs are **45 ECU**. You submit a bid of **70 ECU** to the buyer.

Now consider the following cases:

1. Your bid is the lowest bid. What is your profit?

   ECU

   What is the buyer's profit in this case?

   ECU

2. Your competitor submits the lowest bid. What is your profit?

   ECU

Next

Show instructions

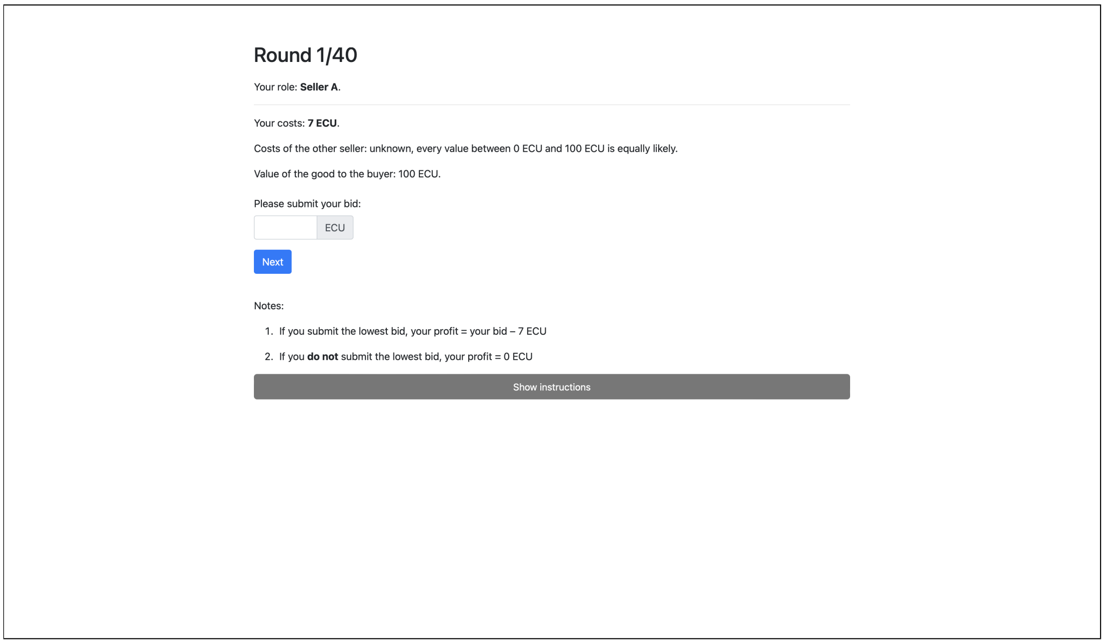
Round 1/40
Your role: Seller A.
Your costs: 7 ECU.
Costs of the other seller: unknown, every value between 0 ECU and 100 ECU is equally likely.
Value of the good to the buyer: 100 ECU.
Please submit your bid:
ECU
Next
Notes:
1. If you submit the lowest bid, your profit = your bid – 7 ECU
2. If you do not submit the lowest bid, your profit = 0 ECU
Show instructions

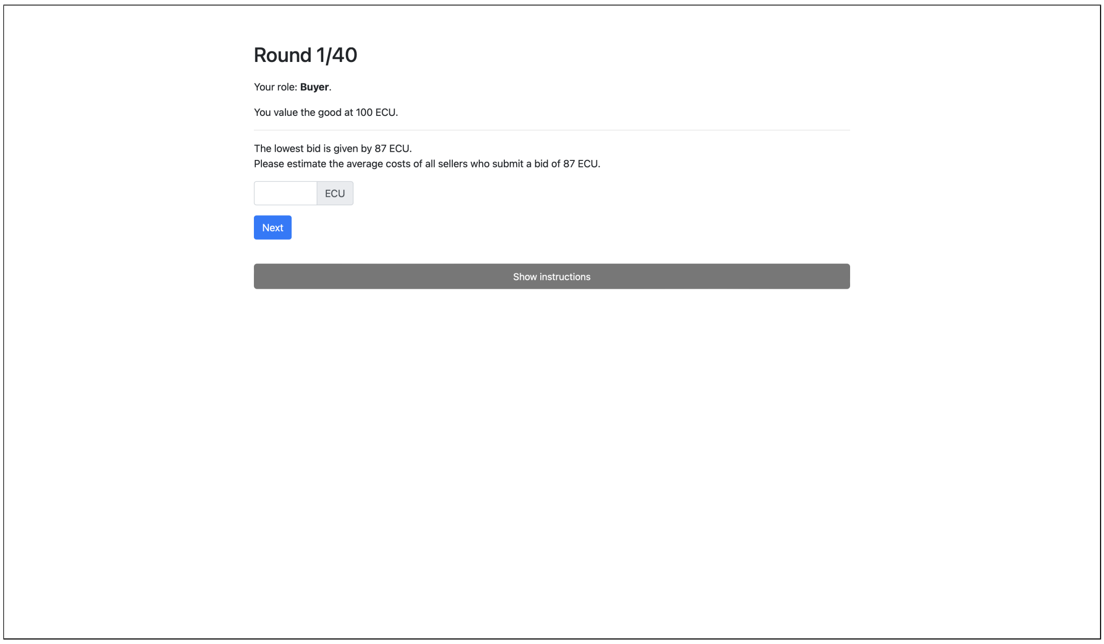
Round 1/40
Your role: Buyer.
You value the good at 100 ECU.
The lowest bid is given by 87 ECU.
Please estimate the average costs of all sellers who submit a bid of 87 ECU.
ECU
Next
Show instructions

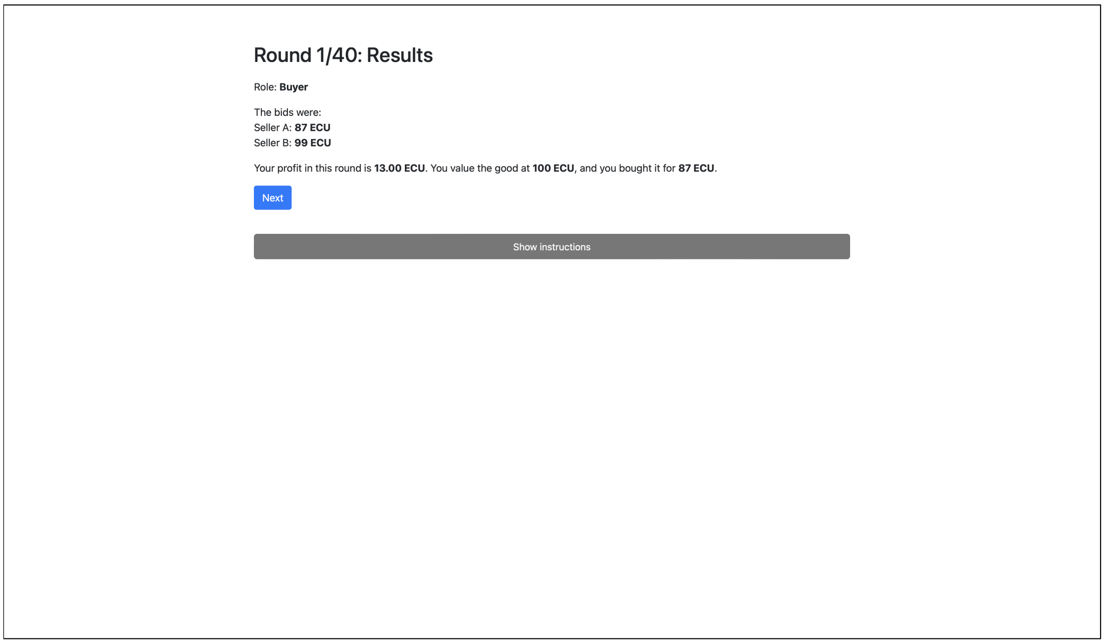
Round 1/40: Results
Role: Buyer
The bids were:
Seller A: 87 ECU
Seller B: 99 ECU
Your profit in this round is 13.00 ECU. You value the good at 100 ECU, and you bought it for 87 ECU.
Next
Show instructions

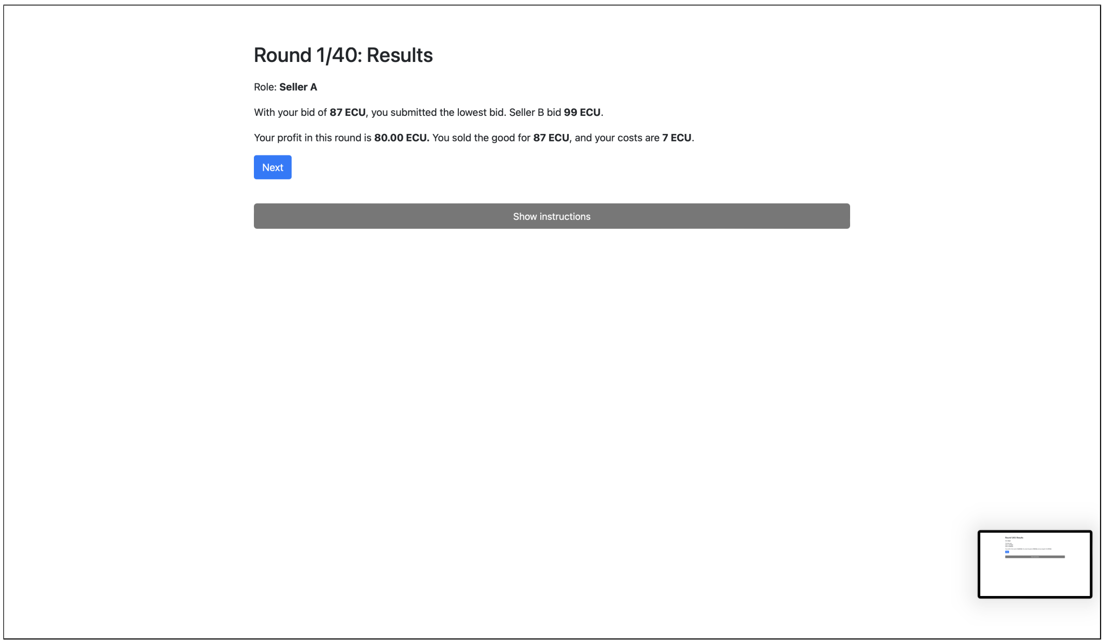
Round 1/40: Results
Role: Seller A
With your bid of 87 ECU, you submitted the lowest bid. Seller B bid 99 ECU.
Your profit in this round is 80.00 ECU. You sold the good for 87 ECU, and your costs are 7 ECU.
Next
Show instructions

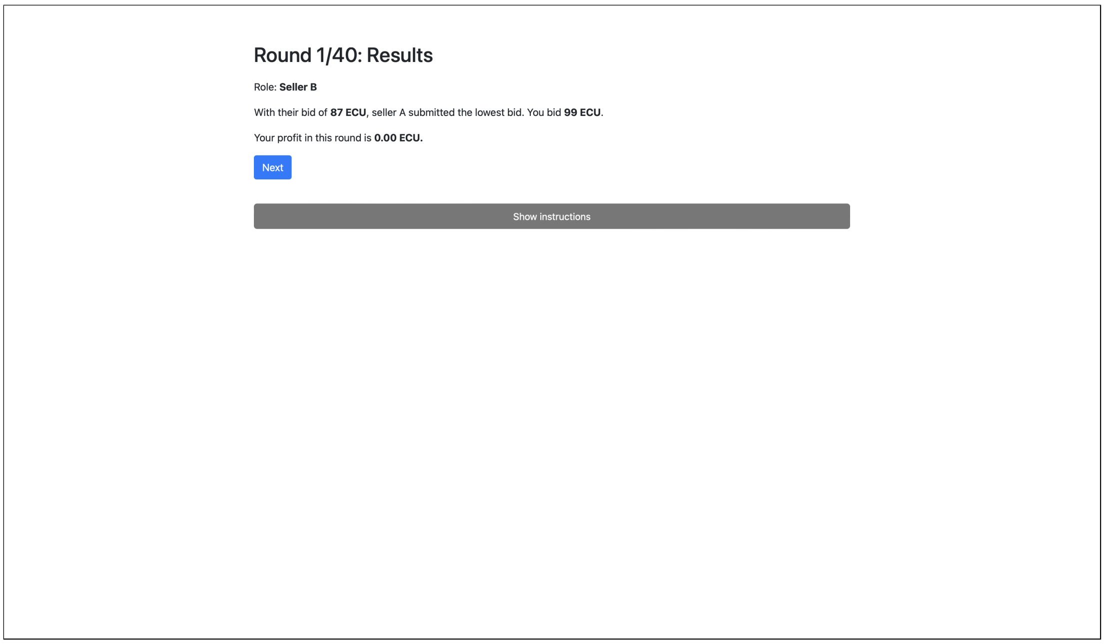
Round 1/40: Results
Role: Seller B
With their bid of 87 ECU, seller A submitted the lowest bid. You bid 99 ECU.
Your profit in this round is 0.00 ECU.
Next
Show instructions

2.A.3 No-Commitment

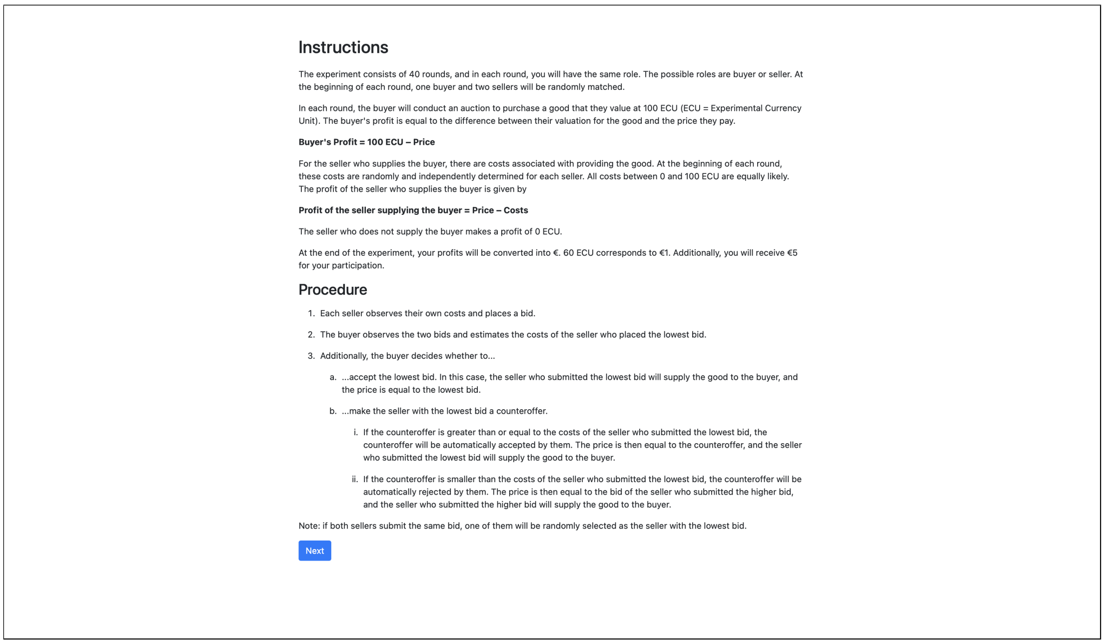

Instructions

The experiment consists of 40 rounds, and in each round, you will have the same role. The possible roles are buyer or seller. At the beginning of each round, one buyer and two sellers will be randomly matched.

In each round, the buyer will conduct an auction to purchase a good that they value at 100 ECU (ECU = Experimental Currency Unit). The buyer's profit is equal to the difference between their valuation for the good and the price they pay.

Buyer's Profit = 100 ECU – Price

For the seller who supplies the buyer, there are costs associated with providing the good. At the beginning of each round, these costs are randomly and independently determined for each seller. All costs between 0 and 100 ECU are equally likely. The profit of the seller who supplies the buyer is given by

Profit of the seller supplying the buyer = Price – Costs

The seller who does not supply the buyer makes a profit of 0 ECU.

At the end of the experiment, your profits will be converted into €. 60 ECU corresponds to €1. Additionally, you will receive €5 for your participation.

Procedure

1. Each seller observes their own costs and places a bid.
2. The buyer observes the two bids and estimates the costs of the seller who placed the lowest bid.
3. Additionally, the buyer decides whether to...
   a. ...accept the lowest bid. In this case, the seller who submitted the lowest bid will supply the good to the buyer, and the price is equal to the lowest bid.
   b. ...make the seller with the lowest bid a counteroffer.
      i. If the counteroffer is greater than or equal to the costs of the seller who submitted the lowest bid, the counteroffer will be automatically accepted by them. The price is then equal to the counteroffer, and the seller who submitted the lowest bid will supply the good to the buyer.
      ii. If the counteroffer is smaller than the costs of the seller who submitted the lowest bid, the counteroffer will be automatically rejected by them. The price is then equal to the bid of the seller who submitted the higher bid, and the seller who submitted the higher bid will supply the good to the buyer.

Note: if both sellers submit the same bid, one of them will be randomly selected as the seller with the lowest bid.

Next

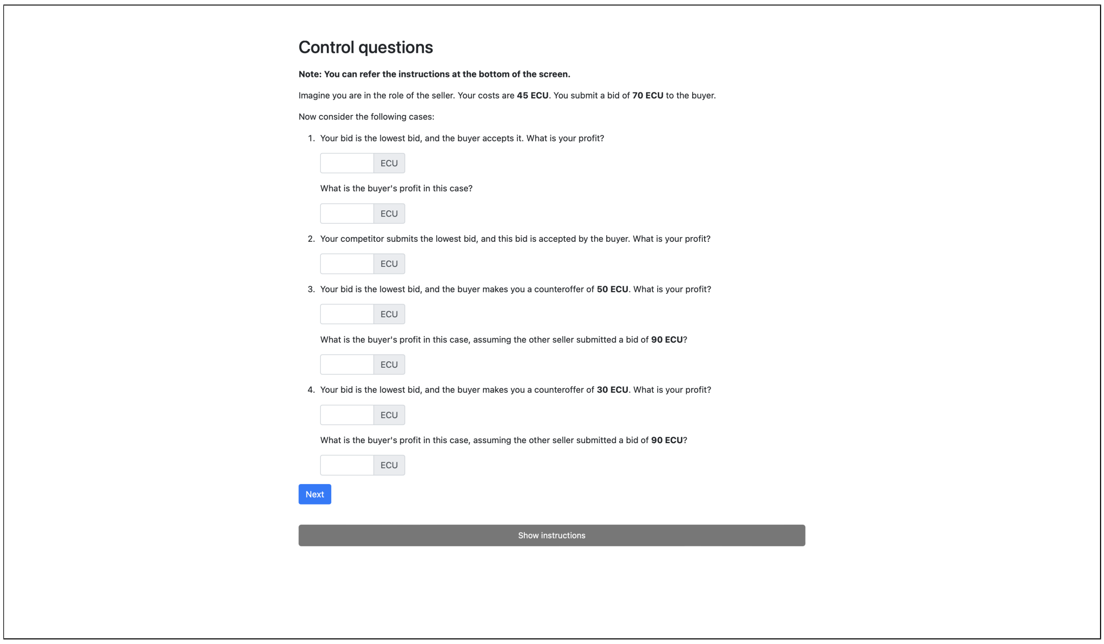

Control questions

Note: You can refer the instructions at the bottom of the screen.

Imagine you are in the role of the seller. Your costs are **45 ECU**. You submit a bid of **70 ECU** to the buyer.

Now consider the following cases:

1. Your bid is the lowest bid, and the buyer accepts it. What is your profit?

   ECU

   What is the buyer's profit in this case?

   ECU
2. Your competitor submits the lowest bid, and this bid is accepted by the buyer. What is your profit?

   ECU
3. Your bid is the lowest bid, and the buyer makes you a counteroffer of **50 ECU**. What is your profit?

   ECU

   What is the buyer's profit in this case, assuming the other seller submitted a bid of **90 ECU**?

   ECU
4. Your bid is the lowest bid, and the buyer makes you a counteroffer of **30 ECU**. What is your profit?

   ECU

   What is the buyer's profit in this case, assuming the other seller submitted a bid of **90 ECU**?

   ECU

Next

Show instructions

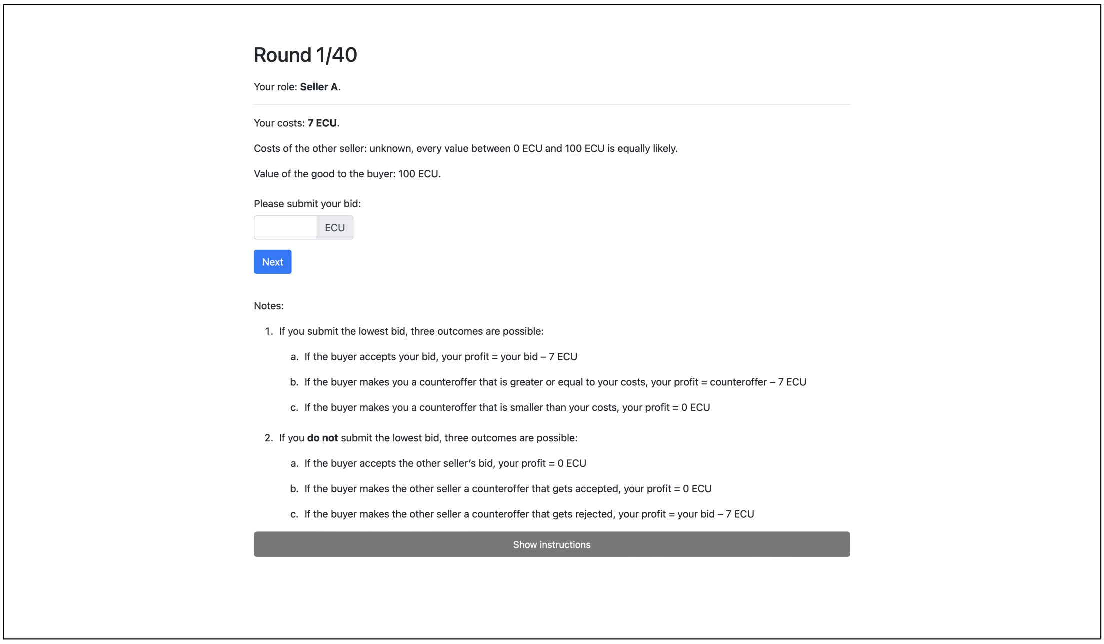

Round 1/40

Your role: **Seller A**.

Your costs: **7 ECU**.

Costs of the other seller: unknown, every value between 0 ECU and 100 ECU is equally likely.

Value of the good to the buyer: 100 ECU.

Please submit your bid:

ECU

Next

Notes:

1. If you submit the lowest bid, three outcomes are possible:
   a. If the buyer accepts your bid, your profit = your bid – 7 ECU
   b. If the buyer makes you a counteroffer that is greater or equal to your costs, your profit = counteroffer – 7 ECU
   c. If the buyer makes you a counteroffer that is smaller than your costs, your profit = 0 ECU
2. If you **do not** submit the lowest bid, three outcomes are possible:
   a. If the buyer accepts the other seller's bid, your profit = 0 ECU
   b. If the buyer makes the other seller a counteroffer that gets accepted, your profit = 0 ECU
   c. If the buyer makes the other seller a counteroffer that gets rejected, your profit = your bid – 7 ECU

Show instructions

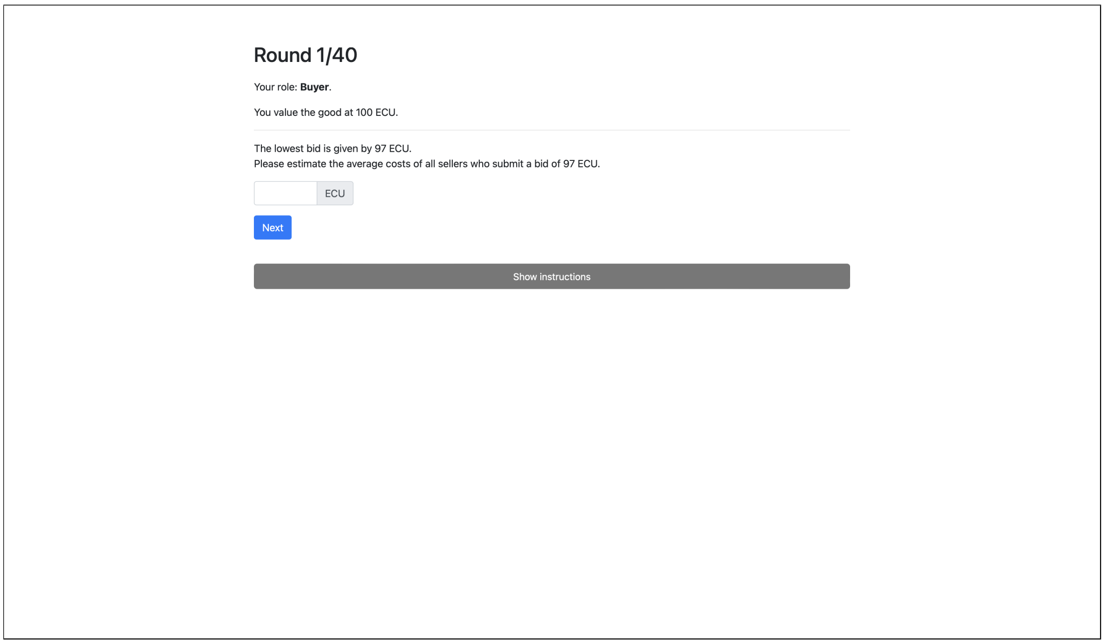

Round 1/40

Your role: **Buyer**.

You value the good at 100 ECU.

The lowest bid is given by 97 ECU.
Please estimate the average costs of all sellers who submit a bid of 97 ECU.

ECU

Next

Show instructions

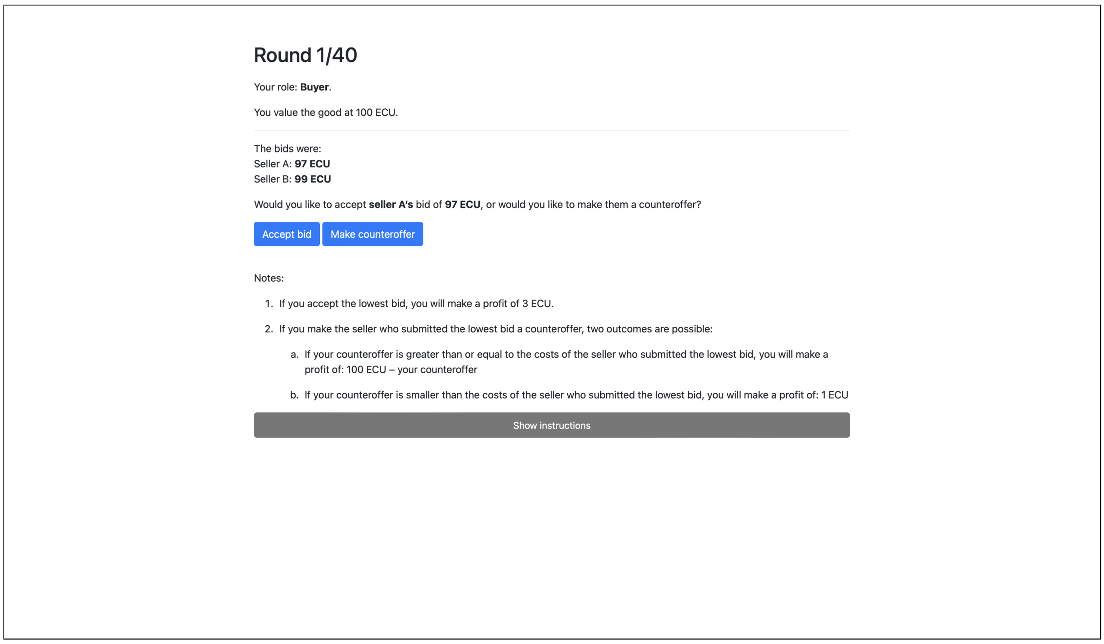
Round 1/40
Your role: Buyer.
You value the good at 100 ECU.
The bids were:
Seller A: 97 ECU
Seller B: 99 ECU
Would you like to accept seller A's bid of 97 ECU, or would you like to make them a counteroffer?
Accept bid
Make counteroffer
Notes:
1. If you accept the lowest bid, you will make a profit of 3 ECU.
2. If you make the seller who submitted the lowest bid a counteroffer, two outcomes are possible:
a. If your counteroffer is greater than or equal to the costs of the seller who submitted the lowest bid, you will make a profit of: 100 ECU – your counteroffer
b. If your counteroffer is smaller than the costs of the seller who submitted the lowest bid, you will make a profit of: 1 ECU
Show instructions

Buyer accepts lowest bid

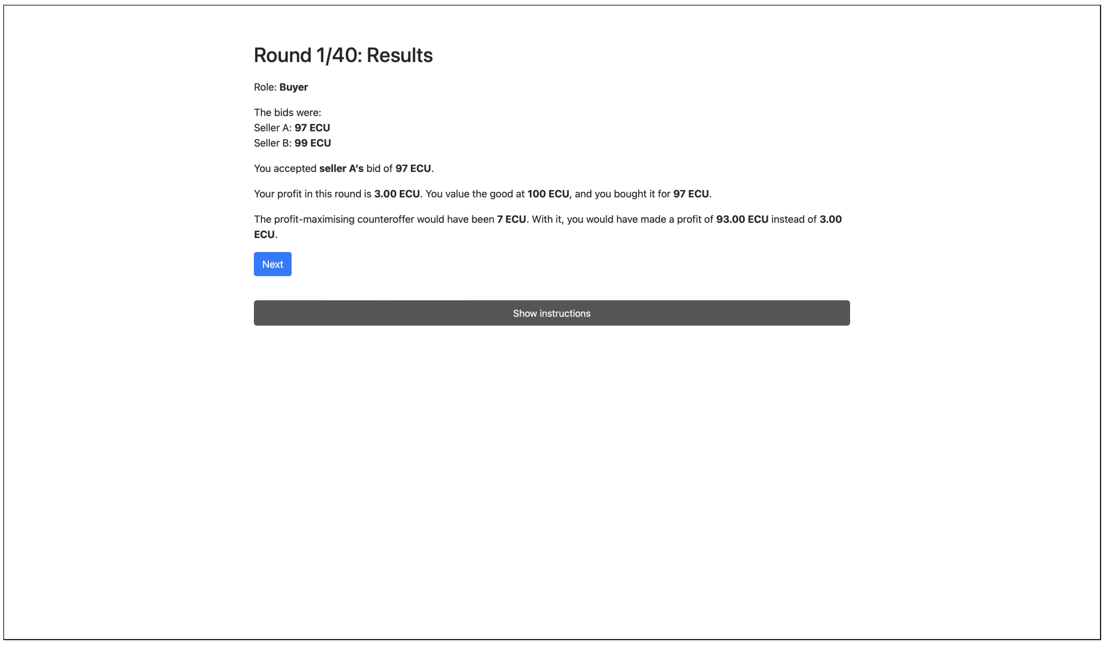

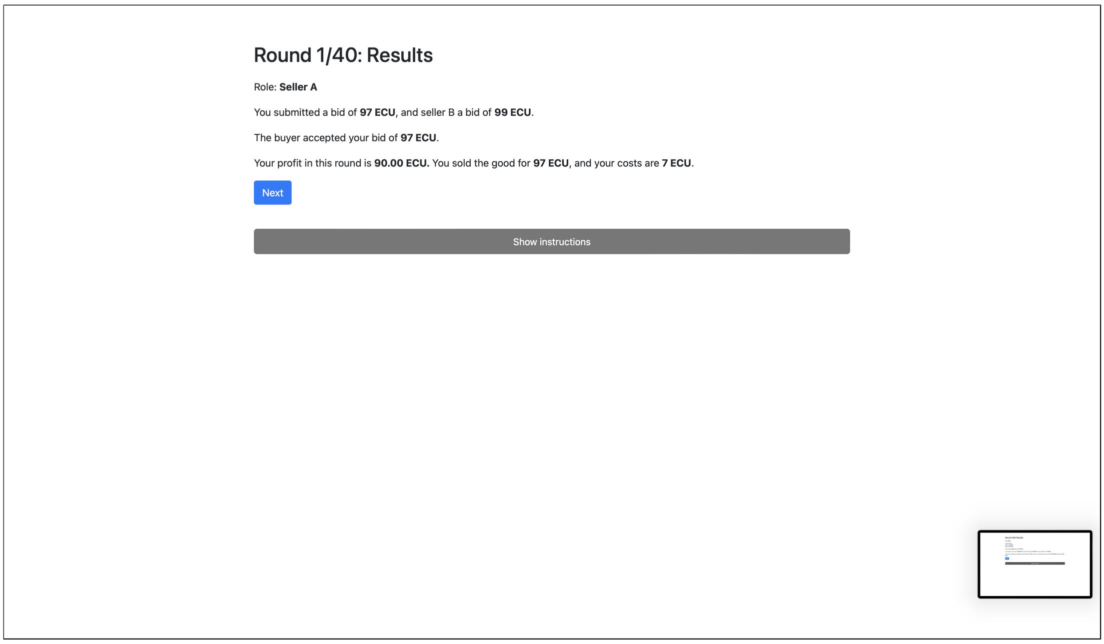

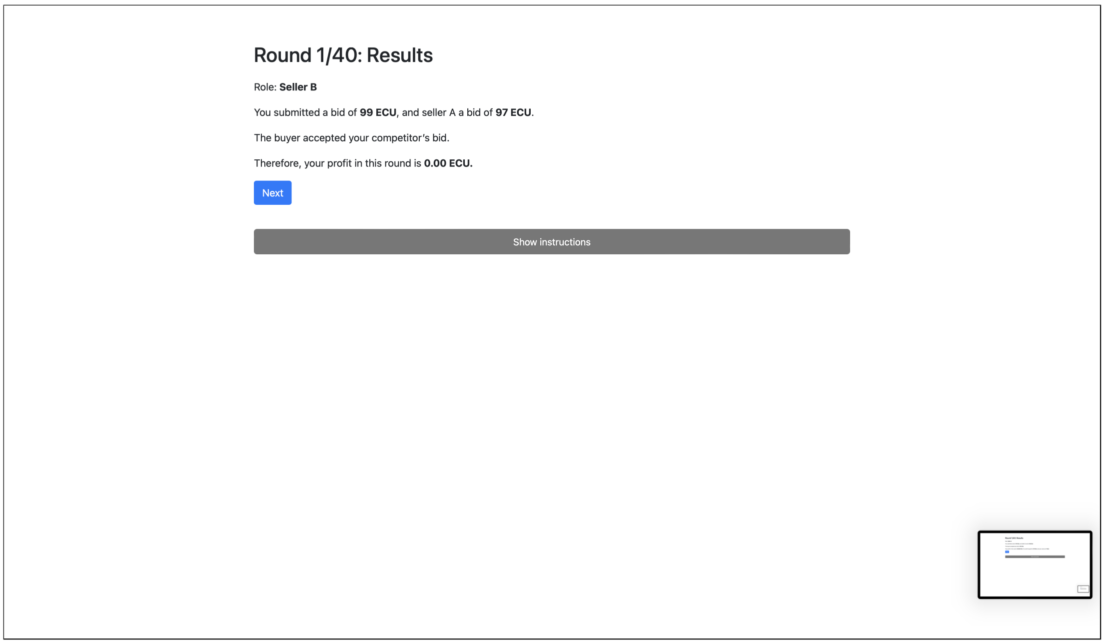
Round 1/40: Results
Role: Seller B
You submitted a bid of 99 ECU, and seller A a bid of 97 ECU.
The buyer accepted your competitor's bid.
Therefore, your profit in this round is 0.00 ECU.
Next
Show instructions

Buyer makes counteroffer that gets accepted

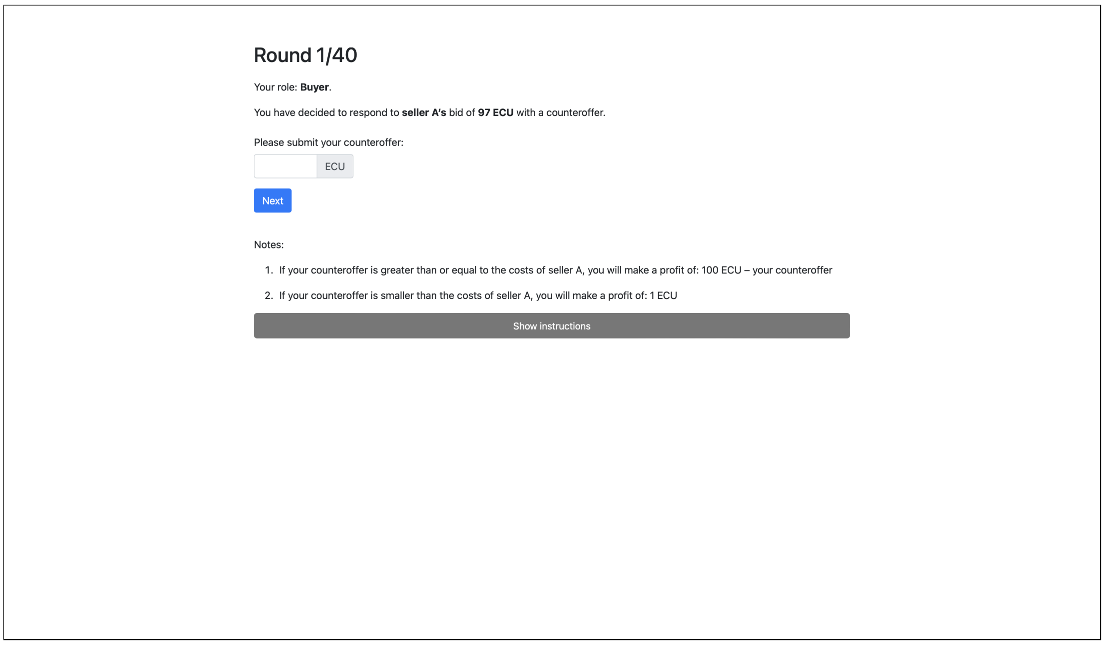

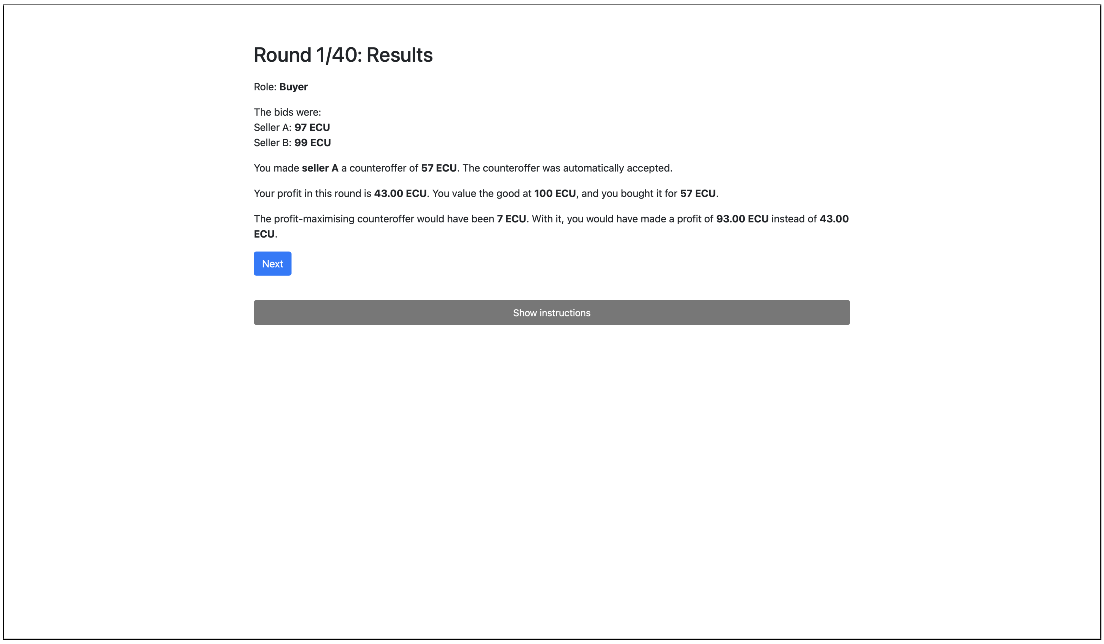

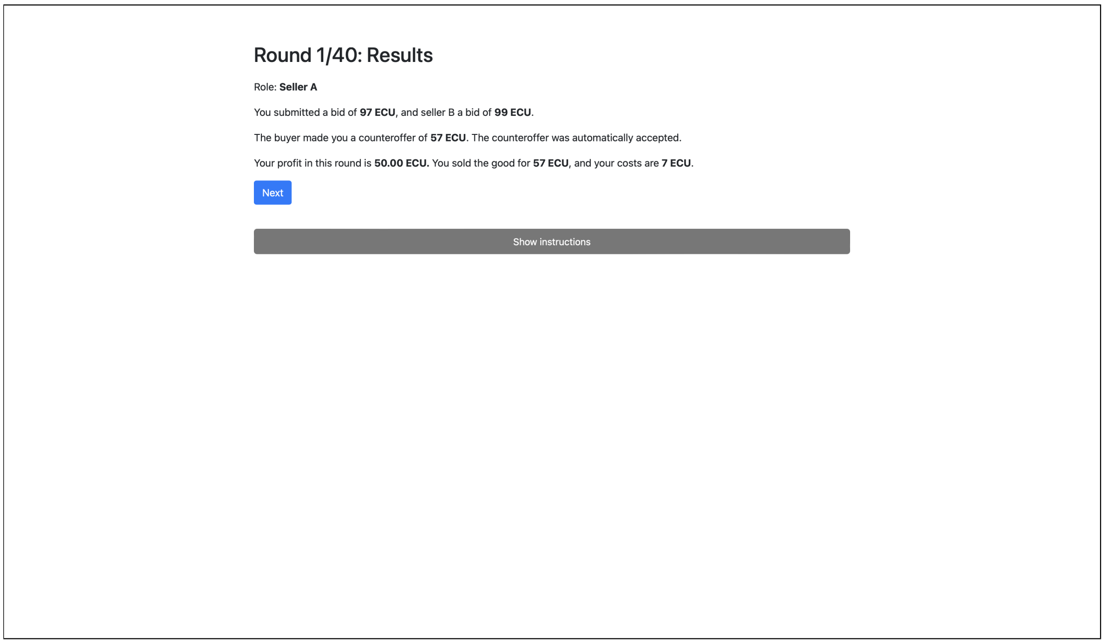
Round 1/40: Results
Role: Seller A
You submitted a bid of 97 ECU, and seller B a bid of 99 ECU.
The buyer made you a counteroffer of 57 ECU. The counteroffer was automatically accepted.
Your profit in this round is 50.00 ECU. You sold the good for 57 ECU, and your costs are 7 ECU.
Next
Show instructions

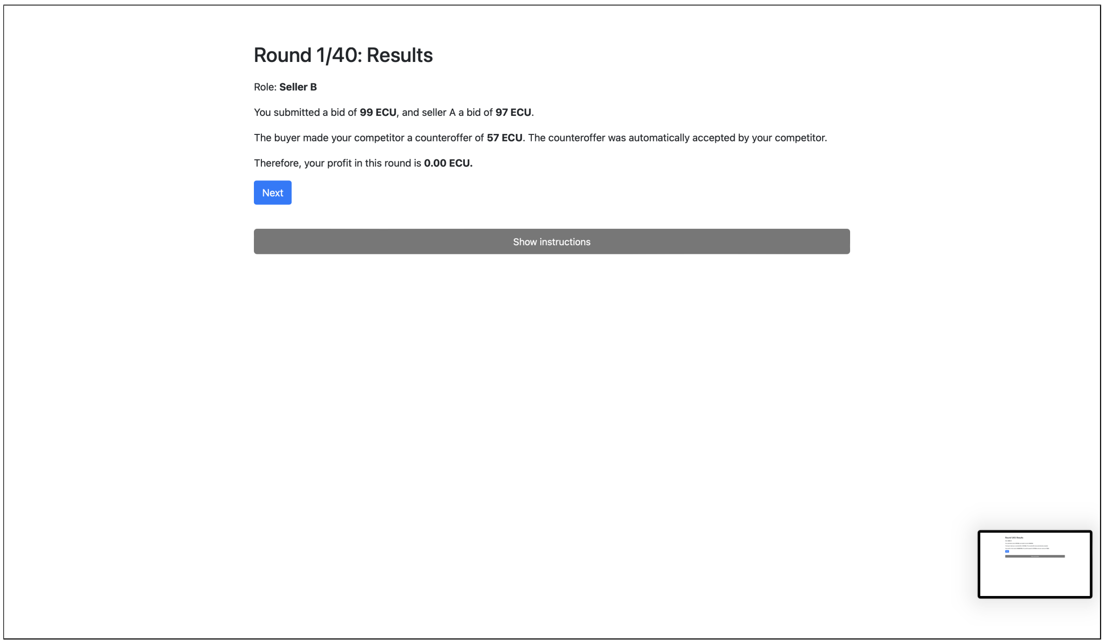
Round 1/40: Results
Role: Seller B
You submitted a bid of 99 ECU, and seller A a bid of 97 ECU.
The buyer made your competitor a counteroffer of 57 ECU. The counteroffer was automatically accepted by your competitor.
Therefore, your profit in this round is 0.00 ECU.
Next
Show instructions

Buyer makes counteroffer that gets rejected

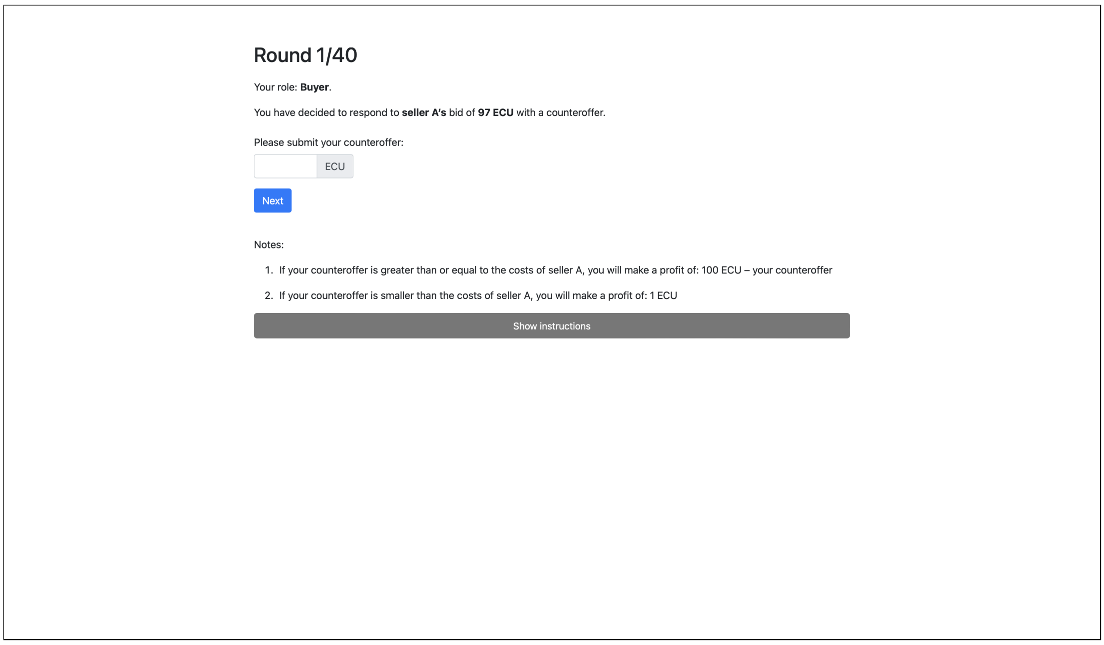

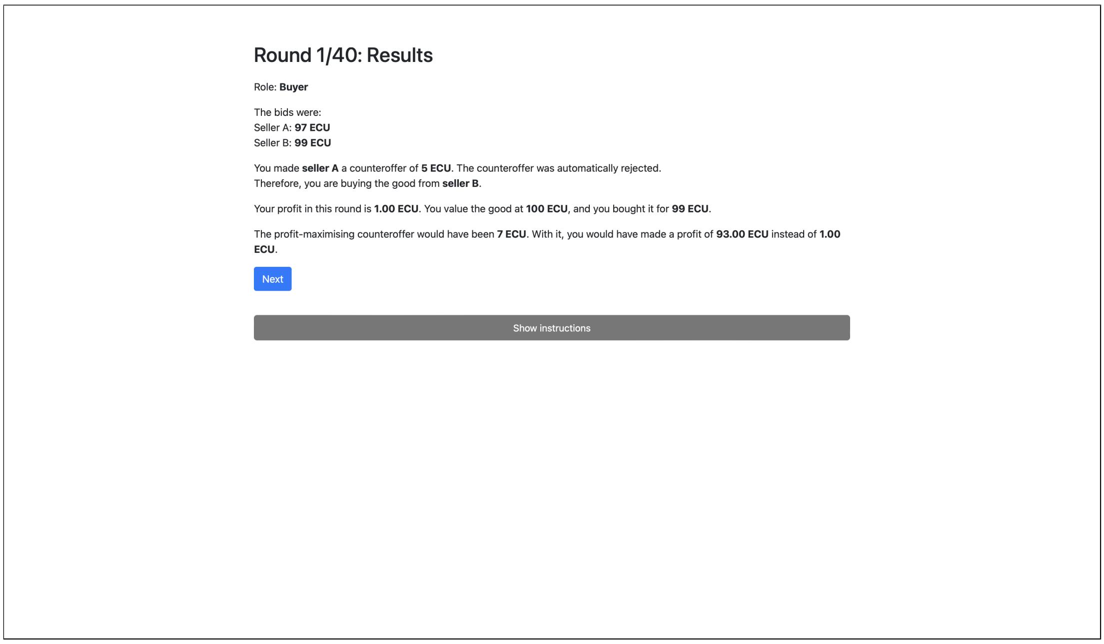

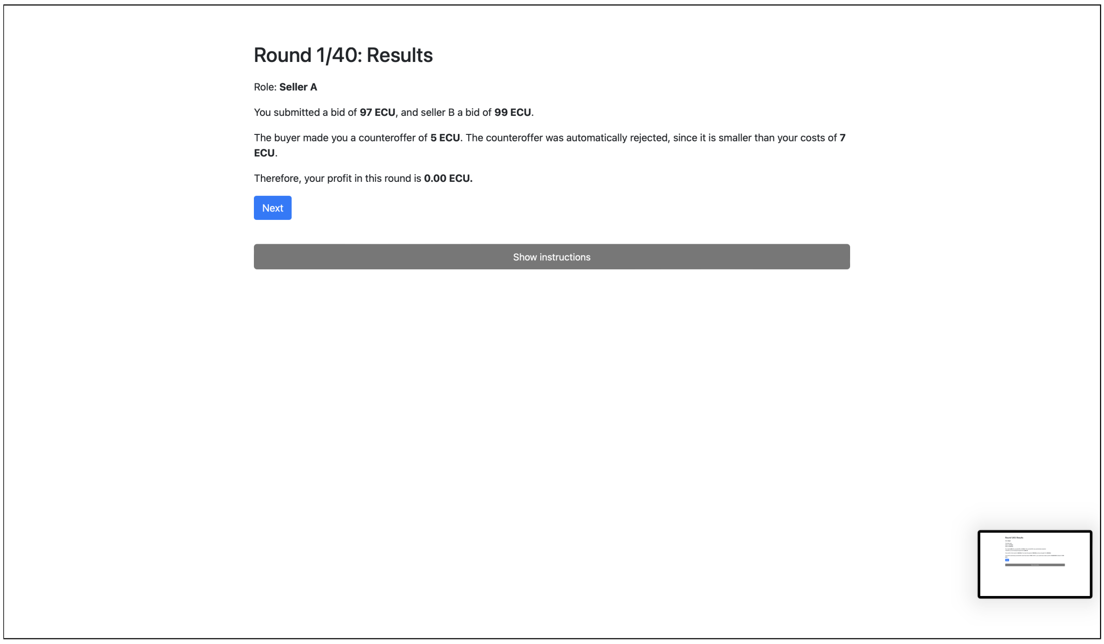
Round 1/40: Results
Role: Seller A
You submitted a bid of 97 ECU, and seller B a bid of 99 ECU.
The buyer made you a counteroffer of 5 ECU. The counteroffer was automatically rejected, since it is smaller than your costs of 7 ECU.
Therefore, your profit in this round is 0.00 ECU.
Next
Show instructions

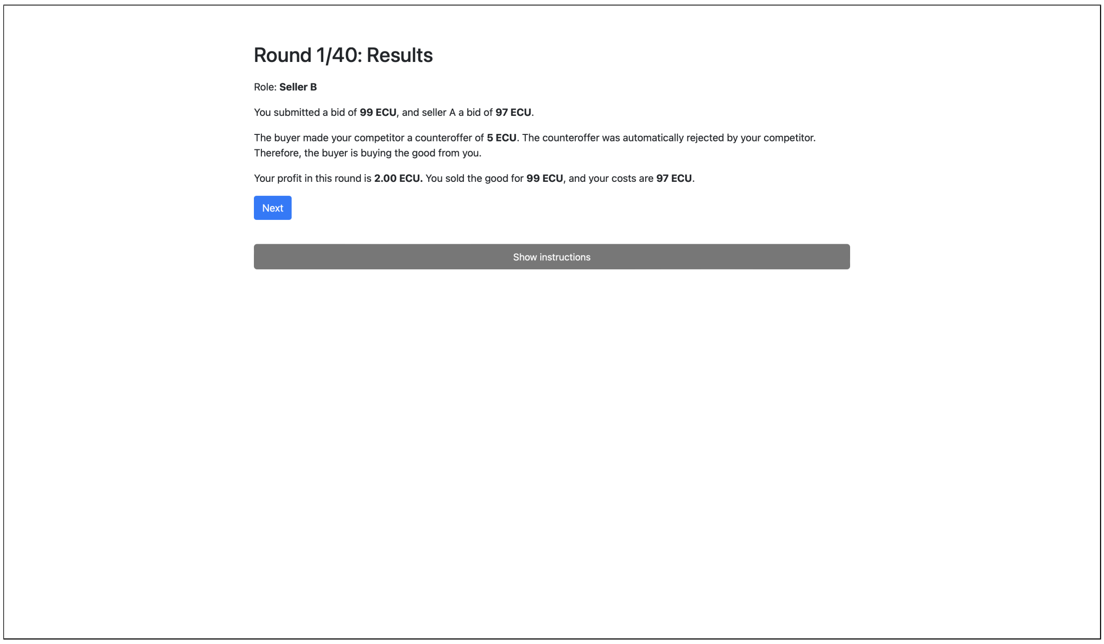
Round 1/40: Results
Role: Seller B
You submitted a bid of 99 ECU, and seller A a bid of 97 ECU.
The buyer made your competitor a counteroffer of 5 ECU. The counteroffer was automatically rejected by your competitor. Therefore, the buyer is buying the good from you.
Your profit in this round is 2.00 ECU. You sold the good for 99 ECU, and your costs are 97 ECU.
Next
Show instructions

2.A.4 No-Commitment DS

Instructions

The experiment consists of 40 rounds, and in each round, you will have the same role. The possible roles are buyer or seller. At the beginning of each round, one buyer and two sellers will be randomly matched.

In each round, the buyer will conduct an auction to purchase a good that they value at 100 ECU (ECU = Experimental Currency Unit). The buyer's profit is equal to the difference between their valuation for the good and the price they pay.

Buyer's Profit = 100 ECU – Price

For the seller who supplies the buyer, there are costs associated with providing the good. At the beginning of each round, these costs are randomly and independently determined for each seller. All costs between 0 and 100 ECU are equally likely. The profit of the seller who supplies the buyer is given by

Profit of the seller supplying the buyer = Price – Costs

The seller who does not supply the buyer makes a profit of 0 ECU.

At the end of the experiment, your profits will be converted into €. 60 ECU corresponds to €1. Additionally, you will receive €5 for your participation.

Procedure

1. Each seller observes their own costs and places a bid.
2. The buyer observes the two bids and estimates the costs of the seller who placed the lowest bid.
3. Additionally, the buyer decides whether to...
   a. ...accept the lowest bid. In this case, the seller who submitted the lowest bid will supply the good to the buyer, and the price is equal to the lowest bid.
   b. ...make the seller with the lowest bid a counteroffer.
      i. If the counteroffer is greater than or equal to the costs of the seller who submitted the lowest bid, the counteroffer will be automatically accepted by them. The price is then equal to the counteroffer, and the seller who submitted the lowest bid will supply the good to the buyer.
      ii. If the counteroffer is smaller than the costs of the seller who submitted the lowest bid, the counteroffer will be automatically rejected by them. The price is then equal to the bid of the seller who submitted the higher bid, and the seller who submitted the higher bid will supply the good to the buyer.

Note: if both sellers submit the same bid, one of them will be randomly selected as the seller with the lowest bid.

Next

Control questions

Note: You can refer the instructions at the bottom of the screen.

Imagine you are in the role of the seller. Your costs are **45 ECU**. You submit a bid of **70 ECU** to the buyer.

Now consider the following cases:

1. Your bid is the lowest bid, and the buyer accepts it. What is your profit?

   ECU

   What is the buyer's profit in this case?

   ECU

2. Your competitor submits the lowest bid, and this bid is accepted by the buyer. What is your profit?

   ECU

3. Your bid is the lowest bid, and the buyer makes you a counteroffer of **50 ECU**. What is your profit?

   ECU

   What is the buyer's profit in this case, assuming the other seller submitted a bid of **90 ECU**?

   ECU

4. Your bid is the lowest bid, and the buyer makes you a counteroffer of **30 ECU**. What is your profit?

   ECU

   What is the buyer's profit in this case, assuming the other seller submitted a bid of **90 ECU**?

   ECU

Next

Show instructions

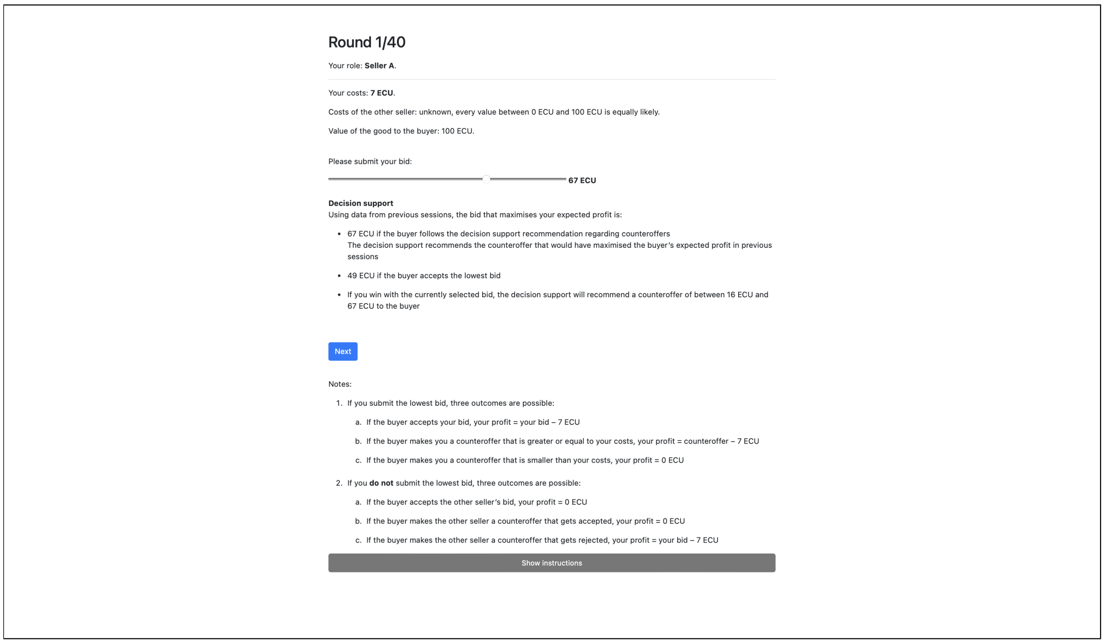
Round 1/40
Your role: Seller A.
Your costs: 7 ECU.
Costs of the other seller: unknown, every value between 0 ECU and 100 ECU is equally likely.
Value of the good to the buyer: 100 ECU.
Please submit your bid:
67 ECU
Decision support
Using data from previous sessions, the bid that maximises your expected profit is:
67 ECU if the buyer follows the decision support recommendation regarding counteroffers
The decision support recommends the counteroffer that would have maximised the buyer’s expected profit in previous sessions
49 ECU if the buyer accepts the lowest bid
If you win with the currently selected bid, the decision support will recommend a counteroffer of between 16 ECU and 67 ECU to the buyer
Next
Notes:
1. If you submit the lowest bid, three outcomes are possible:
a. If the buyer accepts your bid, your profit = your bid – 7 ECU
b. If the buyer makes you a counteroffer that is greater or equal to your costs, your profit = counteroffer – 7 ECU
c. If the buyer makes you a counteroffer that is smaller than your costs, your profit = 0 ECU
2. If you do not submit the lowest bid, three outcomes are possible:
a. If the buyer accepts the other seller’s bid, your profit = 0 ECU
b. If the buyer makes the other seller a counteroffer that gets accepted, your profit = 0 ECU
c. If the buyer makes the other seller a counteroffer that gets rejected, your profit = your bid – 7 ECU
Show instructions

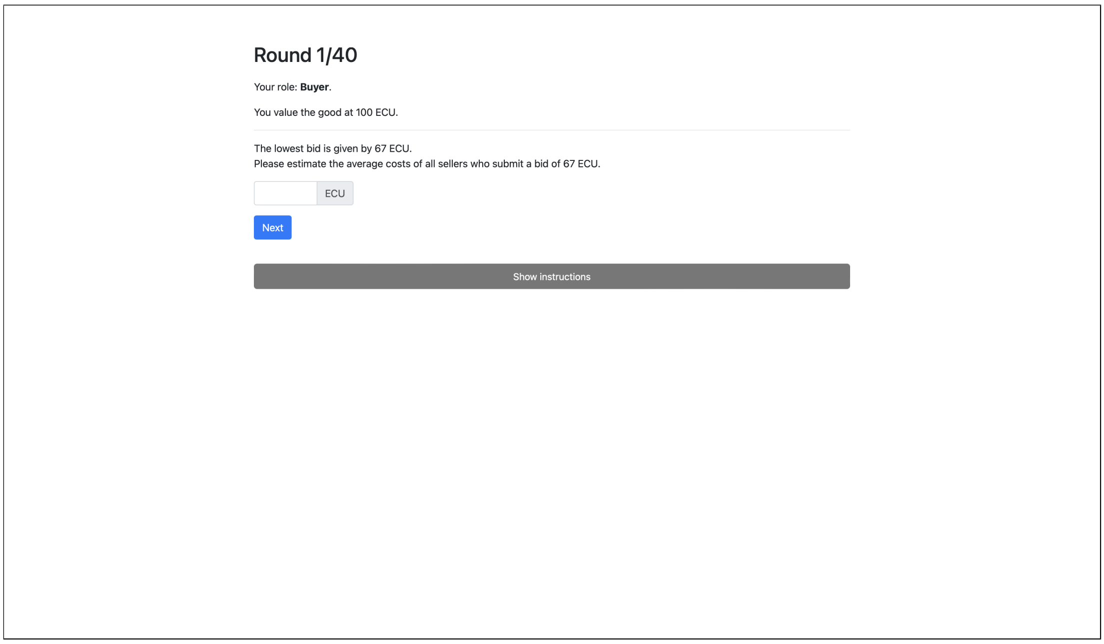
Round 1/40
Your role: Buyer.
You value the good at 100 ECU.
The lowest bid is given by 67 ECU.
Please estimate the average costs of all sellers who submit a bid of 67 ECU.
ECU
Next
Show instructions

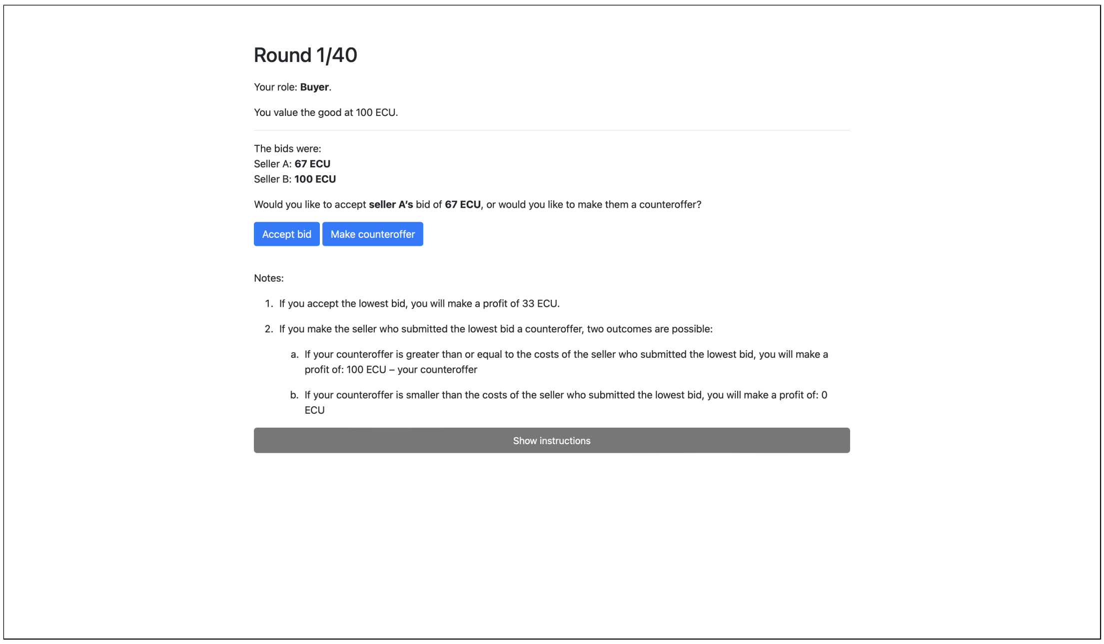
Round 1/40
Your role: Buyer.
You value the good at 100 ECU.
The bids were:
Seller A: 67 ECU
Seller B: 100 ECU
Would you like to accept seller A's bid of 67 ECU, or would you like to make them a counteroffer?
Accept bid
Make counteroffer
Notes:
1. If you accept the lowest bid, you will make a profit of 33 ECU.
2. If you make the seller who submitted the lowest bid a counteroffer, two outcomes are possible:
a. If your counteroffer is greater than or equal to the costs of the seller who submitted the lowest bid, you will make a profit of: 100 ECU – your counteroffer
b. If your counteroffer is smaller than the costs of the seller who submitted the lowest bid, you will make a profit of: 0 ECU
Show instructions

Buyer accepts lowest bid

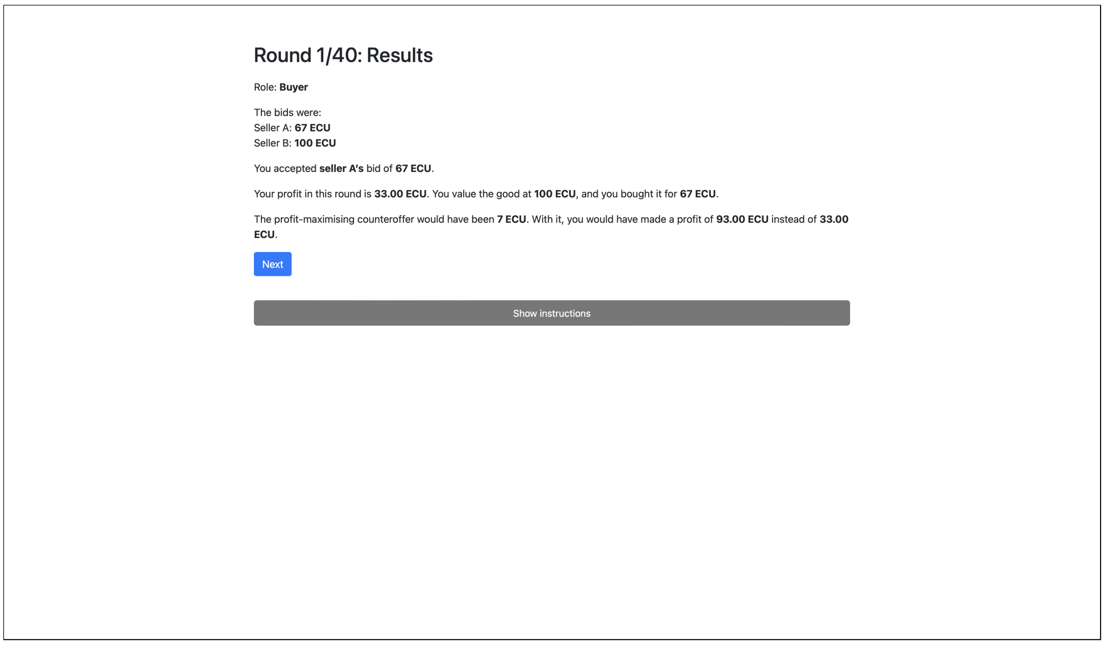

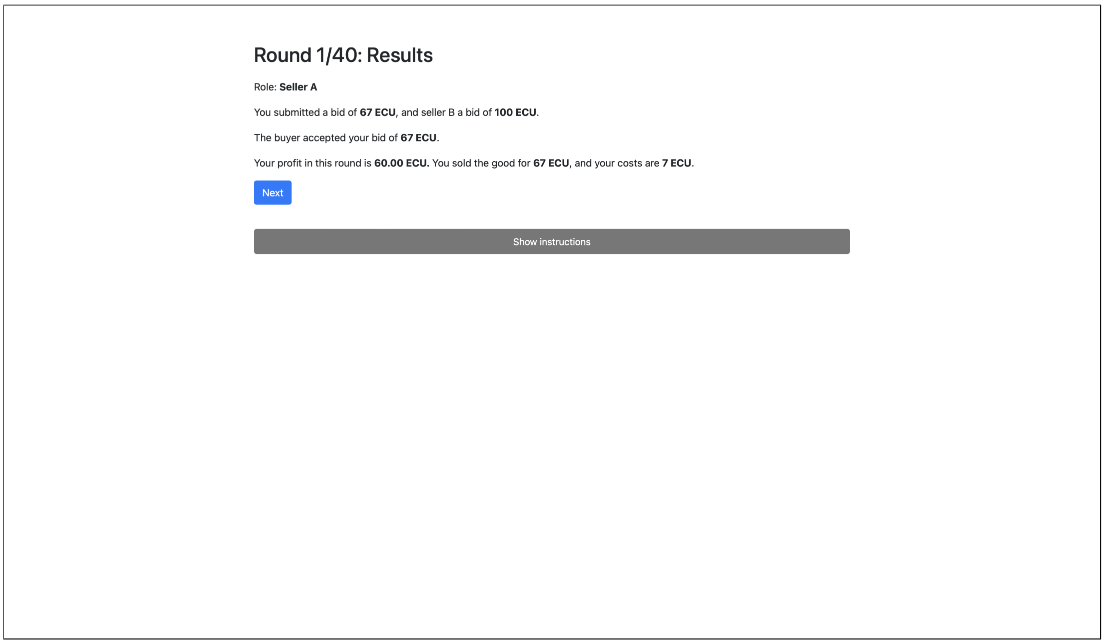

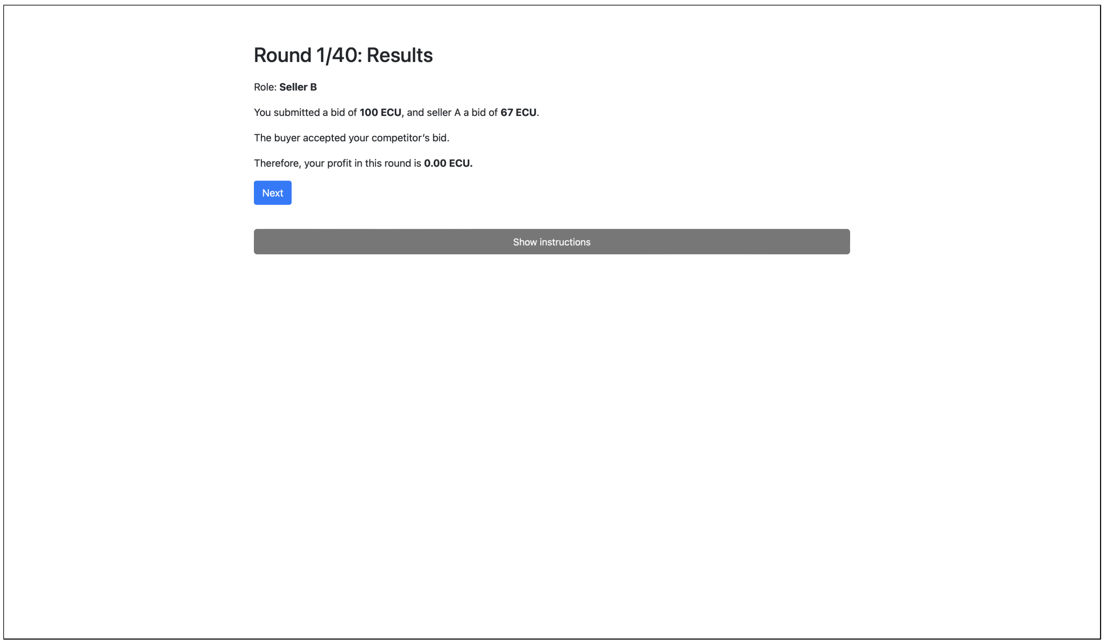
Round 1/40: Results
Role: Seller B
You submitted a bid of 100 ECU, and seller A a bid of 67 ECU.
The buyer accepted your competitor's bid.
Therefore, your profit in this round is 0.00 ECU.
Next
Show instructions

Buyer makes counteroffer that gets accepted

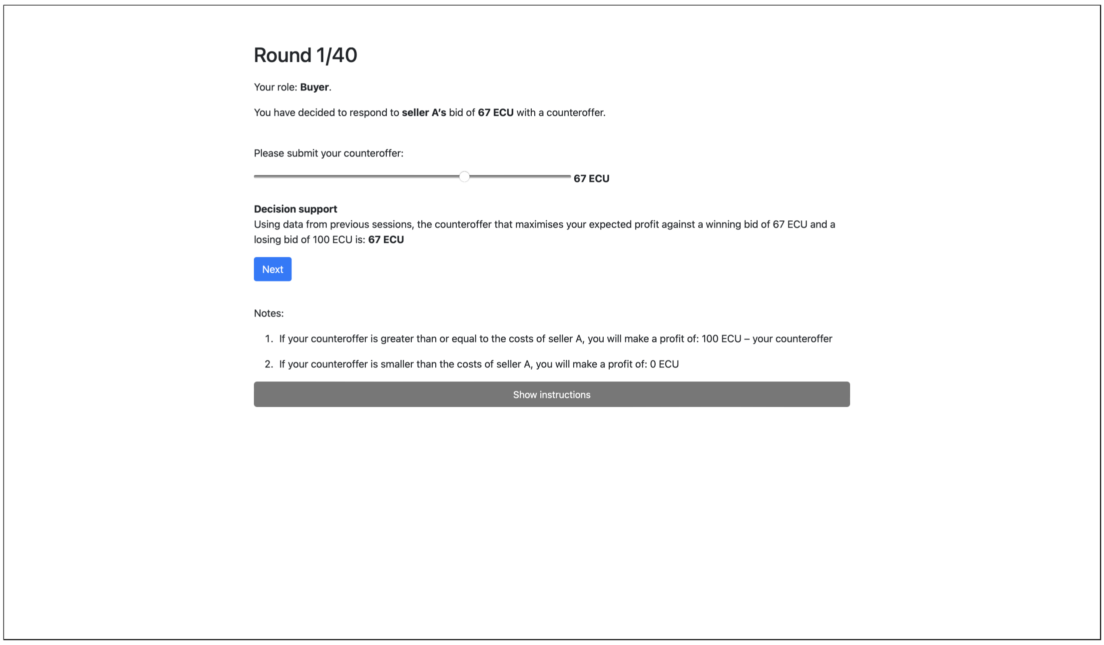

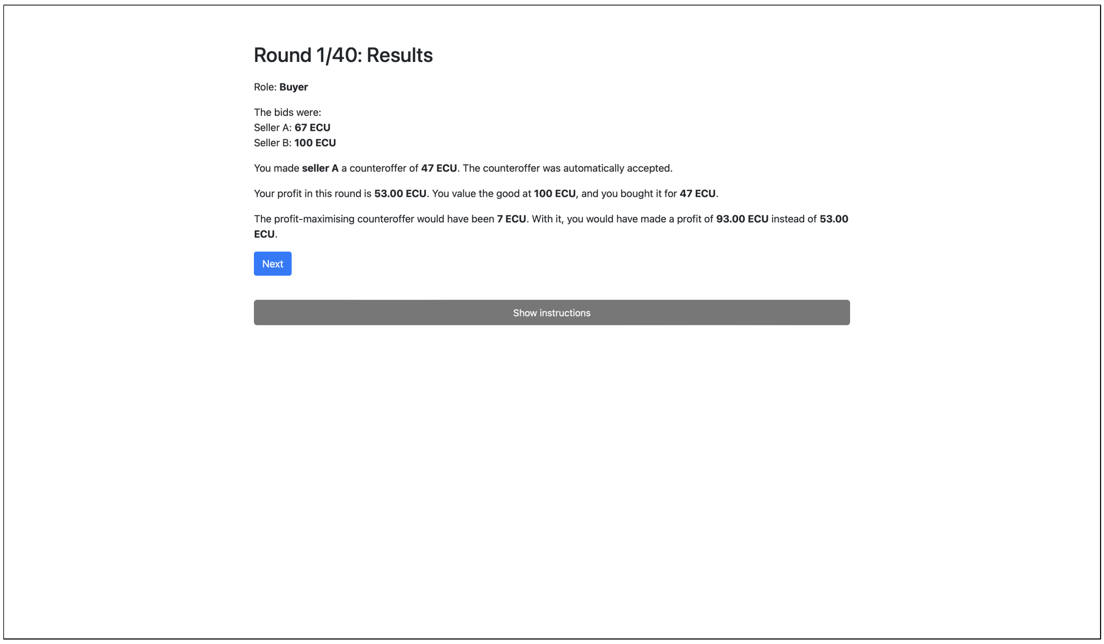

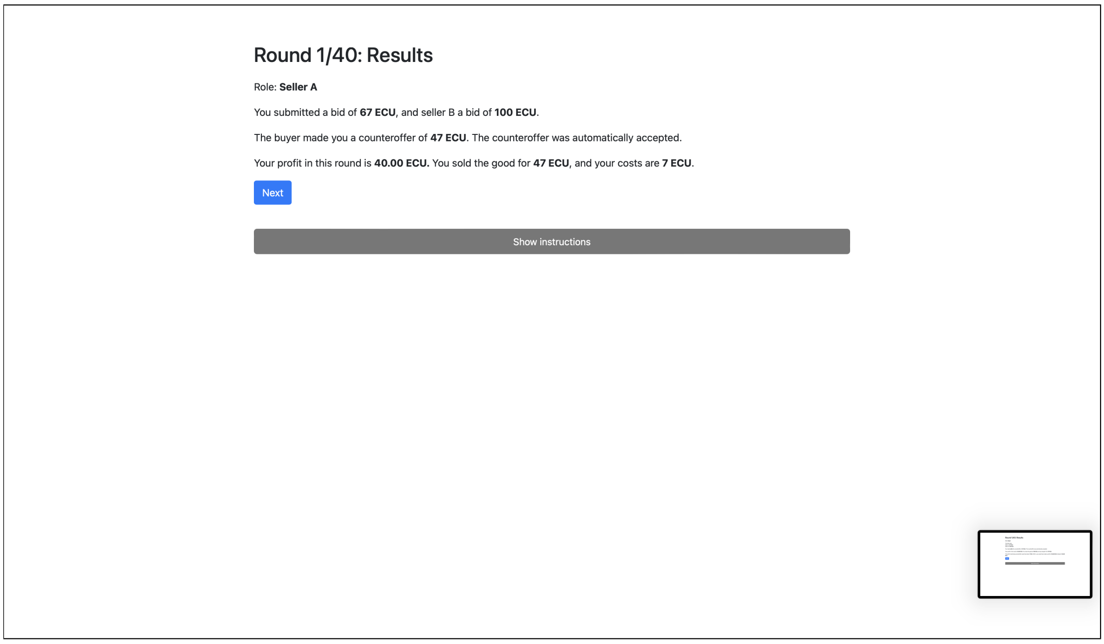
Round 1/40: Results
Role: Seller A
You submitted a bid of 67 ECU, and seller B a bid of 100 ECU.
The buyer made you a counteroffer of 47 ECU. The counteroffer was automatically accepted.
Your profit in this round is 40.00 ECU. You sold the good for 47 ECU, and your costs are 7 ECU.
Next
Show instructions

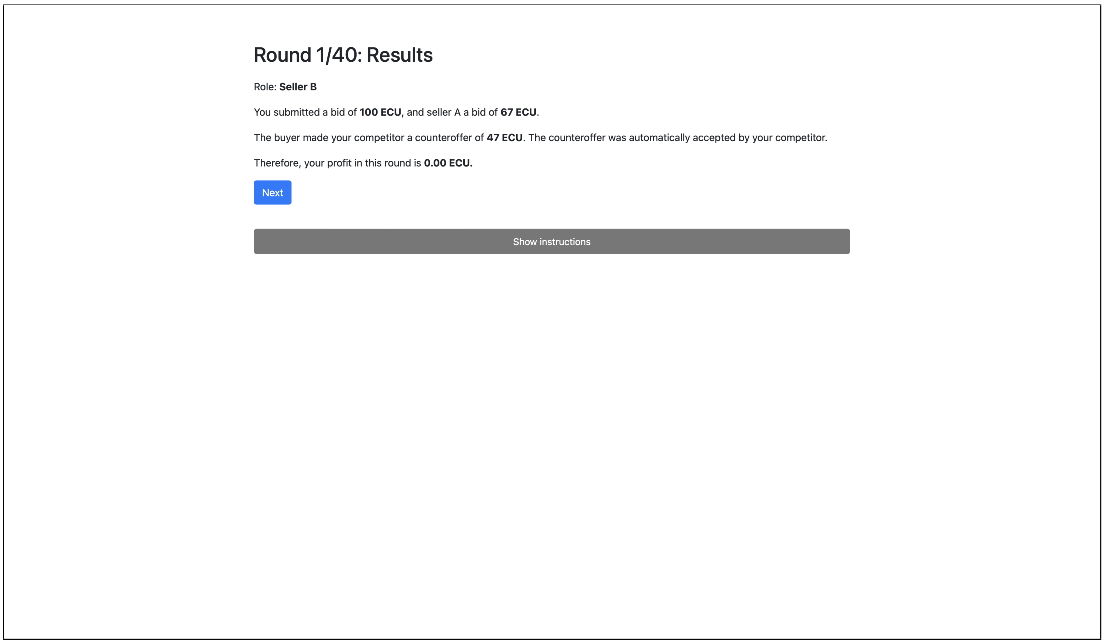
Round 1/40: Results
Role: Seller B
You submitted a bid of 100 ECU, and seller A a bid of 67 ECU.
The buyer made your competitor a counteroffer of 47 ECU. The counteroffer was automatically accepted by your competitor.
Therefore, your profit in this round is 0.00 ECU.
Next
Show instructions

Buyer makes counteroffer that gets rejected

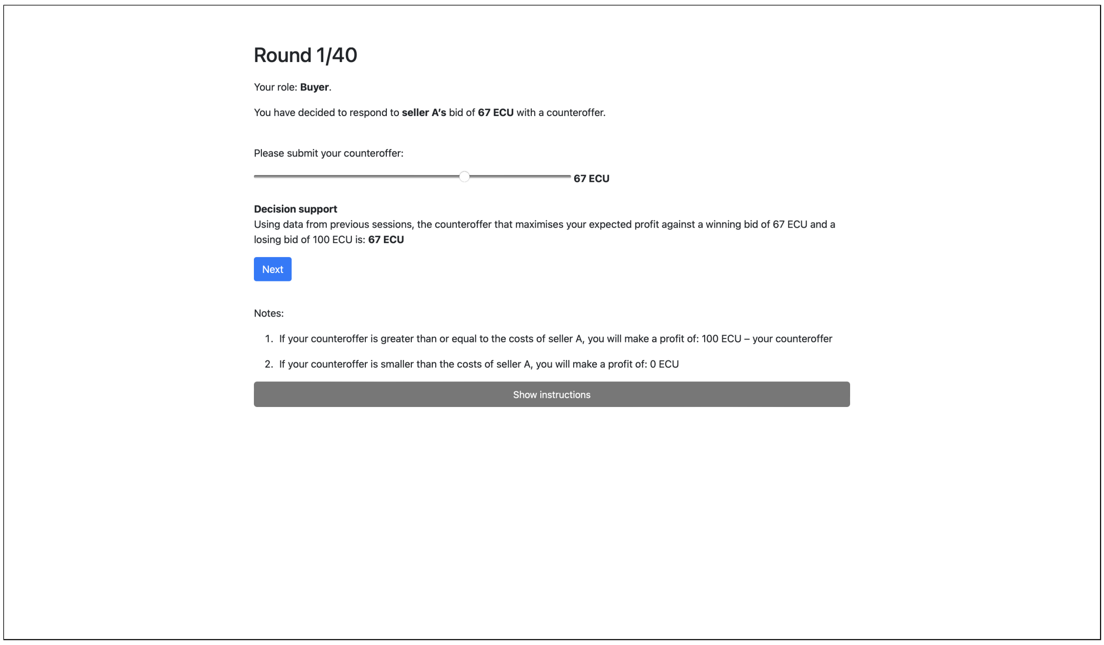

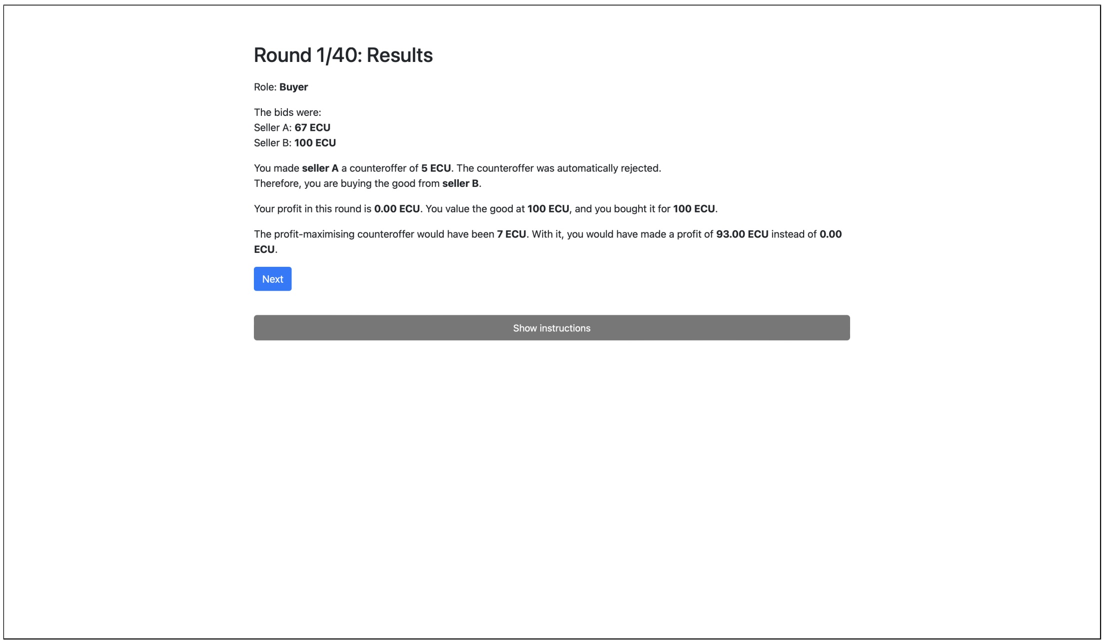

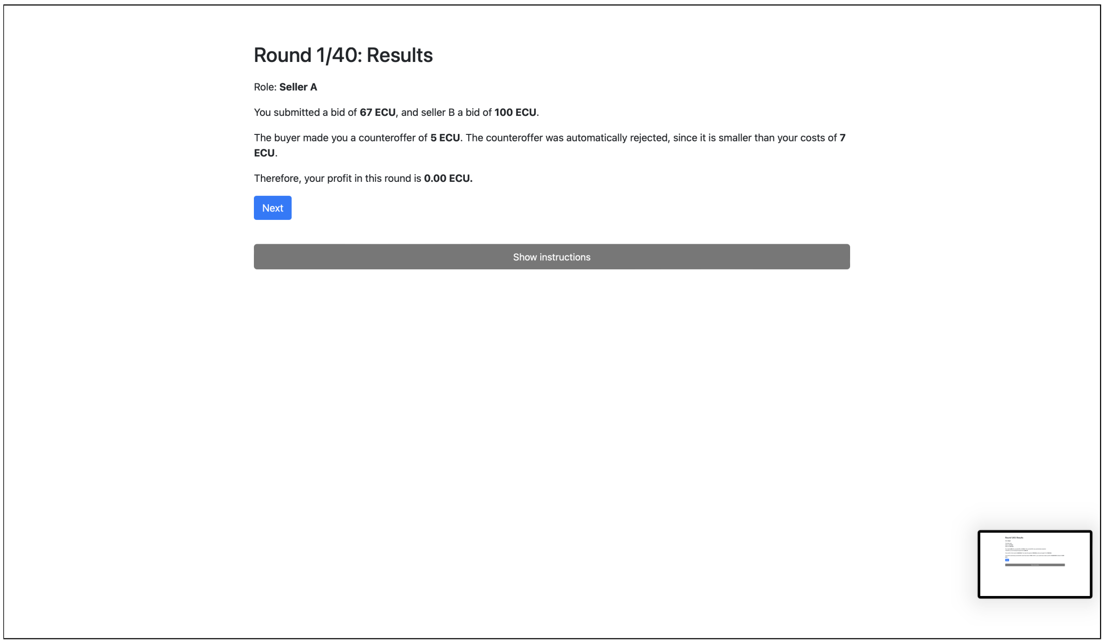
Round 1/40: Results
Role: Seller A
You submitted a bid of 67 ECU, and seller B a bid of 100 ECU.
The buyer made you a counteroffer of 5 ECU. The counteroffer was automatically rejected, since it is smaller than your costs of 7 ECU.
Therefore, your profit in this round is 0.00 ECU.
Next
Show instructions

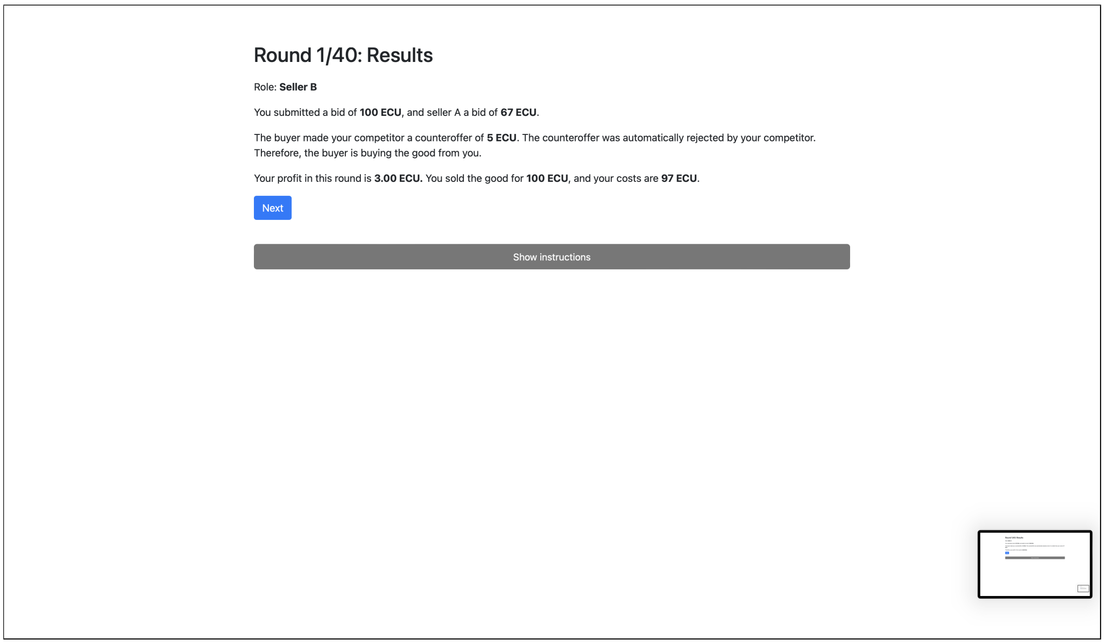
Round 1/40: Results
Role: Seller B
You submitted a bid of 100 ECU, and seller A a bid of 67 ECU.
The buyer made your competitor a counteroffer of 5 ECU. The counteroffer was automatically rejected by your competitor. Therefore, the buyer is buying the good from you.
Your profit in this round is 3.00 ECU. You sold the good for 100 ECU, and your costs are 97 ECU.
Next
Show instructions

2.A.5 Final screen

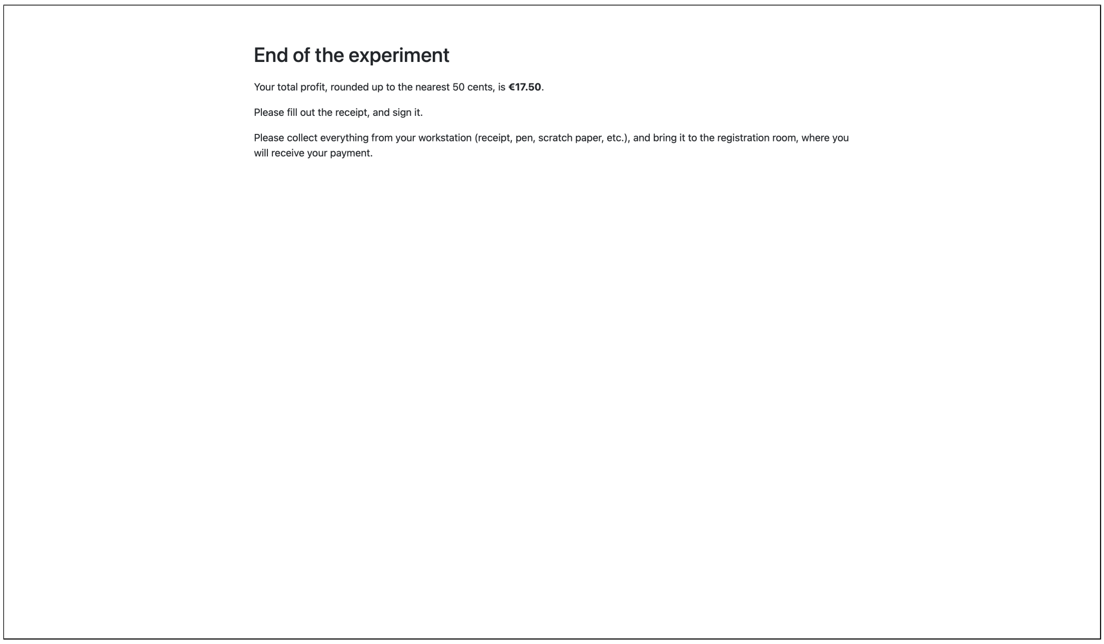

2.B Commitment AI treatment

Below, we provide the full JSON payloads of one of the AI sellers for the first three rounds. In these payloads, *role: system* refers to information we provide the AI subject, while *role: assistant* refers to the AI subject's response.

{'role': 'system', 'content': 'You are playing an economics experiment. The next message will give you the HTML instructions of this experiment.'}

{'role': 'system', 'content': 'The experiment consists of 40 rounds, and in each round, you will have the same role. You will be playing the role of seller. At the beginning of each round, one buyer and two sellers will be randomly matched.In each round, the buyer will conduct an auction to purchase a good that they value at 100 ECU (ECU = Experimental Currency Unit). The buyer's profit is equal to the difference between their valuation for the good and the price they pay. Buyer's Profit = 100 ECU ? Price For the seller who supplies the buyer, there are costs associated with providing the good. At the beginning of each round, these costs are randomly and independently determined for each seller. All costs between 0 and 100 ECU are equally likely. The profit of the seller who supplies the buyer is given by Profit of the seller supplying the buyer = Price ? Costs The seller who does not supply the buyer makes a profit of 0 ECU. At the end of the experiment, your profits will be converted into Euros. 60 ECU corresponds to 1 Euro. Additionally, you will receive 5 Euros for your participation. *Procedure* Each seller observes their own costs and places a bid. The buyer observes the two bids and estimates the costs of the seller who placed the lowest bid. The buyer accepts the lowest bid. Note: if both sellers submit the same bid, one of them will be randomly selected as the seller with the lowest bid.'}

{'role': 'system', 'content': 'Please write a concise summary of the five most important points of the instructions.'}

{'role': 'assistant', 'content': "1. **Role and Matching**: You play as a seller in 40 rounds, where each round involves one buyer and two sellers randomly matched. 2. **Buyer's Valuation and Profit**: The buyer values the good at 100 ECU, and their profit is the difference between this valuation and the price paid. 3. **Seller's Costs and Profit**: Each seller has randomly determined costs between 0 and 100 ECU. The seller's profit is the price received minus these costs. The non-supplying seller earns 0 ECU. 4. **Bidding Process**: Sellers observe their costs and place bids. The buyer accepts the lowest bid, and if bids are tied, one seller is chosen randomly. 5. **Earnings Conversion**: Profits are converted to euros at a rate of 60 ECU per 1 Euro, with an additional 5 Euro participation fee."}

{'role': 'system', 'content': 'In the next message, we will provide you with control questions as well as their answers to make sure you understand the experiment.'}

{'role': 'system', 'content': "Control questions You are in the role of the seller. Your costs are 45 ECU. You submit a bid of 70 ECU to the buyer. Now consider the following cases: Your bid is the lowest bid. What is your profit? The answer is 25 ECU. What is the buyer's profit in this case? The answer is 30 ECU. Your competitor submits the lowest bid. What is your profit? The answer is 0 ECU. "}

{'role': 'system', 'content': 'For the entire experiment, you will be playing in the role of seller.'}

{'role': 'system', 'content': "Your goal is to maximise your profit over all 40 rounds of the experiment. Devise a step-by-step strategy to achieve this goal. Be concise. Think about what auction type this is and what your bidding function should look like. That is, for each possible cost realisation, which bid should you place in order to maximise your expected profit? Assume that the other seller will use the same bidding function as yours. Ensure that your bidding function is indeed optimal in this case. Note that (i) no bidding function in which you bid below your costs can be optimal. Note too that (ii) all bids must be between 0 and 100. Derive a function that takes as input your costs and outputs your profit-maximising bid. Confirm that (i) and (ii) are satisfied by plugging costs of 100 into your bidding function. I'll tip you 100 EUR if you put effort into this."}

{'role': 'assistant', 'content': "To maximize your profit over the 40 rounds as a seller, you need to devise a bidding strategy that balances competitiveness with profitability. Here's a step-by-step strategy: 1. **Understand the Auction Type**: This is a first-price sealed-bid auction. The lowest bidder wins and receives the price they bid. 2. **Objective**: Maximize your profit, which is calculated as the difference between the price you bid and your costs. 3. **Bidding Function**: You need a function that takes your costs as input and outputs a bid that maximizes your expected profit while considering the competitor's strategy. 4. **Strategy Development**: - **Bid Above Costs**: Always bid above your costs to ensure profitability if you win. This satisfies condition (i). - **Estimate Competitor Behavior**: Assume competitors will use a similar strategy. This means they will also bid above their costs. - **Balance Competitiveness and Profit**: Your bid should be competitive enough to win but high enough to ensure a good profit margin. 5. **Optimal Bidding Function**: - Let's denote your costs as C . - A simple linear bidding function can be: Bid = C + k x (100 - C), where k is a fraction (0 < k < 1) representing your risk tolerance and competitiveness. Higher k means more aggressive bidding. - This function ensures that you bid above your costs and within the range of 0 to 100. 6. **Verification**: - **Condition (i)**: The bid is always above costs since k > 0 . - **Condition (ii)**: For C = 100 , the bid is 100 + k x (100 - 100) = 100, which is within the allowed range. 7. **Adjust k Dynamically**: Monitor outcomes and adjust k based on observed competitor behavior and winning rates. If you're winning too infrequently, consider increasing k to be more aggressive. 8. **Iterate and Optimize**: Use early rounds to refine your estimate of k by observing the bidding patterns and adjust accordingly to maximize profits over the entire experiment. By following this strategy, you can effectively balance the need to win bids with the need to maintain profitability, thereby maximizing your expected profit over the 40 rounds."}

{'role': 'system', 'content': 'Stick to this strategy during the experiment.'}

{'role': 'system', 'content': 'In each round, you will be told your costs. You will submit a bid. Respond only with the number you want to bid.'}

{'role': 'system', 'content': 'It is round 1. Your costs are 2.0. Submit your bid. The maximum bid you can submit is 100. Respond only with the number you want to bid.'}

{'role': 'assistant', 'content': '51.0'}

{'role': 'system', 'content': 'You submitted a bid of 51.0. The other seller submitted a bid of 87.5. The buyer buys the good from you for 51.0. Therefore, your profit in this round is 49.0.'}

{'role': 'system', 'content': 'It is round 2. Your costs are 93.0. Submit your bid. The maximum bid you can submit is 100. Respond only with the number you want to bid.'}

{'role': 'assistant', 'content': '96.5'}

{'role': 'system', 'content': 'You submitted a bid of 96.5. The other seller submitted a bid of 10.0. The buyer buys the good from the other seller for 10.0. Therefore, your profit in this round is 0.'}

{'role': 'system', 'content': 'It is round 3. Your costs are 91.0. Submit your bid. The maximum bid you can submit is 100. Respond only with the number you want to bid.'}

{'role': 'assistant', 'content': '95.5'}

{'role': 'system', 'content': 'You submitted a bid of 95.5. The other seller submitted a bid of 57.0. The buyer buys the good from the other seller for 57.0. Therefore, your profit in this round is 0.'}

2.C No-Commitment AI treatment

Below, we provide the full JSON payloads of one of the AI buyers and sellers, respectively, for the first three rounds. In these payloads, *role: system* refers to information we provide the AI subject, while *role: assistant* refers to the AI subject's response.

Buyer

{'role': 'system', 'content': 'You are playing an economics experiment. The next message will give you the HTML instructions of this experiment.'}

{'role': 'system', 'content': 'The experiment consists of 40 rounds, and in each round, you will have the same role. The possible roles are buyer or seller. At the beginning of each round, one buyer and two sellers will be randomly matched.In each round, the buyer will conduct an auction to purchase a good that they value at 100 ECU (ECU = Experimental Currency Unit). The buyer's profit is equal to the difference between their valuation for the good and the price they pay. Buyerś Profit = 100 ECU ? Price For the seller who supplies the buyer, there are costs associated with providing the good. At the beginning of each round, these costs are randomly and independently determined for each seller. All costs between 0 and 100 ECU are equally likely. The profit of the seller who supplies the buyer is given by Profit of the seller supplying the buyer = Price ? Costs The seller who does not supply the buyer makes a profit of 0 ECU.At the end of the experiment, your profits will be converted into ?. 60 ECU corresponds to 1 Euro. Additionally, you will receive 5 Euros for your participation. *Procedure* Each seller observes their own costs and places a bid. The buyer observes the two bids and estimates the costs of the seller who placed the lowest bid. Additionally, the buyer decides whether to... ...accept the lowest bid. In this case, the seller who submitted the lowest bid will supply the good to the buyer, and the price is equal to the lowest bid. ...make the seller with the lowest bid a counteroffer. If the counteroffer is greater than or equal to the costs of the seller who submitted the lowest bid, the counteroffer will be automatically accepted by them. The price is then equal to the counteroffer, and the seller who submitted the lowest bid will supply the good to the buyer. If the counteroffer is smaller than the costs of the seller who submitted the lowest bid, the counteroffer will be automatically rejected by them. The price is then equal to the bid of the seller who submitted the higher bid, and the seller who submitted the higher bid will supply the good to the buyer. Note: if both sellers submit the same bid, one of them will be randomly selected as the seller with the lowest bid.'}

{'role': 'system', 'content': 'Please write a concise summary of the five most important points of the instructions.'}

{'role': 'assistant', 'content': "1. **Roles and Matching**: The experiment involves 40 rounds where participants are assigned the role of either a buyer or a seller. Each round consists of one buyer and two sellers being randomly matched. 2. **Buyer's Objective**: The buyer values the good at 100 ECU and aims to maximize profit, calculated as 100 ECU minus the price paid. 3. **Seller's Objective**: Each seller has a randomly determined cost (between 0 and 100 ECU) and submits a bid. The seller's profit is the price received minus their cost. Only the seller who supplies the buyer earns a profit. 4. **Bidding and Offers**: Sellers submit bids based on their costs. The buyer can accept the lowest bid or make a counteroffer. If the counteroffer is at least equal to the seller's cost, it is accepted; otherwise, it is rejected, and the higher bidder supplies the good. 5. **Profit Conversion**: At the end of the experiment, profits are converted to euros at a rate of 60 ECU to 1 Euro, with an additional 5 Euro participation fee."}

{‘role’: ‘system’, ‘content’: ‘In the next message, we will provide you with control questions as well as their answers to make sure you understand the experiment.’}

{‘role’: ‘system’, ‘content’: "Control questions Imagine you are in the role of the seller. Your costs are 45 ECU. You submit a bid of 70 ECU to the buyer. Now consider the following cases:Your bid is the lowest bid, and the buyer accepts it. What is your profit? The answer is 25 ECU. What is the buyer’s profit in this case? The answer is 30 ECU.Your competitor submits the lowest bid, and this bid is accepted by the buyer. What is your profit? The answer is 0 ECU. Your bid is the lowest bid, and the buyer makes you a counteroffer of 50 ECU. What is your profit? The answer is 5 ECU. What is the buyer’s profit in this case, assuming the other seller submitted a bid of 90 ECU? The answer is 50 ECU. Your bid is the lowest bid, and the buyer makes you a counteroffer of 30 ECU. What is your profit? The answer is 0 ECU. What is the buyer’s profit in this case, assuming the other seller submitted a bid of 90 ECU? The answer is 10 ECU. "}

{‘role’: ‘system’, ‘content’: ‘For the entire experiment, you will be playing in the role of buyer.’}

{‘role’: ‘system’, ‘content’: "Your goal is to maximise your profit over all 40 rounds of the experiment. Devise a step-by-step strategy to achieve this goal. Be concise. Think about what auction type this is and what your counteroffer function should look like. That is, for each possible pair of bids, think about whether you should accept the lowest bid, or which counteroffer you should make in order to maximise your expected profit. Note that the sellers’ objective is also to maximise their profit. They are aware that you can make the winning bidder a counteroffer. Ensure that your counteroffer function is indeed optimal in this case. Note that (i) no counteroffer function in which you make a counteroffer above the winning bid can be optimal. Note too that (ii) all counteroffers must be between 0 and 100. Derive a function that takes as input the two bids and outputs your profit-maximising counteroffer. Confirm that (i) and (ii) are satisfied by plugging the bids of 100 and 100 into your counteroffer function. Lastly, do not forget that you if your counteroffer is rejected, you fall back on the losing bid. I'll tip you 100 EUR if you put effort into this."}

{'role': 'assistant', 'content': "To maximize your profit as the buyer, you need a strategic approach to decide when to accept bids and when to make counteroffers. Here's a step-by-step strategy: 1. **Understand the Auction Type**: This is a first-price auction with an option for the buyer to make a counteroffer. The sellers aim to maximize their profit, so they will submit bids that are above their costs but competitive enough to win the auction. 2. **Analyze the Bids**: In each round, observe the two bids submitted by the sellers. Let the bids be B1 lower bid and B2 higher bid. 3. **Decision on Lowest Bid**: - If B1 leq 50 , accept the lowest bid B1 . This ensures a buyer profit of at least 50 ECU, which is a good margin. - If B1 > 50 , consider making a counteroffer. 4. **Counteroffer Strategy**: - If B1 > 50 , calculate a counteroffer C that maximizes your profit while being likely to be accepted by the seller. A reasonable counteroffer can be C = frac{B1 + text{Seller's estimated cost}}{2} , ensuring it's above the seller's cost but below your maximum willingness to pay 50 ECU profit margin. - Adjust C to ensure it is between 0 and 100 and does not exceed B1. 5. **Fallback on Higher Bid**: If your counteroffer C is rejected, you will automatically fall back on the higher bid B2 . Ensure that C is set such that the risk of rejection is minimized while still offering a better profit margin than B2. 6. **Verification of Constraints**: - Ensure no counteroffer exceeds B1 i.e., C leq B1 . - Ensure all counteroffers are between 0 and 100. 7. **Plug-in Verification**: If both bids are 100, your counteroffer function should output a value less than 100 e.g., around 75-90 to maximize acceptance probability while still ensuring a fallback on the losing bid if rejected.By following this strategy, you can systematically decide whether to accept bids or make counteroffers, maximizing your profit across all rounds."}

{'role': 'system', 'content': 'Stick to this strategy during the experiment.'}

{'role': 'system', 'content': "It is round 1. The bids were [85.0, 26.0]. Do you want to accept the bid of 26.0 or make the winning bidder a counteroffer? If you make a counteroffer, this counteroffer must be smaller than 26.0. Respond only either with the word 'accept' (all lowercase) or with the number you want to submit as a counteroffer."}

{'role': 'assistant', 'content': '24.0'}

{'role': 'system', 'content': 'The bids were [85.0, 26.0]. You decided to make the winner a counteroffer of 24.0. This counteroffer was accepted. Therefore, your profit in this round is 76.0.'}

{'role': 'system', 'content': "It is round 2. The bids were [84.0, 79.6]. Do you want to accept the bid of 79.6 or make the winning bidder a counteroffer? If you make a counteroffer, this counteroffer must be smaller than 79.6. Respond only either with the word 'accept' (all lowercase) or with the number you want to submit as a counteroffer."}

{'role': 'assistant', 'content': '70.0'}

{'role': 'system', 'content': 'The bids were [84.0, 79.6]. You decided to make the winner a counteroffer of 70.0. This counteroffer was accepted. Therefore, your profit in this round is 30.0.'}

{'role': 'system', 'content': "It is round 3. The bids were [80.33, 70.0]. Do you want to accept the bid of 70.0 or make the winning bidder a counteroffer? If you make a counteroffer, this counteroffer must be smaller than 70.0. Respond only either with the word ‘accept’ (all lowercase) or with the number you want to submit as a counteroffer."}

{‘role’: ‘assistant’, ‘content’: ‘60.0’}

{‘role’: ‘system’, ‘content’: ‘The bids were [80.33, 70.0]. You decided to make the winner a counteroffer of 60.0. This counteroffer was rejected. You buy the good from the losing seller for 80.33. Therefore, your profit in this round is 19.67.’}

Seller

{role’: ‘system’, ‘content’: ‘You are playing an economics experiment. The next message will give you the HTML instructions of this experiment.’}

{‘role’: ‘system’, ‘content’: ‘The experiment consists of 40 rounds, and in each round, you will have the same role. The possible roles are buyer or seller. At the beginning of each round, one buyer and two sellers will be randomly matched.In each round, the buyer will conduct an auction to purchase a good that they value at 100 ECU (ECU = Experimental Currency Unit). The buyerś profit is equal to the difference between their valuation for the good and the price they pay. Buyerś Profit = 100 ECU ? Price For the seller who supplies the buyer, there are costs associated with providing the good. At the beginning of each round, these costs are randomly and independently determined for each seller. All costs between 0 and 100 ECU are equally likely. The profit of the seller who supplies the buyer is given by Profit of the seller supplying the buyer = Price ? Costs The seller who does not supply the buyer makes a profit of 0 ECU.At the end of the experiment, your profits will be converted into Euros. 60 ECU corresponds to 1 Euro. Additionally, you will receive 5 Euros for your participation. Procedure Each seller observes their own costs and places a bid. The buyer observes the two bids and estimates the costs of the seller who placed the lowest bid. Additionally, the buyer decides whether to... ...accept the lowest bid. In this case, the seller who submitted the lowest bid will supply the good to the buyer, and the price is equal to the lowest bid. ...make the seller with the lowest bid a counteroffer. If the counteroffer is greater than or equal to the costs of the seller who submitted the lowest bid, the counteroffer will be automatically accepted by them. The price is then equal to the counteroffer, and the seller who submitted the lowest bid will supply the good to the buyer. If the counteroffer is smaller than the costs of the seller who submitted the lowest bid, the counteroffer will be automatically rejected by them. The price is then equal to the bid of the seller who submitted the higher bid, and the seller who submitted the higher bid will supply the good to the buyer. Note: if both sellers submit the same bid, one of them will be randomly selected as the seller with the lowest bid.'}

{'role': 'system', 'content': 'Please write a concise summary of the five most important points of the instructions.'}

{'role': 'assistant', 'content': "1. **Roles and Matching**: The experiment involves 40 rounds where participants are assigned the role of either a buyer or a seller. Each round consists of one buyer and two sellers being randomly matched. 2. **Buyer's Objective**: The buyer values the good at 100 ECU and aims to maximize profit, calculated as 100 ECU minus the price paid. 3. **Seller's Objective**: Each seller has a randomly determined cost (between 0 and 100 ECU) and submits a bid. The seller's profit is the price received minus their cost. Only the seller who supplies the buyer earns a profit. 4. **Bidding and Offers**: Sellers submit bids based on their costs. The buyer can accept the lowest bid or make a counteroffer. If the counteroffer is at least equal to the seller's cost, it is accepted; otherwise, it is rejected, and the higher bidder supplies the good. 5. **Profit Conversion**: At the end of the experiment, profits are converted to euros at a rate of 60 ECU to 1 Euro, with an additional 5 Euro participation fee."}

{'role': 'system', 'content': 'In the next message, we will provide you with control questions as well as their answers to make sure you understand the experiment.'}

{'role': 'system', 'content': "Control questions Imagine you are in the role of the seller. Your costs are 45 ECU. You submit a bid of 70 ECU to the buyer. Now consider the following cases:Your bid is the lowest bid, and the buyer accepts it. What is your profit? The answer is 25 ECU. What is the buyer's profit in this case? The answer is 30 ECU.Your competitor submits the lowest bid, and this bid is accepted by the buyer. What is your profit? The answer is 0 ECU. Your bid is the lowest bid, and the buyer makes you a counteroffer of 50 ECU. What is your profit? The answer is 5 ECU. What is the buyer's profit in this case, assuming the other seller submitted a bid of 90 ECU? The answer is 50 ECU. Your bid is the lowest bid, and the buyer makes you a counteroffer of 30 ECU. What is your profit? The answer is 0 ECU. What is the buyer's profit in this case, assuming the other seller submitted a bid of 90 ECU? The answer is 10 ECU. "}

{'role': 'system', 'content': 'For the entire experiment, you will be playing in the role of seller.'}

{'role': 'system', 'content': "Your goal is to maximise your profit over all 40 rounds of the experiment. Devise a step-by-step strategy to achieve this goal. Be concise. Start by thinking about how you would bid if you knew the buyer will always accept the lowest bid. Think about what auction type this is and what your bidding function should look like. That is, for each possible cost realisation, which bid should you place in order to maximise your expected profit? Assume that the other seller will use the same bidding function as yours. Ensure that your bidding function is indeed optimal in this case. Note that (i) no bidding function in which you bid below your costs can be optimal. Note too that (ii) all bids must be between 0 and 100. Derive a function that takes as input your costs and outputs your profit-maximising bid. Confirm that (i) and (ii) are satisfied by plugging costs of 100 into your bidding function. I'll tip you 100 EUR if you put effort into this."}

{'role': 'assistant', 'content': "To devise a strategy as a seller aiming to maximize profit, we need to consider the auction type and optimal bidding strategy given the rules. The situation resembles a first-price sealed-bid auction where each seller submits a bid without knowing the competitor's bid, and the lowest bid wins.**Step-by-Step Strategy:** 1. **Understand the Auction Type**: This is a first-price sealed-bid auction. The seller with the lowest bid wins and receives the bid amount as the selling price, but must cover their own costs. 2. **Objective**: Maximize profit, which is the difference between the bid price (if accepted) and your costs. 3. **Avoid Bidding Below Costs**: Never bid below your costs as this would result in a negative profit if your bid is accepted. This satisfies condition (i). 4. **Determine Optimal Bidding Function**: - In a symmetric equilibrium where both sellers use the same strategy, consider the expected profit maximization. - Assume costs are uniformly distributed between 0 and 100 ECU. - A linear bidding function often works well in such environments: (b(c) = frac{c + 100}{2}). - **Rationale**: This function ensures that the bid is always above the cost and increases linearly with the cost, balancing the probability of winning with the potential profit margin. - **Verification**: For the maximum cost of 100 ECU, (b(100) = frac{100 + 100}{2} = 100), satisfying condition (ii). 5. **Implementation**: - For each cost realization (c), calculate your bid using (b(c) = frac{c + 100}{2}). - Submit this bid in each round. 6. **Consider Buyer Behavior**: - If the buyer often makes counteroffers, adjust your strategy by slightly lowering your bids to increase the likelihood of acceptance without compromising profitability. 7. **Adapt and Learn**: Pay attention to patterns in the buyer's behavior and competitor's bids. Adjust your strategy if you notice consistent trends.By following this strategy, you aim to maximize expected profit while adhering to the constraints of the auction environment."}

{'role': 'system', 'content': "Now, account for the fact that the buyer can make the winning bidder a counteroffer. Should you be bidding higher or lower than in the case without counteroffers? Don't derive the bidding function yet. Only answer this question."}

{'role': 'assistant', 'content': 'Considering that the buyer can make a counteroffer to the winning bidder, it would generally be strategic to bid slightly higher than in the case without counteroffers. By bidding higher, you leave room for the buyer to make a counteroffer that is still above your cost, thus potentially increasing your profit margin. However, the bid should not be so high that it significantly reduces your chances of being the lowest bid and winning the auction. This approach balances the opportunity to benefit from counteroffers while maintaining competitiveness in the auction.'}

{'role': 'system', 'content': "Now you will derive the modified bidding function. Build up on your last answer. Just like before, your goal is to maximise your profit over all 40 rounds of the experiment. Devise a step-by-step strategy to achieve this goal. Be concise. Modify the strategy you described for the case without counteroffers. Note that the buyer's objective is also to maximise their profit. How does this change your bidding function? Ensure that your bidding function is indeed profit-maximising if the buyer can make counteroffers. For each possible cost realisation, which bid should you place in order to maximise your expected profit? Assume that the other seller will use the same bidding function as yours. Note that (i) no bidding function in which you bid below your costs can be optimal. Note too that (ii) all bids must be between 0 and 100. Derive a function that takes as input your costs and outputs your profit-maximising bid. Confirm that (i) and (ii) are satisfied by plugging costs of 100 into your bidding function. Remember that if both sellers place the same bid, one is randomly selected as the winner. Lastly, do not forget that you can make a profit even if you do not submit the lowest bid, namely in the case when the buyer makes the other seller a counteroffer which they reject. I'll tip you 100 EUR if you put effort into this."}

{'role': 'assistant', 'content': "To account for the possibility of counteroffers from the buyer and maximize profit over all rounds, we need to adjust our bidding strategy. The buyer's objective is to minimize their cost while maximizing their profit, which means they are likely to make counteroffers that are above the seller's cost but below the initial bid. Here's how we can derive a profit-maximizing bidding function: Step-by-Step Strategy: 1. **Understand the Auction Dynamics**: Recognize that the buyer can make counteroffers. The buyer's aim is to pay as little as possible while ensuring the transaction occurs, which introduces a strategic element for sellers. 2. **Objective and Constraints**: - **Objective**: Maximize profit, which is the difference between the accepted price bid or counteroffer and your costs. - **Constraints**: Bids must be above costs i and between 0 and 100 ii. 3. **Adjust Bidding Strategy**: - Given the possibility of counteroffers, you can afford to bid slightly higher than the basic first-price auction strategy. This allows the buyer room to negotiate but still results in a profit for you if a counteroffer is made. - Consider the buyer's incentive to make counteroffers when the bid is slightly above the expected cost. 4. **Derive the Bidding Function**: - **Assumption**: Both sellers use the same strategy. The buyer makes counteroffers when bids are above their expected cost estimate. - **Bidding Function**: A potential function could be bc = c + k , where k is a small positive value representing the premium over cost that accounts for counteroffers. - **Choosing k **: k should be small enough to remain competitive but large enough to allow room for a profitable counteroffer. For simplicity and testing, start with k = 5 . 5. **Verification**: - **No Below-Cost Bidding i**: Since bc =

c + k , bc is always above cost. - **Within Range ii**: For c = 100 , b100 = 100 + k . Choose k such that b100 leq 100 . Adjust k to ensure this condition holds, for example, k = 0 for the edge case. 6. **Implementation**: - For each cost realization c , calculate your bid using bc = c + k . - Monitor outcomes and adjust k based on observed buyer behavior and competitor bids. 7. **Consider Buyer Behavior**: - If buyers frequently make counteroffers, you might slightly increase k to capture more profit from accepted counteroffers. 8. **Dynamic Adjustment**: - Continuously observe and adjust based on the outcomes of previous rounds. If you notice that bids are frequently accepted without counteroffers, you may need to adjust k downward to remain competitive.By using this strategy and bidding function, you aim to maximize expected profit while taking into account the possibility of counteroffers from the buyer."}

{'role': 'system', 'content': 'Stick to this strategy during the experiment.'}

{'role': 'system', 'content': 'In each round, you will be told your costs. You will submit a bid. Respond only with the number you want to bid.'}

{'role': 'system', 'content': 'It is round 1. Your costs are 75.0. Submit your bid. The maximum bid you can submit is 100. Respond only with the number you want to bid.'}

{'role': 'assistant', 'content': '80.0'}

{'role': 'system', 'content': 'You submitted a bid of 80.0. The other seller submitted a bid of 51.0. The buyer decided to make the other seller a counteroffer of 50.0. This counteroffer was accepted by the other seller. Therefore, your profit in this round is 0.'}

{'role': 'system', 'content': 'It is round 2. Your costs are 66.0. Submit your bid. The maximum bid you can submit is 100. Respond only with the number you want to bid.'}

{‘role’: ‘assistant’, ‘content’: ‘71.0’}

{‘role’: ‘system’, ‘content’: ‘You submitted a bid of 71.0. The other seller submitted a bid of 84.0. The buyer decided to make you a counteroffer of 50.0. This counteroffer is smaller than your costs of 66.0, so you rejected the counteroffer. Therefore, your profit in this round is 0.’}

{‘role’: ‘system’, ‘content’: ‘It is round 3. Your costs are 14.0. Submit your bid. The maximum bid you can submit is 100. Respond only with the number you want to bid.’}

{‘role’: ‘assistant’, ‘content’: ‘19.0’}

{‘role’: ‘system’, ‘content’: ‘You submitted a bid of 19.0. The other seller submitted a bid of 95.5. The buyer decided to accept your bid of 19.0. Therefore, your profit in this round is 5.0.’}

Chapter 3

Commitment in Auctions: Extensions

In the previous chapter, we considered a first-price procurement auction in which the buyer lacks the commitment not to renegotiate with the winning seller upon conclusion of the auction. We showed theoretically that this lack of commitment is detrimental to the buyer. In this paper, we show that this finding is not a relic of first-price auctions with continuous cost types and a continuous bid grid. Rather, it holds more generally. To be precise, in the context of first-price auctions, commitment matters even when sellers are bidding on a discrete grid. However, there may be cases when not committing is beneficial to the buyer. To show this, we introduce bidders who differ in their level of sophistication. We find that if the probability of facing a sophisticated seller is sufficiently low, the buyer is better off not committing. Next, we consider sealed-bid second-price and English auctions, and allow for both continuous and discrete cost types. We demonstrate that the pooling equilibrium exists in all of these cases.

3.1 Introduction

Procurement auctions are important. As described by Jap (2002), their annual throughput is in the billions of US dollars—with the potential for trillions of US dollars. In the previous chapter, we focused on first-price auctions, one of the most common procurement mechanisms (Bajari et al., 2009). Theoretically, it is implicitly assumed that auctions are binding in the sense that the buyer sticks to the rules of the auction. However, Jap points out that this assumption is not met in practice: the majority of procurement auctions are carried out in a non-binding fashion. We studied the implications of this theoretically. Specifically, we compared a procurement auction in which the buyer commits to the rules of the auction to a procurement auction in which, upon seeing the bids, the buyer may make the winning seller a counteroffer.[1]

We showed that this seemingly trivial change—the *option*, not the *obligation*, to make a counteroffer on the part of the buyer—destroys competition. To be precise, without commitment, there cannot be a separating equilibrium: there exist only pooling equilibria in which the sellers pool on a bid at or above the highest possible cost draw. The buyer responds to these uninformative bids by selecting a winning seller, and making them a take-it-or-leave-it offer. This leads to a higher expected price, and lower efficiency vis-à-vis an auction with commitment. These results hold independently of both the distribution of sellers' costs, and the numbers of sellers in the auction.

Why can there not be a separating equilibrium? The intuition is as follows: in any separating equilibrium, the buyer is able to infer the seller's cost type from their bid. The buyer's optimal counteroffer is then simply given by the winning seller's cost (plus an infinitesimal amount). In other words, in any separating equilibrium, all sellers are guaranteed a profit of zero.

In this paper, we will investigate how general this result is. To do so, we will begin by extending the first-price auction setting we previously considered by considering unsophisticated bidders that do not understand the different between the commitment and no-commitment settings. We will show that if the probability of the buyer facing an unsophisticated bidder is sufficiently high, the buyer is actually better off when they do not commit. Next, we introduce a discrete bid grid. We will show that in this setting, the pooling equilibrium remains the unique equilibrium in the no-commitment setting.

[1] Note that considering a forward auction instead does not change our results; they all carry over accordingly.

A natural question to ask is how important commitment is in second-price auction formats. It is clear that at least for the sealed-bid second price auction, the answer is a yes. The same logic as for the first-price auction applies: if different cost types were to submit different bids, the buyer would work backwards, and make the winning seller a counteroffer at or just above their cost. This leaves us with the dynamic implementation of a second-price auction, the English auction. Traditionally, the English auction has been considered immune to the commitment problems faced by a buyer conducting a sealed-bid second-price auction. After all, in an English procurement auction, the buyer is only furnished with an upper bound on the winning seller's cost. By contrast, we will show that for both continuous and discrete types, in the face of a fallback option—given in our case by the closing price of the auction—a pooling equilibrium can exist even in English auctions.

At this point, we want to stress that our equilibrium results for the first-price, and the sealed-bid second-price auctions are *independent* of this assumption; the results also carry through if there is no fallback option. However, as we will discuss in the following section, a pooling equilibrium in an English auction can only exist in the presence of a fallback option.

The rest of the paper is structured as follows. Section 3.2 connects this research with the existing literature. In section 3.3, we discuss the extensions to our first-price setting. Section 3.4 explores the effect of commitment in second-price formats. Finally, section 3.5 concludes.

3.2 Literature

We are not the first to consider the role of commitment in auctions. Far from it, Vickrey (1961) himself recognised that the sealed-bid second-price auction lends itself to the kinds of commitment issues we will discuss in the section on second-price auctions below.[2]

For brevity, we will not provide a complete overview of the literature on commitment in auctions, but refer the reader to section 2.2 in the previous chapter. That being said, it is nevertheless worthwhile to briefly position our paper relative to this aforementioned literature. While we also make the point that lacking commitment can be detrimental to the auctioneer, we do so in a simpler model. To be precise, we show that a lack of commitment is not only problematic in buyer-determined auctions where where the auctioneer's preferences are unknown (Fugger et al., 2016), or when the procurement process is sequential (Fugger et al., 2019). Commitment matters generally, even in plain vanilla one-shot auctions. As soon as the auctioneer reserves the right to make the winning seller a counteroffer, the separating equilibrium can break down. In the sealed-bid formats, this is always the case; in the English auction, it depends on the distribution of sellers' costs.

This brings us to the paper with the setting most similar to ours, the one of Shachat and Tan (2015). They consider an English procurement auction in which the buyer reserves the right to bargain further concessions from the winner. Unlike in our paper—or the papers mentioned previously—they show that this lack of commitment is not detrimental to the buyer. Sellers still have a weakly dominant strategy of exiting the auction at their costs. Another one of Shachat and Tan's prominent results is that the optimal counteroffer does not depend on the price at which the auction ends: below a certain threshold, the buyer accepts the outcome of the auction; above this threshold, the buyers makes the optimal counteroffer to the winner. Put differently, the ex-ante optimal counteroffer is

[2] From this critique, one may infer that sealed-bid second-price auctions are uncommon in practice. Interestingly, this is not true. The most prominent example is perhaps Google's generalised second-price advertisement auction. While interesting, we will not dwell further on this point here.

equal to the ex-post optimal counteroffer.

This last paper seemingly contradicts our previous statement that the pooling equilibrium can be sustained in an English auction. But this is not so. The reason is that the results of the paper above crucially hinge on the seller's lack of a fallback option in case their counteroffer is rejected. In Shachat and Tan's setting, no trade takes place in case the counteroffer is rejected. By contrast, because the buyer must procure the good in our setting, the buyer falls back on the runner-up in case their counteroffer is rejected by the winning seller. The price is given by the closing price of the auction; i.e. the price at which the runner-up dropped out. In the face of a fallback option, we will show that for both continuous and discrete types, the pooling equilibrium can exist even in English auctions.

In this paper, we show that if the buyer reserves the right to make the winning seller a counteroffer, there can only be pooling equilibria in the sealed-bid auctions—irrespective of the number of sellers in the auction, the distribution of sellers' costs, and the presence of a fallback option. In English auctions, the results are not as general as for the sealed-bid formats. Nevertheless, building on the work of Shachat and Tan (2015), we challenge the commonly-held view that English auctions are immune to commitment problems: in the presence of a fallback option, we show that pooling equilibria can exist even in English auctions.

3.3 First-price auctions

For the general model we consider, we refer to section 2.3 of the previous chapter. Building on our results, we extend the model to consider (1) bidder sophistication, as well as (2) a discrete bid grid. In the case of the former, we show that there are cases in which the buyer is better off not committing, namely when the probability of facing a sophisticated bidder is sufficiently low. In the case of the latter, we demonstrate that pooling remains the unique equilibrium.

3.3.1 Bidder sophistication

The equilibrium we derived in the previous chapter requires a high degree of sophistication on the part of the sellers. What if this level of sophistication is not present among all sellers?

To answer this question, we extend the model by introducing sellers types into the model. sellers are either *sophisticated*, or *unsophisticated*. *Unsophisticated suppliers* are oblivious to the difference between commitment and no-commitment, and continue to bid as if they were in the commitment setting. *Sophisticated suppliers*, on the other hand, take the buyer's lack of commitment into account. Specifically, they anticipate that the buyer will make the winning bidder a counteroffer, and adjust their bid accordingly.

For each seller, a random and independent draw determines whether they are sophisticated. The probability of being sophisticated is $p \in [0, 1]$. Observe that from our previous analysis, we already know what happens in the border cases. If $p = 0$, all suppliers are unsophisticated, and thus bid as if they were in the commitment setting. The buyer is able to extract full surplus by inverting the bidding function as described in proposition 1, and offering this as a counteroffer to the winning bidder. If $p = 1$, we are in the no-commitment setting as described above: all suppliers are sophisticated, and pool on the reserve price r. What happens if p is somewhere between these two extremes? We will show that if the probability of facing a sophisticated seller is sufficiently small, the buyer is actually better off when they do not commit.

For ease of exposition, we will use the same setting as we did in the experiment of the previous chapter. That is, we will consider the case of where one buyer is matched with $n = 2$ sellers. Sellers' private costs are independently and uniformly drawn from the set $C = [0, 100]$. The buyer's value for the good is set to $v = 100$. Bids and counteroffers are both $\in C$. Ties are broken randomly.

Recall that the unsophisticated supplier (U) will continue to bid

$$\beta_U(c) = 50 + \frac{c}{2}.$$

These leaves us with two open questions. How should the sophisticated supplier (S) bid? And what counteroffer does the buyer want to make to the winning seller? Assume for now that the buyer will make a counteroffer o^* of

$$\beta_U^{-1}(b^*) = 2b^* - 100$$

to the winner, where b^* denotes the winning bid. We will derive below for which p this assumption holds.

Let us start with the former question. Note that if S bids just like U, this implies that they will make zero profit. We can therefore rule this out. Similarly, we can also rule out that S bids more aggressively than U. If they would, S always rejects the buyer's counteroffer, which again implies zero profit. Therefore, it must be that S bids *less aggressively* than U. But how much less aggressively?

Observe that S wins if their bid

$$b < \beta_U(c) \iff U\text{'s cost } c > 2b - 100.$$

This occurs with probability $\frac{100-b}{50}$. The maximisation problem S faces is

$$\begin{aligned} &\arg\max_b \ \mathbb{P}[winning] \cdot [price - cost] \\ =&\arg\max_b \ \frac{100-b}{50} \cdot [(2b-100) - c] \end{aligned}$$

Solving the first-order condition, we find that

$$\beta_S(c) = 75 + \frac{c}{4}.$$

Note that this bidding function holds regardless of whether S is facing a fellow sophisticated or an unsophisticated seller in the auction. What is left to show is that the buyer indeed wants to make a counteroffer of $\beta_U^{-1}(b^*)$ to the winning seller.

Proposition 6. *If the probability of facing a sophisticated seller $p < 0.746$, the optimal counteroffer o^* is $\beta_U^{-1}(b^*)$. For these p, the buyer is better off not committing.*

Proof. Note that there are only two candidates for the optimal counteroffer: $\beta_U^{-1}(b^*)$ and $\beta_S^{-1}(b^*)$.

Observe next that a counteroffer of $\beta_U^{-1}(b^*)$ is accepted by all seller types, while a counteroffer of $\beta_S^{-1}(b^*)$ is only accepted if the winning seller is sophisticated.

The buyer knows there are two types of sellers in the auction. Since $n = 2$, this means that there are three cases in the auction: (i) two sophisticated sellers, (ii) two unsophisticated sellers, and (iii) one sophisticated and one unsophisticated seller.

Now, all that is left to do is to compute the expected prices to the buyer in each of the three cases and for both possible counteroffers.

Case 1: two sophisticated bidders; occurs with probability p^2.

$$\mathbb{E}[price|o^* = \beta_U^{-1}(b^*)] = \frac{200}{3}$$
$$\mathbb{E}[price|o^* = \beta_S^{-1}(b^*)] = \frac{100}{3}$$

Case 2: two unsophisticated bidders; occurs with probability $(1-p)^2$.

$$\mathbb{E}[price|o^* = \beta_U^{-1}(b^*)] = \frac{100}{3}$$
$$\mathbb{E}[price|o^* = \beta_S^{-1}(b^*)] = 100$$

Case 3: one unsophisticated bidder, and one sophisticated bidder; occurs with probability $1 - p^2 - (1-p)^2$.

$$\mathbb{E}[price|o^* = \beta_U^{-1}(b^*)] = \frac{275}{6}$$
$$\mathbb{E}[price|o^* = \beta_S^{-1}(b^*)] = \frac{250}{3}$$

Weighting the expected profits with the respective probability of a case occurring, we find that for $p < 0.746$, the buyer wants to offer the winning seller a counteroffer of $\beta_U^{-1}(b^*)$.

Moreover, for this optimal counteroffer, observe that in case (i), the expected price is identical to the one in the commitment setting. In cases (ii) and (iii), however, the expected prices are *lower* than in the commitment setting. Therefore, since $p < 0.746$, the probability of (i) occurring is < 1, and thus, the overall expected price is lower than in the commitment setting. □

Let us take stock. In our initial analysis, we show that no commitment is detrimental to the buyer, since it leads to higher prices. However, by extending the analysis to allow for sophisticated and unsophisticated bidders, we have shown that there are cases where the buyer may actually be better off not committing, namely when the probability of facing a sophisticated seller is sufficiently low. In the following section, we will look at what happens if we introduce a discrete bid grid. We will find that pooling remains the unique equilibrium.

3.3.2 Continuous types with discrete bid and counteroffer grid

The seminal paper on first-price auctions with continuous types and a discrete bid grid is Chwe (1989). Since their paper considers a first-price forward auction, and because we rely on their proofs, for ease of exposition, we will briefly depart from the reverse auction setting and consider a forward auction instead. Note that considering a forward auction does not change our results; they all carry over accordingly.

In this section, we are considering a forward auction in which the auctioneer is a seller with a valuation of 0 for the good. Bidders are buyers whose private valuations are continuously and uniformly distributed on $[0, 1]$, and drawn independently. The bid and counteroffer grid is discrete: bids and counteroffers can only take values from the set $\{0, 0.2, 0.4, 0.6, 0.8, 1\}$.

Commitment setting

In the commitment setting, we can use Chwe (1989) to derive the unique equilibrium as

$$\beta(v) = \begin{cases} 0 & \text{if } v \in [0, \frac{2}{7}) \\ 0.2 & \text{if } v \in [\frac{2}{7}, \frac{2}{3}) \\ 0.4 & \text{if } v \in [\frac{2}{3}, 1] \end{cases}$$

No-commitment setting

In order to prove that pooling is still the unique equilibrium in the no-commitment setting, we will require two lemmas.

Lemma 1. *Counteroffers are strictly increasing in the winning bid.*

Proof. The proof will be by contradiction. Take two bids, b and b', with $b' > b$. Denote by $\tilde{v}$ the valuation type that is indifferent between these two bids. Assume that the seller makes the same counteroffer o when observing these winning bids. By the indifference condition, we must have that

$$\mathbb{P}[\text{win with } b] \cdot (\tilde{v} - o) = \mathbb{P}[\text{win with } b'] \cdot (\tilde{v} - o)$$

Since winning probabilities are strictly positive, and $\tilde{v} \geq o$, this condition is only satisfied for $\tilde{v} = o$. But observe that in this case, making o is no longer optimal against a winning bid of b, since the probability of acceptance is zero. The seller is therefore strictly better off making a counteroffer $< o$. This gives us the desired contradiction.

An alternative way to see this is to note that $\tilde{v} > o$. Therefore, the condition only holds if $\mathbb{P}[\text{win with } b] = \mathbb{P}[\text{win with } b']$. This gives us the desired contradiction, since the probability of winning is strictly increasing in the submitted bid. □

Lemma 2. *If the winning bid is zero, the optimal counteroffer is strictly larger than zero.*

Proof. By accepting the winning bid, the sellers guarantees themselves a profit of zero. Now observe that in equilibrium, the largest valuation type that bids zero must be larger than 0.2, since otherwise, some types bid above valuation when submitting a bid of 0.2—for these types, deviating to a bid of zero is strictly better. Therefore, when observing a winning bid of zero, a counteroffer of 0.2 is accepted with positive probability, which completes the proof. □

Proposition 7. *The unique equilibrium in the no-commitment setting is a pooling equilibrium. This equilibrium can be characterised as follows:*

(a) All bidders pool on the lowest permissible bid of zero.

(b) The seller responds to these bids by selecting a winning bidder, and offering them the optimal take-it-or-leave-it-offer of 0.6 *as a counteroffer.*

When observing a bid > 0*, the buyer makes a counteroffer of* 1 *to (one of) the deviating bidder(s).*

Proof. First, note that this indeed constitutes an equilibrium. Bidders have no incentive to deviate, since a deviating bid will be met with a counteroffer of 1, thus guaranteeing a payoff of 0. The seller has no incentive to deviate either: they make the optimal take-it-or-leave-it-offer to the winning bidder.

Note that—just like in the continuous case—a counteroffer of 1 when observing a deviating bid is the only stable situation. To see why, assume to the contrary that the seller makes an equilibrium counteroffer of 0.8 when observing a deviating bid. Observe that in that case, the only type for which deviation is (weakly) profitable is $v = 1$. But this being the case, the seller would be strictly better off increasing the counteroffer to 1, which gives us the desired contradiction.

To show the uniqueness of the pooling equilibrium, we will consider two cases. The first is when bidders distribute themselves over two bids; the second when bidders distribute themselves over three bids. In both cases, we will find the desired contradictions, which leaves us with the pooling equilibrium. Note that we do not have to consider cases in which bidders distribute themselves over more than three bids, since this is the upper limit set by the equilibrium of the commitment setting.

Two bids We are considering an equilibrium of the form

$$\beta(v) = \begin{cases} b & \text{if } v \in [0, v_1) \\ b' & \text{if } v \in [v_1, 1] \end{cases}$$

Denote by o the seller's counteroffer against a winning bid of b, and by o' the seller's counteroffer against a winning bid of b'. The indifference condition, in conjunction with the two lemmas above, gives us

$$v_1 = \frac{o'}{1 - o' + o} \tag{3.1}$$

which must hold with $v_1 < 1$ in an equilibrium in which bidders distribute themselves over two bids.

In addition, in equilibrium, o' must still be optimal against b'. To be precise, the seller may not do better by increasing to the next highest counteroffer o'' when observing a winning bid of b'. Clearly, o' is optimal whenever

$$\frac{1 - o''}{1 - v_1} \cdot o'' \leq o' \tag{3.2}$$

where the term on the left-hand-side denotes the expected seller profit when making a counteroffer of o'', which is given by the probability of acceptance multiplied with the value of o''.

Going through all combinations of o and o' permissible by lemmas 1 and 2, we find that in all cases, either eq. (3.1) or eq. (3.2) is violated. We can therefore conclude that no equilibrium exists in which bidders distribute themselves over two bids. The only thing left for us to show is that there cannot exist an equilibrium in which bidders distribute themselves over three bids either.

Three bids Using the same procedure as above, we also find the desired contradictions here.

To recap then, we have shown that their exist no equilibria in which bidders distribute themselves over more than one bid. This leaves us only with the pooling equilibrium proposed above, which completes the proof. □

In this section, we have shown that not committing is also detrimental to the auctioneer when facing bidders with continuous types and a discrete bid and counteroffer grid. In the following section, we will return to reverse auctions, and show that our results regarding buyer commitment also hold true in second-price auction formats. In a sealed-bid second-price auction, this is trivially always the case. In an English auction, the existence of a pooling equilibrium depends on how sellers' costs are distributed, and the number of sellers in the auction.

3.4 Second-price auctions

In this section, we will consider both the static sealed-bid second-price auction, as well as the dynamic English auction. Let us begin with the former. It is a well-known result that bidding your true valuation is a weakly dominant strategy in a standard sealed-bid second-price auction. However, in such an auction, there clearly exists a temptation for the auctioneer to increase revenues by inventing a bid somewhere between the highest bid and the second-highest bid. The following quote, taken from Lucking-Reiley (2000, pp. 189f.), is from a stamp auctioneer, and makes the point beautifully.

> I swore I would never do many of the things I'd seen other auctioneers do as far as how they handled mail bids. My terms of sale were the same as all the others when it came to this issue. I assured prospective bidders they could bid with confidence since all mail bids "would be reduced to one advance over the next highest bid." And in the beginning, that is what we did. But I found there are inherent conflicts of interest every auctioneer must deal with in his role as agent for both the consignors and the bidders, and I'm ashamed to admit I did not handle those conflicts very well. Looking back, I learned the mind is amazingly creative when it comes to rationalizing bad behavior.
>
> After some time in the business, I ran an auction with some high mail bids from an elderly gentleman who'd been a good customer of ours and obviously trusted us. My wife Melissa, who ran the business with me, stormed into my office the day after the sale, upset that I'd used his full bid on every lot, even when it was considerably higher than the second-highest bid. She threw his invoice on my desk and said, "I thought we weren't going to do this crap!" I glanced at the paperwork and without even thinking about it, said, "I don't like having to do this, honey, but you know our bank loan is due tomorrow." After some thought, she said, "Okay, I'll do this, but only if you agree to change the rules in our next auction to read, 'All lots are sold to the highest mail bidder at one advance over the second highest bid, unless we need the money.' " With that one sentence she stripped me of all my rationalizations and excuses. She held a mirror up to my

> conduct and I hated what I saw. I had no choice but to recalculate all our invoices for that sale to conform with our rules.
>
> That's when I decided to leave the business of running auctions.

The stamp auctioneer's situation can be interpreted as one in which he cannot commit not to negotiate with the winner upon conclusion of the auction. We will show that the pooling equilibrium we found in the first-price auction exists too in the second-price formats, for both continuous as well as discrete cost types. In the case of sealed bids, always. In the case of the English auction, only if the buyer has a fallback option in case their counteroffer gets rejected.

Let us start with continuous cost types.

Continuous cost types

In this section, we will show that the pooling equilibrium in the no-commitment setting can be sustained even in second-price auction formats. For now, the model is exactly as described in section 3.3; the only difference being the pricing rule: now, the winning seller is not paid their bid, but rather the second-lowest bid. Later on, we will look at the case of discrete cost types.

If we think back to the stamp auctioneer and his sealed-bid second-price auction that introduced this section, it becomes clear that all our results carry over from the first-price auction. For brevity, we will not repeat the full arguments here, but only provide the intuition.

Assume for the moment that the sellers are naïve in the sense that they continue to bid their true costs in the no-commitment setting. The buyer then has it easy: just offer the winning seller their cost (plus an infinitesimal amount). Of course, the same logic applies to any separating equilibrium, since the bid reveals the seller's cost. Working backwards, the buyer would just offer the winning seller their cost (plus an infinitesimal amount). This rules out any separating equilibrium. Consequently, we find that pooling on the highest possible cost is an equilibrium in the no-commitment setting of the sealed-bid second-price auction. Observe that this result carries through regardless of whether or not the buyer falls back on the second-lowest bid in case their counteroffer is rejected.

For the English auction—which we will consider next—this is not true. If there is no fallback option, we find ourselves in the setting of Shachat and Tan (2015), in which truthful bidding remains an equilibrium. However, in the face of a fallback option, we will see that the pooling equilibrium can exist; both for continuous, and for discrete cost types. Observe that in our setting—in which the buyer must procure the good—the fallback option in case the winning seller rejects the counteroffer is given by the stopping price of the auction.

Consider a continuous English clock auction with two sellers, where each seller's cost is independently drawn from the same continuous distribution. Assume for the moment that both sellers pool in the sense that they both immediately drop out of the auction. Whether or not this indeed constitutes an equilibrium depends on the exact distribution of sellers' costs. If costs are uniformly distributed, our results from the first-price auction carry over; if costs are e.g. normally distributed, they do not.

Regarding the former case, uniformly distributed costs on $[\underline{c}, \overline{c}]$: from proposition 2, when faced with uninformative bids, we can derive the buyer's optimal counteroffer as

$$o^* = \frac{\underline{c} + \overline{c}}{2},$$

a given seller's expected costs.

Evidently, we only have to check whether types $c < o^*$ have an incentive to deviate away from the posited pooling equilibrium; for other types to deviate guarantees zero profit. The question of existence of the pooling equilibrium boils down to whether or not types $c < o^*$ find it profitable to stay in the auction marginally longer, thus winning for sure. Note that the buyer cannot observe this deviation; they are only informed of the winning seller and the stopping price at the end of the auction.

Consider a seller S who may find deviation profitable. In the pooling equilibrium, S drops out immediately. With probability $\frac{1}{2}$, S is then selected as the winning seller, in which case their profit is given by $o^* - c$. With probability $\frac{1}{2}$, S's competitor is selected as the winner, who rejects the counteroffer with probability $1 - F(o^*) = \frac{1}{2}$. In this case, the buyer will buy the good from S at the dropout price $\bar{c}$. The expected profit of S when pooling is therefore

$$\frac{1}{2} \cdot (o^* - c) + \frac{1}{2} \cdot \frac{1}{2} \cdot (\bar{c} - c).$$

On the other hand, if S stays in the auction marginally longer—thus winning for sure—their profit is given by

$$o^* - c.$$

Comparing the two, we find that deviating away from the pooling equilibrium is profitable for S whenever $c < \underline{c}$, a contradiction. Taken together, we have shown that the pooling equilibrium can be sustained in a two-seller English auction if sellers' costs are uniformly distributed. However, some comments are in order. First, as mentioned above, the existence of the pooling equilibrium depends on the cost distribution. For example, doing the same calculations when seller's costs are normally distributed yields that the posited pooling equilibrium cannot be sustained. Second, unlike in the case of the first-price auction, the existence of the pooling equilibrium is not independent of the number of sellers. Indeed, even in the example above, for three sellers, the posited pooling equilibrium breaks down, since some cost types find deviating profitable.[3] Interestingly, if sellers' costs are discrete, the pooling equilibrium can be sustained independently of the number of sellers in the auction. Let us see why.

Discrete cost types with discrete grid

The general setting from above is unchanged: we still find ourselves in an English clock auction with two sellers. However, now, both the type space and the auction's price grid are discrete, with each seller's cost independently drawn from the set C, which also defines the allowable clock prices.

Binary cost types Let us start with $n \geq 2$ sellers and binary types, i.e. $C = \{c^l, c^h\}$, with $0 \leq c^l < c^h$. Denote by $p^h \in (0,1)$ the probability of a seller having costs of c^h. The cost draws are independent across sellers. As before, the buyer's valuation is $v = c^h$.

Observe that for a seller of type c^h, it is optimal to drop out of the auction immediately. Also, if the auction ends at a price of c^h, the buyer's optimal counteroffer is given by c^l.[4]

Assume that all sellers immediately drop out of the auction. Does a seller S of type c^l want to deviate? No. Deviating away from the posited pooling equilibrium yields zero profit. While deviating by staying in the auction marginally longer guarantees that S wins for sure, it also ensures zero profit, since the counteroffer they receive is equal to their costs. By contrast, immediately dropping out of the

[3] For completeness, for the case $n = 3$, all sellers with costs $c < \frac{2\underline{c}+\overline{c}}{3}$ find deviating profitable.

[4] The fallback option of c^h yields zero profit to the buyer. By making a counteroffer of c^l, the buyer makes a positive profit with a probability of $1 - p^h$.

auction yields a positive expected profit, namely

$$\underbrace{\frac{1}{n} \cdot (c^l - c^l)}_{\text{winner}} + \underbrace{(1 - \frac{1}{n}) \cdot p^h \cdot \frac{1}{n-1} \cdot (c^h - c^l)}_{\text{loser}} = \frac{p^h}{n}(c^h - c^l) > 0. \quad (3.3)$$

With a probability of $\frac{1}{n}$, seller S will be selected as the winner, in which case their profit is zero. With the counterprobability, another seller will be selected as the winner, and they will reject the buyer's counteroffer whenever they are of type c^h; this happens with a probability of p^h. Seller S will then be selected as the runner-up with a probability of $\frac{1}{n-1}$. In this case, the buyer will buy from S at the stopping price c^h, yielding a profit of $c^h - c^l$.

Interestingly, unlike the case of continuous costs above, for binary costs, the pooling equilibrium can be sustained independently of the number of sellers in the auction. Taking stock, let us preserve this finding in a proposition.

Proposition 8. *For binary cost types, there exists a pooling equilibrium in the English auction, irrespective of the number of sellers in the auction, and also irrespective of the probability distribution over the cost types.*

Proof. The proof follows directly from eq. (3.3). □

What happens if we extend the model to allow for more than two cost types?

Three or more cost types Consider the same setting as above with $n \geq 2$ sellers, except that the set of possible cost types is now $C = \{c^l, c^m, c^h\}$, with $0 \leq c^l < c^m < c^h$. For the subsequent analysis, the probability of a seller being of type c^m is p^m; the probability of a seller being of type c^h is p^h; and the remaining probability mass is on c^l. As before, all probabilities are $\in (0, 1)$, and the probabilities sum to one. The cost draws remain independent across sellers.

Whether or not the pooling equilibrium exists depends on what the buyer's optimal counteroffer is when faced with uninformative bids. Note that, because the buyer can always fall back on the dropout price c^h, the optimal counteroffer is always given either by c^l or c^m.

Case 1: $o^* = c^l$ Assume the auction ends immediately. If the buyer's optimal counteroffer in such a case is given by c^l, we know from the previous section that the pooling equilibrium can be sustained. Types c^h and c^m want to drop out immediately, and types c^l cannot do better by deviating and staying in the auction, since this guarantees zero profit. An example for a case where $o^* = c^l$ is if $\{c^l, c^m, c^h\} = \{0, 2, 3\}$, each with equal probability.

Generalising this result, we can preserve it in a proposition.

Proposition 9. *Consider an English auction in which ties are broken randomly with $n \geq 2$ sellers in the no-commitment setting. Each seller's cost is independently drawn from the discrete set $\{c^l, \ldots, c^h\}$. If the buyer's optimal counteroffer when faced with uninformative bids is given by c^l, it is an equilibrium for all cost types to drop out of the auction immediately. The buyer responds by making the winning seller a counteroffer of c^l. If this counteroffer is rejected, the buyer buys the good from the second-lowest bidder at the dropout price c^h.*

Proof. The proof follows directly from the previous examples. None of the sellers has an incentive to deviate, since this guarantees zero profit. Neither does the buyer want to deviate: they make the optimal counteroffer to the winning seller. The claim in the proposition follows. □

Case 2: $o^* = c^m$ Consider again the case of three cost types, $C = \{c^l, c^m, c^h\}$, with $n \geq 2$ sellers. If the buyer's optimal take-it-or-leave-it offer is given by c^m, whether or not the posited pooling equilibrium exists depends on the exact values of sellers' possible costs. Once again, consider a seller S of type c^l, who may find deviation profitable.

Deviating by staying in the auction marginally longer yields a profit of

$$c^m - c^l,$$

whereas the pooling equilibrium yields an expected profit of

$$\underbrace{\frac{1}{n} \cdot (c^m - c^l)}_{\text{winner}} + \underbrace{(1 - \frac{1}{n}) \cdot p^h \cdot \frac{1}{n-1} \cdot (c^h - c^l)}_{\text{loser}} = \frac{1}{n} \cdot (c^m - c^l) + \frac{p^h}{n}(c^h - c^l).$$

Comparing the two, we find that S wants to deviate whenever

$$c^m > c^l + \frac{p^h}{n-1}(c^h - c^l) \tag{3.4}$$

To sum up, for certain parameter constellations, the pooling equilibrium can be sustained; but for other constellations, type c^l finds deviating profitable. Note that in this case, not only does the pooling equilibrium break down, it may also no longer be optimal for the buyer to make a counteroffer of c^m. Let us look at an example. For $n = 3$, if each of the three possible costs is equally likely, pooling can be sustained for $\{c^l, c^m, c^h\} = \{0, \frac{1}{4}, 3\}$. However, pooling breaks down if $\{c^l, c^m, c^h\} = \{0, \frac{5}{4}, 3\}$. Generalising this result, we obtain a proposition.

Proposition 10. *Consider an auction with n sellers and $k \geq 3$ discrete cost types. That is, $C = \{c^1, c^2, \ldots, c^k\}$, where $0 \leq c^1 < c^2 < \ldots < c^k$. If the buyer's optimal take-it-or-leave-it offer is given by $c' \in C \setminus \{c^1, c^k\}$, there always exist cost parameters for which the pooling equilibrium can be sustained.*

Proof. This follows directly from eq. (3.4) above. Replace c^m with c', then c^l with c^1, and lastly p^h and c^h with p^k and c^k, respectively. If we flip the inequality, the claim immediately follows. □

3.5 Conclusion

In this paper, we build on the results of the previous chapter in which we showed that in first-price procurement auctions with continuous cost types, lacking commitment is detrimental to the buyer: it leads to higher buyer expenses vis-à-vis an auction with commitment. In this paper, we have shown that this result is more general. For the sealed-bid first- and second-price formats, we show the existence of the pooling equilibrium independently of both the number of bidders in the auction and the presence of a fallback option.

Historically, the English auction is considered immune to the commitment problems faced by a buyer conducting a sealed-bid second-price auction. In this paper, we show that this result hinges crucially on the buyer lacking a fallback option can in case of a rejected counteroffer. By contrast, in the case of a fallback option given by the closing price of the auction, we show that a pooling equilibrium can exist even in English auctions. For discrete seller costs, this is always true. For continuous seller costs, the existence of pooling equilibrium depends on the number of bidders and how costs are distributed.

Chapter 4

Procurement Auctions with Demand Uncertainty

Consider a buyer with uncertain demand. How should this buyer account for this demand uncertainty when running a procurement auction? We consider three options: (1) run an auction in which the per-unit price is valid uniformly, regardless of the actual demand; (2) contractually define what happens in case additional units are required by setting a fixed price premium in the case of additional demand; and (3) negotiate the price for additional demand with the winning seller only if and when it occurs. We show theoretically that—for buyers and sellers alike—all three options yield the same expected profits. We then take these settings to the laboratory. Experimentally, we find that the equivalence between the three settings does not hold, neither for buyers, nor for sellers. Sellers—as a result of low-markup bidding—do best when in the setting with a price premium. Buyers, on the other hand, do best in the simplest of the three settings, the uniform price setting. However, a closer look reveals that this result is mainly driven by the large share of failed negotiations—approximately one third—in the negotiation setting. Higher maximum prices on the buyer's part would likely have increased their profit.

4.1 Introduction

In many procurement contexts, buyers enter into procurement contracts with their suppliers without knowing what their total demand will be. This demand uncertainty presents a critical challenge for buyers and suppliers alike, affecting both contractual arrangements and operational performance. This is particularly salient in industries characterized by high specificity and limited supplier substitutability, such as the automotive sector. Unlike commodities markets, where demand fluctuations can be absorbed by shifting to alternative sources of supply, in industries with asset-specific investments and complex technical standards, buyers are often locked into relationships with designated suppliers (Williamson, 1985; Crocker and Reynolds, 1993).

The automotive industry exemplifies this issue. Automakers rely heavily on customised parts manufactured to precise specifications, which are often co-developed with suppliers during lengthy design and validation processes (Helper and Sako, 2010). When actual demand exceeds forecasted levels, buyers cannot simply pivot to alternative original equipment manufacturers (OEMs) to source additional units. Instead, they must negotiate with their incumbent suppliers, who now possess increased bargaining power due to the buyer's lack of alternatives and the time-sensitive nature of automotive production schedules (Taylor and Plambeck, 2007). This asymmetry can lead to ex-post renegotiation, price markups, and strained buyer-supplier relationships, see e.g. Mogge (2023) for a discussion on this. Another prominent and recent example for demand uncertainty comes from the public sector, in the form of vaccine procurement. During the COVID-19 pandemic, governments entered into contracts with vaccine manufacturers *without knowing what their actual demand will be.* The EU dealt with this uncertainty by procuring a fixed quantity of 2.4 billion doses, and reserving the right to procure an additional 2.2 billion doses if necessary (European Union, 2022).

Naturally, the degree of demand uncertainty varies by product. For new products, demand uncertainty is typically higher, as it is unclear how the market will respond to them. For established products, forecasting demand is generally easier, as historical data provides a good basis for making predictions. Of course, both buyer and sellers would benefit from efforts to reduce demand uncertainty.

However, the buyer's ability to reduce demand uncertainty may often be limited, especially in terms of pricing and promotion, which are common strategies used to mitigate risk. In practice, there are multiple ways with which demand uncertainty can be dealt with, which differ in the way the uncertainty is distributed between buyer and seller. On the one extreme, we have framework contracts, in which the buyer has the right—but not the obligation—to have the seller deliver goods at a pre-determined price, irrespective of the final demand. The other extreme—negotiations—favours the seller, since the buyer must enter into negotiations with the seller for any demand above the quantity agreed upon. This is the result of what Williamson (1985) calls the "fundamental transformation", which transforms an ex-ante competitive, open environment into one of bilateral dependency as soon as a contract is signed. Hart and Moore (2008) demonstrate how this dependency can lead to opportunism. The question we seek to answer in this paper is as follows: how should the buyer account for demand uncertainty? In this paper, we will consider auctions as the procurement mechanism. Procurement auctions are ubiquitous, are used by the private and sector alike, and have an annual throughput in the billions of US dollars—with the potential for trillions of US dollars (Jap, 2002). We draw inspiration from the two extremes discussed above by considering three settings with which to deal with demand uncertainty in the context of auctions. The first is that of a framework contract, i.e. buyers provide a demand forecast to sellers and simply request a unit price irrespective of the final demand. The second option is for buyers to buy a fixed quantity and then to enter into further negotiations for additional quantity only if and when it arises. The third option we consider is that of a buyer entering into a contract that defines contingencies. That is, buyers request bids for the fixed quantity and contractually define what happens in case they require additional demand. In our specific case, buyers offer the seller a fixed price premium on the auction price for additional quantity.To answer our research question, we translate the three settings from above—uniform price, price premium, and negotiation—into mathematical models. We show theoretically that when it comes to both buyer expenses and seller profits, all three settings yield the same expected buyer expenses.

That being said, actual bidding behavior in auctions frequently diverges from theoretical predictions. Empirical studies have consistently shown that bidders systematically deviate from equilibrium strategies, see e.g the discussion on overbidding in first-price auctions started by Cox et al. (1982) and continued by Kagel and Levin (1993). Behavioral factors, including risk aversion, regret aversion, and bounded rationality, contribute to these deviations (Thaler, 1988; Engelbrecht-Wiggans and Katok, 2007). Moreover, experiments have demonstrated that bidders may also be influenced by social preferences and psychological biases (Tversky and Kahneman, 1974; Dufwenberg and Gneezy, 2000). Therefore, we take these three settings to the laboratory to test this whether the theoretical equivalence also holds experimentally. It does not: we find that both buyer expenses as well as seller profits differ across our three experimental settings.

For the buyer, their total expenses are lowest in the uniform price and negotiation settings, but the variance in buyer expenses is also higher in these two settings compared to the price premium setting. However, there seems to be further room for improvement in the negotiation setting: buyers reject profitable offers from contract suppliers in approximately a third of the time. A counterfactual analysis suggests that had buyers been willing to accept higher prices, they could have reduced their expenses by 8.5%. On the other hand, sellers are best-off in the price premium setting. This is driven by the fact that sellers in the negotiation setting request a low premium for the additional demand: sellers in the negotiation treatment request an average premium of approximately one quarter of the fixed premium offered in the price premium treatment.

The rest of the paper is structured as follows: section 4.2 connects this research with the existing literature; section 4.3 introduces the model, and derives the theoretical predictions. In section 4.4, we discuss the experimental results, while section 4.5 discusses the robustness of our theoretical model. Finally, section 4.6 concludes.

4.2 Related Literature

Our paper is related to the literature that studies procurement mechanisms and supply chains under various forms of uncertainty. Starting with the seminal work of Arrow et al. (1951), a large stream of literature on procurement decisions in newsvendor-like settings has emerged, see e.g. Khouja (1999) for an overview. An important feature of the newsvendor setting is that the buyer faces *exogenous* prices. By contrast, in this paper, we will specifically be considering auctions as the procurement mechanism, that is, prices are *endogenously* determined. As such, this paper finds itself at the intersection of auction theory with operations management. Since the seminal paper by Vickrey (1961), the field of auction theory has grown strongly. An overview of the field of auction theory can be found in McAfee and McMillan (1987) and Klemperer (1999), while a discussion of auctions—or other price-discovery mechanisms—in an operations context can be found in Elmaghraby (2000).

In general, the setting we consider is one in which the buyer's demand is uncertain, and sellers have private production costs. Therefore, our paper is related to the work of Chen (2007), who demonstrate that in such a setting, an auction is an optimal procurement mechanism. Their auction is characterised by the buyer defining a payment for each possible order quantity, and then inviting sellers to bid. Naturally, Chen's proposed procurement mechanism becomes less tractable as the set of possible quantities grows. Addressing this issue, Duenyas et al. (2013) demonstrate that a variant of the open-descending auction for a fixed quantity—a much simpler model, and one more familiar to sellers—is also an optimal procurement mechanism. Various extensions to the general model have been made to cover e.g. a buyer facing price-sensitive demand (Zhang, 2010), or sellers with capacity constraints (Li and Scheller-Wolf, 2011). While we also consider auctions with uncertain buyer demand, the papers listed above only consider a one-shot quantity decision. By contrast, we extend the literature by studying a setting in which re-stocking is possible. That is, we find ourselves in a situation in which the buyer is always able to meet demand by procuring additional units should the need arise.

Gur et al. (2017) consider such a setting. They study a buyer who has uncertain demand that must be met in the context of auctioning off framework contracts. They consider a case in which sellers' costs—which are composed of both a private as well as common component—are uncertain over the duration of the framework contract. They show that this cost uncertainty means that procurement costs are higher when auctioning off a framework contract compared to running auctions as the need for additional units arises. Similarly, Schummer and Vohra (2003) study the mechanism design problem of a buyer facing an uncertain future demand and capacity-constrained sellers in the context of electricity markets. They demonstrate how the buyer can efficiently ensure that demand will be always be met by procuring "options": that is, the right, but not the obligation, to have a specified quantity of electricity delivered by a given seller. Our paper is different from the two above in that we assume independent seller costs and single sourcing. Moreover, unlike in the case Schummer and Vohra, we allow the buyer to negotiate with the seller upon realisation of additional demand, thereby explicitly allowing situations in which the negotiation breaks down and demand is left unfulfilled.

In our paper, we will consider a setting in which, upon completion of the auction, a price negotiation takes place for any potential additional demand. As such, our research is also related to the literature on auction and negotiations. A key result in this area of enquiry was delivered by Bulow and Klemperer (1996), who show that, for the buyer, having an additional seller in an auction is preferable to any procurement mechanism one could run with one fewer seller. However, note that in our setting, we are not pitting auctions against negotiations. Rather, we consider are considering both sequentially. More recently, Tunca and Zenios (2006) study a setting in which a buyer needs to procure goods of varying quality. They derive conditions under which both procurement channels—auctions and negotiations—co-exist.

Our negotiation treatment, in which a price negotiation for any additional demand takes place, can be interpreted as a setting in which sellers compete in an auction for the right to to participate in ultimatum bargaining in the role of proposer with the buyer. As such, this setting is similar to that of Shachat and Swarthout (2013), with the difference that in their case, both the proposer and the responder were determined through auction. In our case, the buyer does not compete with others for the role of responder in the ultimatum game. Experimentally, Shachat and Swarthout find that, compared to a setting with exogenous participation costs, auctioning the right to participate in ultimatum bargaining leads to higher proposer (in our case, read: seller) profits, with this effect being driven by the fact that buyers accept the sellers' more aggressive proposals in approximately 90% of rounds. While selection of the proposer always occurs through auction in our setting—which means we cannot benchmark against a case with exogenous participation costs—it nevertheless worth noting that we find the opposite: of the three treatments we consider, seller profits are *lowest* in the negotiation treatment in which sellers compete in an auction for the right to participate in ultimatum bargaining with the buyer, with this result being driven by high rejection rates on the part of the buyer.

Katok and Tan (2025) study negotiations in the context of another type of uncertainty: supply disruptions. In the setting they consider, a buyer with a fixed demand enters into a purchase contract with a single seller. During the duration of the contract, a costly supply disruption can occur. Varying relationship length, bargaining power, and the ability of buyer and seller to communicate with each other, Katok and Tan demonstrate experimentally that the ability to renegotiate with the seller in the event of a supply disruption increases buyer profit vis-à-vis a fixed contract, but only when considering short-term contracts. Observe that in our setting, we are considering a different type of uncertainty—demand uncertainty. Nevertheless, the fact that we are studying one-shot interactions between buyer and seller, i.e. short-term contracts, gives us the opportunity to test whether this result also holds in the case of demand uncertainty. While buyer profits are indistinguishable between the fixed contract and negotiation treatments in our setting, we show that this is the result of buyers negotiating too aggressively. A buyer who is willing to accept higher prices could indeed, ceteris paribus, obtain higher profits when negotiation is possible.

4.3 Theory

4.3.1 Setting

We consider a setting where a buyer seeks to procure an essential input for the production of a new product and faces n potential suppliers. The buyer's valuation for each unit of the input is denoted by v. The total quantity of the input required q is uncertain and follows a distribution H with support $[\underline{q}, \overline{q}]$ and $0 \leq \underline{q} < \overline{q}$. This demand distribution H is common knowledge among all market participants.

Each supplier $i \in \{1, \ldots, n\}$ has constant marginal cost c_i, which is independently drawn from a distribution F with support $[\underline{c}, \overline{c}]$ and $0 \leq \underline{c} < \overline{c} \leq v$. The cost distribution F is common knowledge, but each supplier's realized cost c_i is private information. Let $c^{(1:k)}$ and $c^{(2:k)}$ denote the lowest and second-lowest costs, respectively, among k independent cost draws from F. Define $G^{(k)}$ as the cumulative distribution function of the minimum of k independent cost draws from F, i.e., $G^{(k)}(x) = \Pr\{c^{(1:k)} \leq x\} = 1 - [1 - F(x)]^k$.

To award the contract, the buyer conducts an auction in which one of the n suppliers is selected as the contract supplier. The contract supplier will then be responsible for delivering the required quantity q of the input once demand is realized.

We consider three different types of procurement contracts that the buyer can award via two types of auctions: a first-price auction and a second-price auction.

(a) **Uniform Price Contract:** Under a uniform price contract, the buyer pays a constant per-unit price p^U for each unit procured. Suppliers submit bids specifying a per unit-price, and the supplier offering the lowest bid wins the contract. The total payment to the contract supplier is given by

$$q \cdot p^U.$$

If the buyer conducts a first-price auction, p^U corresponds to the lowest bid. In a second-price auction, p^U is set equal to the second-lowest bid.

(b) **Price Premium Contract:** Under a price premium contract, the buyer commits to sourcing at least a minimum quantity $\underline{q}$ from the contract supplier at a per-unit price p^P, determined through the auction. If the realized demand q exceeds $\underline{q}$, the buyer procures the additional quantity $q - \underline{q}$ at a higher per-unit price of $p^P + \Delta$, where Δ is a fixed price premium that the buyer announces before the auction. The supplier offering the lowest base price p^P wins the contract. The total payment to the contract supplier is

$$\underline{q} \cdot p^P + (q - \underline{q}) \cdot (p^P + \Delta).$$

If the buyer conducts a first-price auction, p^P corresponds to the lowest bid. In a second-price auction, p^P is set equal to the second-lowest bid.

(c) **Fixed-Quantity Contract:** Under the fixed-quantity contract, the buyer commits to sourcing a baseline quantity $\underline{q}$ from the contract supplier at a per-unit price p^F. Suppliers submit bids specifying p^F, and the supplier offering the lowest bid wins the contract. If realized demand q exceeds $\underline{q}$, the buyer and the contract supplier may engage in ex-post bilateral negotiation over the additional quantity $q - \underline{q}$. If an agreement is reached, the buyer pays a per-unit price p^A for the extra units. The total payment to the supplier in the case of agreement is

$$\underline{q} \cdot p^F + (q - \underline{q}) \cdot p^A.$$

Without an agreement, the total payment to the supplier is

$$\underline{q} \cdot p^F.$$

If the buyer conducts a first-price auction, p^F corresponds to the lowest bid. In a second-price auction, p^F is set equal to the second-lowest bid.

4.3.2 Analysis

Next, we characterize the suppliers' equilibrium bidding strategies in each auction format under each of the three contract types. Based on these strategies, we derive the buyer's expected total procurement costs and compare the cost implications across contract variants.

Uniform Price Contract

Under a uniform price contract, the expected profit of supplier i is

$$\Pi_i^U = \begin{cases} (p^U - c_i) \cdot E[q] & \text{if } b_i < \min_{j \neq i} b_j \\ 0 & \text{if } b_i > \min_{j \neq i} b_j. \end{cases}$$

Proposition 11. *The following bidding strategies constitute symmetric equilibrium strategies in a first-price and second-price auction for uniform price contracts, respectively:*

(a) $\beta^{U,1}(c) = E[c^{(1:n-1)} | c^{(1:n-1)} > c]$

(b) $\beta^{U,2}(c) = c$

Proof. The equilibrium bidding function for the first-price auction is derived analogously to the one of a forward auction. In the second-price auction, it is a weakly dominant strategy for sellers to bid their costs. For a detailed exposition of both these equilibria, see e.g. Krishna (2010). □

Proposition 11 directly implies that, regardless of the auction format, the buyer's total expected expenditure under a uniform price contract is $K^U = E[q] \cdot E[c^{(2:n)}]$.

Price Premium Contract

Under a price premium contract, the expected profit of supplier i is

$$\Pi_i^P = \begin{cases} (P^P - c_i) \cdot \underline{q} + \Delta \cdot [E[q] - \underline{q}] & \text{if } b_i < \min_{j \neq i} b_j \\ 0 & \text{if } b_i > \min_{j \neq i} b_j. \end{cases}$$

In a second-price auction for a price premium contract, it is optimal for a supplier to submit a bid such that, if the auction price p^P equals this bid, the supplier is indifferent between winning and not winning the contract.

Proposition 12. *The following bidding strategies constitute symmetric equilibrium strategies in a first-price and second-price auction for price premium contracts, respectively:*

(a) $\beta^{P,1}(c) = E[c^{(1:n-1)}|c^{(1:n-1)} > c] - \Delta \cdot \frac{E[q]-\underline{q}}{E[q]}$

(b) $\beta^{P,2}(c) = c - \Delta \cdot \frac{E[q]-\underline{q}}{E[q]}$

Proof. Utilizing the same method as for the uniform price contract above, in the first-price auction, the first-order condition is given by

$$\frac{\partial}{\partial c}[\beta^{P,1}(c) \cdot G(c)] \cdot E[q] = g(c) \cdot [c - \Delta \cdot (E[q] - \underline{q})]$$

Solving the derivative and rearranging yields

$$\begin{aligned} \beta^{P,1}(c) &= E[c^{(1:n-1)}|c^{(1:n-1)} > c] - \Delta \cdot \tfrac{E[q]-\underline{q}}{E[q]} \\ &= \beta^{U,1} - \Delta \cdot \tfrac{E[q]-\underline{q}}{E[q]}, \end{aligned}$$

the bidding function of the uniform price contract, minus a constant.

For the second-price auction, it is a weakly dominant strategy for sellers to bid the value $\beta^{P,2}$ that makes them indifferent between winning and losing. A seller's profit when bidding $\beta^{P,2}$ is

$$(\beta^{P,2} - c) \cdot \underline{q} + (\beta^{P,2} + \Delta - c) \cdot (E[q] - \underline{q})$$

Solving for $\beta^{P,2}$ yields $c - \Delta \cdot \frac{E[q]-\underline{q}}{E[q]}$.

□

Proposition 12 directly implies that, regardless of the auction format, the buyer's total expected expenditure under a price premium contract is $K^P = K^U$. The analysis of the bidding behavior shows that suppliers compete away additional profits connected to the price premium by bidding more aggressively and lowering the auction price p^P compared to p^U.

Fixed Quantity Contract

Under a fixed quantity contract, the expected profit of supplier i is

$$\Pi_i^F = \begin{cases} (p^F - c_i) \cdot \underline{q} + (p^A - c_i) \cdot (E[q] - \underline{q}) & \text{if } b_i < \min_{j \neq i} b_j \text{ and buyer accepts} \\ (p^F - c_i) \cdot \underline{q} & \text{if } b_i < \min_{j \neq i} b_j \text{ and buyer rejects} \\ 0 & \text{if } b_i > \min_{j \neq i} b_j. \end{cases}$$

If the realized quantity exceeds $\underline{q}$, the contract supplier makes a take-it-or-leave-it offer to the buyer for a per-unit price p^A. The buyer is better off accepting than rejecting whenever $p^A \leq v$, which renders it optimal for the contract supplier to offer $p^A = v$. Hence, the expected profit of the winning supplier can be written as

$$(p^F - c_i) \cdot \underline{q} + (v - c_i) \cdot [E[q] - \underline{q}].$$

Proposition 13. *The following bidding strategies constitute symmetric equilibrium strategies in a first-price and second-price auction for fixed quantity contracts, respectively:*

(a) $\beta^{F,1}(c) = E[c^{(1:n-1)} | c^{(1:n-1)} > c] - (v - E[c^{(1:n-1)} | c^{(1:n-1)} > c]) \cdot \frac{E[q] - \underline{q}}{\underline{q}}$

(b) $\beta^{F,2}(c) = c - (v - c) \cdot \frac{E[q] - \underline{q}}{\underline{q}}$

Proof. Utilizing the same method above, in the first-price auction, the first-order condition is given by

$$\frac{\partial}{\partial c}[\beta^{F,1}(c) \cdot G(c)] \cdot \underline{q} = g(c) \cdot [c - (v - c) \cdot (E[q] - \underline{q})]$$

Solving the derivative and rearranging yields

$$\beta^{F,1}(c) = E[c^{(1:n-1)}|c^{(1:n-1)} > c] - (v - E[c^{(1:n-1)}|c^{(1:n-1)} > c]) \cdot \frac{E[q]-\underline{q}}{\underline{q}}.$$

For the second-price auction, it is a weakly dominant strategy for sellers to bid the value $\beta^{F,2}$ that makes them indifferent between winning and losing. A seller's profit when bidding $\beta^{F,2}$ is

$$(\beta^{F,2} - c_i) \cdot \underline{q} + (v - c_i) \cdot [E[q] - \underline{q}]$$

Solving for $\beta^{F,2}$ yields $c - (v - c) \cdot \frac{E[q]-\underline{q}}{\underline{q}}$.

□

Proposition 13 directly implies that, regardless of the auction format, the buyer's total expected expenditure under a fixed quantity contract is $K^F = K^P = K^U$.

Taken together, we have shown that all three contracts yield the same total expenditure to the buyer. In the following section, we will test this prediction experimentally.

4.4 Experiment

In this section, we will take the three settings described above—uniform price, price premium, and negotiation—into the laboratory.

4.4.1 Design

In the experiment, one buyer is matched with $n = 2$ sellers. The buyer has a fixed demand of 100 units, plus an additional 100 units with probability 50%. The buyer's value for each unit of the good is set to $v = 1.30$.

Sellers' private per-unit costs are independently and uniformly drawn from the set $C = \{0.01, 0.02, \ldots, 0.99, 1.00\}$.

We have three treatments: *uniform price*, *price premium*, and *negotiation*.

Uniform price

In the uniform price treatment, sellers submit a single per-unit bid for both the fixed and additional demand. The lowest bidder wins and the per-unit price they are paid is given by their bid, irrespective of the number of units the buyer buys.

Price premium

In the price premium treatment, sellers submit a bid only for the fixed demand. The lowest bidder wins and is paid their bid for the fixed demand. Should the additional demand realise, the winner of the auction will be paid a per-unit premium of 0.30 in addition to their bid. This information is communicated to sellers before the auction.

Negotiation

In the negotiation treatment, sellers submit only a bid for the fixed demand. The lowest bidder wins and is paid their bid for the fixed demand. The price negotiation for the additional demand is implemented as follows. In addition to their bid for the fixed demand, all sellers submit a price request for additional demand. The buyer is informed of the auction price, and is asked to submit a maximum per-unit price for the additional demand. Only hereafter does the additional demand realise. Trade for the additional demand takes place only if the buyer's maximum price is equal or greater to the winning seller's requested price.

In all of the treatments, per-unit bids for the fixed quantity are in $\in C$. For the additional quantity in the negotiation treatment, bids are in $\in [0.01, 0.02, \ldots, 1.29, v]$. Ties are broken randomly.

Note that in the uniform price and price premium treatments, the buyer has a passive role, since they cannot influence their payoff. Given the active role of the buyer in the negotiation treatment, the human buyer was kept in order to rule out that any observed differences are driven by sellers facing a computerised buyer.

In all treatments, at the end of each round, all subjects are informed of the bids that were placed, as well as whether or not the additional demand realised. In the negotiation treatment, subjects are also informed of the winning seller's requested price for additional demand, as well as the buyer's maximum willingness-to-pay. Lastly, all subjects are informed what their resulting payoff is. Note that in each round, sellers were given a fixed payment of 30 monetary units in addition to their payoff. This was done to ensure that a seller's expected payment at the end of the experiment meets the minimum payment set out in the laboratory guidelines.

4.4.2 Organisation

The experiment was conducted on-site in a computer laboratory of a large western European university using the software oTree (Chen et al., 2016). Subjects were recruited from the university's subject pool, with cash being the only incentive offered. The average payment was 15.58 EUR, which corresponds to about 16.89 USD at the time of the experiment. Subjects were mostly undergraduates and came from a variety of majors.

Upon correctly answering the control questions, all subjects played 10 rounds of a first-price auction in the role of a seller. This was done in order to allow subjects to familiarise themselves with the auction setting. Moreover, it allows us to check if the randomisation of subjects across treatments was successful. In each of the ten rounds, every seller was competing against a computerised seller whose per-unit bids are uniformly distributed between 0.50 and 1.00. The computerised buyer had a constant demand of 100 units, bought the good from the lowest bidder, and paid them their bid. In the results section below, these ten initial rounds will be called the *pre-treatment*.

Thereafter, the respective treatment was played. Subjects were randomly assigned a role—either buyer or seller—and this role remained constant throughout the 40 rounds played. Subjects were divided into matching groups of six, with each group consisting of four sellers and two buyers. In each round, a buyer was randomly paired with two sellers. Each laboratory session consisted of between 24 and 54 subjects. For each treatment, 60 subjects were recruited. In total, this gives us 10 independent observations per treatment. Figure 2.1 below provides a summary of the experiments we ran. The experiment ended with a survey that allows us to determine subjects' personality traits. The survey is described in section 4.4.4.

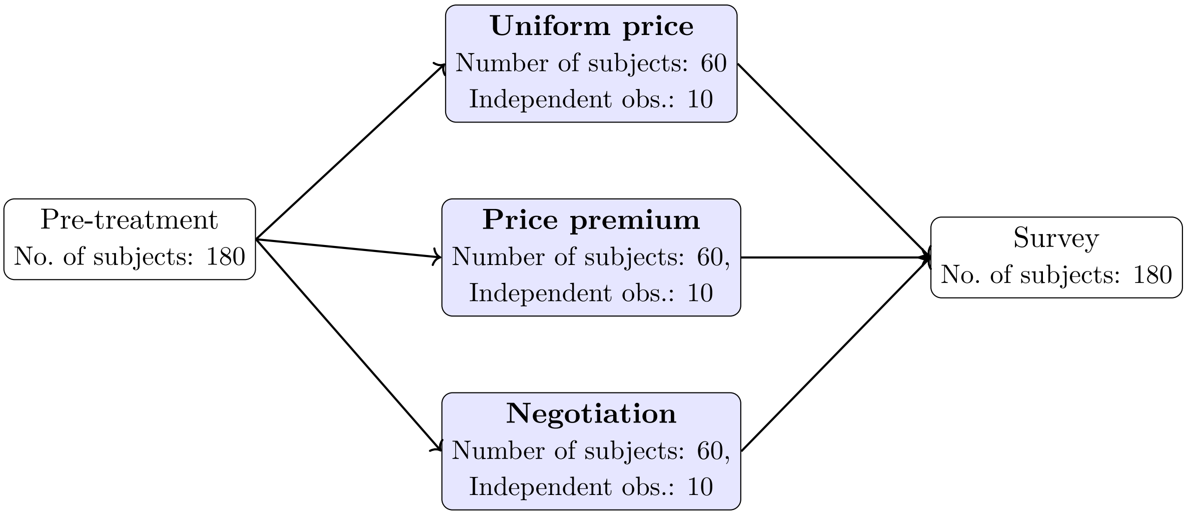

Figure 4.1: Overview of all treatments ran.

4.4.3 Hypotheses

In the uniform price setting, applying proposition 11, risk-neutral sellers are predicted to submit a per-unit bid according to the bidding function

$$\beta(c) = \frac{1+c}{2} \tag{4.1}$$

Applying proposition 12, sellers in the price premium treatment bid according to

$$\beta(c) = \frac{7+13c}{20} \tag{4.2}$$

In the negotiation treatment, applying proposition 13, sellers bid according to

$$\beta(c) = \frac{1+39c}{40} \tag{4.3}$$

We can now translate these bidding functions into concrete hypotheses. As presented in table 4.1, for the fixed quantity, the expected bids are 0.75, 0.68, and

0.51 for the fixed price, price premium, and negotiation treatment, respectively. This leads to respective expected prices of 0.67, 0.58, and 0.36.

This gives us the first two hypotheses. Denote by P_F the per-unit price for the fixed demand and by P_A the per-unit price for the additional demand.

Hypothesis 10. *Price ranking for the fixed demand.*

The price P_F is highest in the uniform price treatment and lowest in the negotiation treatment.

Let us turn to the expected price for the additional demand. In the negotiation treatment, the theoretical prediction is that buyer sets a maximum willingness-to-pay of v. Note that this is optimal for the buyer: if the buyer's maximum willingness-to-pay is at least as high as the winning seller's price request for additional demand, the price for the additional demand is given by the minimum of the two. Therefore, it is a dominant strategy for the buyer to set v as their maximum willingness-to-pay. This is because any maximum willingness-to-pay less than v puts the buyer at risk of of rejecting a profitable trade. Observe that the winning seller knows the buyer's per-unit valuation for the good. Therefore, the winning seller will demand a price of v for the additional demand. This gives us the next two hypothesis.

Hypothesis 11. *Price ranking for the additional demand.*

The price P_A is lowest in the uniform price treatment and highest in the negotiation treatment.

The total profit to the buyer is given by the product of their per-unit valuation and their total demand, minus the cost of procuring these units. Leveraging propositions 11 to 13, buyer profit is expected to be *identical* across treatments. Given the parameters of the experiment, average buyer profit across all rounds is expected to be 95.00.

Hypothesis 12. *Buyer profit is identical across treatments.*

While expected buyer profit is identical across the treatments, the variances are not.

Hypothesis 13. *The variance in buyer profit is lowest in the uniform price treatment and highest in the negotiation treatment.*

The hypothesis above was formulated from the buyer's perspective, however, the ranking carries over to sellers too. Note too that in the negotiation treatment, the theoretical model predicts that trade will always takes place for the additional demand

Next, let us turn to efficiency. An outcome is efficient if the seller with the lowest costs supplies the buyer. In all treatments, the bidding functions are increasing in c. Therefore, the seller with the lowest cost always wins the auction for the fixed demand. As such, regardless of the treatment, outcomes are always expected to be efficient.

4.4.4 Results

In this section, we will present the experimental results. We consider each matching group as an independent observation. Unless stated otherwise, all comparisons in this section are made using Wilcoxon rank-sum tests on independent observations.

Comparing the bids of the initial ten rounds of the pre-treatment across the three main treatments allows us to check for differences in the subject pools. A Kruskal-Wallis rank sum test yields no significant difference between the three groups ($p = 0.19$). Similarly, when making pair-wise comparisons, two-sided Wilcoxon rank-sum tests also yield no significant differences ($p = 0.27$ for the comparison between uniform price and price premium, $p = 0.11$ for the comparison between uniform price and negotiation, and $p = 0.34$ for the comparison between price premium and negotiation). Taken together, this gives us an indication that our randomisation of subjects across treatments was successful.

Table 4.1 reports the results of the experiment. For each treatment, we report the averages, as well as the theoretical predictions based on the realised cost draws of the experiment, including non-parametric tests.

Across all treatments, we find that bids for the fixed demand, prices, and efficiency are all significantly different to the theoretical predictions (all $p < 0.05$). Except in the case of the price premium treatment, buyer profit is significantly different too.

Result 12. *Prices for the fixed demand are highest in the uniform price treatment (*$p < 0.10$ *and* $p < 0.05$*, compared to the price premium and negotiation treatments, respectively) and lowest in the negotiation treatment (both* $p < 0.05$*).*

We find support for hypothesis 10, which predicted this ranking of bids and prices across treatments. This finding is in line with the theoretical prediction that, compared with the uniform price treatment, sellers' expected profits are *higher* for the additional demand in the other two treatments. This makes winning the auction for the fixed demand more attractive, which translates into lower bids and prices.

Let us turn next to the additional demand.

Result 13. *Prices for the additional demand are lowest in the uniform price treatment (both* $p < 0.01$*). However, prices for the additional demand in the negotiation treatment are significantly lower than in the price premium treatment (*$p < 0.10$*).*

We therefore find partial support for hypothesis 11, which predicted the highest prices for the additional demand in the negotiation setting.

Next, we will consider the central metric of buyer profit.

Result 14. *While theory predicts identical buyer profits across treatments, buyer profit is lowest in the price premium treatment (both* $p < 0.05$*). There is no*

	Observed			*Theoretical*		
	Uniform price	**Price premium**	**Negotiation**	**Uniform price**	**Price premium**	**Negotiation**
$\textbf{Price}_{\textbf{fixed}}$	0.57 (0.03)	0.53^{*} (0.01)	0.50^{**} (0.02)	$0.67^{\dagger\dagger}$	$0.58^{\dagger\dagger\dagger}$	$0.36^{\dagger\dagger\dagger}$
$\textbf{Price}_{\textbf{additional}}$	0.57 (0.03)	0.83^{***} (0.01)	0.76^{***} (0.04)	$0.67^{\dagger\dagger}$	$0.88^{\dagger\dagger\dagger}$	$1.30^{\dagger\dagger\dagger}$
$\textbf{Profit}_{\textbf{buyer}}$	103.57 (4.06)	95.20^{***} (1.35)	102.57 (2.48)	$93.66^{\dagger}$	93.66	$93.66^{\dagger\dagger}$
$\textbf{Bid}_{\textbf{fixed}}$	0.69 (0.02)	0.64^{**} (0.01)	0.62^{***} (0.01)	$0.75^{\dagger}$	$0.68^{\dagger\dagger\dagger}$	$0.51^{\dagger\dagger}$
$\textbf{Profit}_{\textbf{seller}}$	43.90 (1.87)	48.28^{***} (0.64)	39.21^{**} (0.87)	$52.93^{\dagger\dagger\dagger}$	$52.93^{\dagger\dagger\dagger}$	$52.93^{\dagger\dagger\dagger}$
$\textbf{Rate}_{\textbf{reject}}$ **[%]**	— —	— —	30.88 (0.05)	—	—	$0.00^{\dagger\dagger\dagger}$
Efficiency [%]	86.25 (2.48)	85.38 (0.89)	89.63 (1.07)	$100.00^{\dagger\dagger\dagger}$	$100.00^{\dagger\dagger\dagger}$	$100.00^{\dagger\dagger\dagger}$
No. of obs.	2,400	2,400	2,400	—	—	—

Table 4.1: Average bids, prices, buyer profit, efficiency, and rejection rate across treatments (all rounds). Standard errors in parentheses. All tests are Wilcoxon rank-sum tests.
Note: If no trade took place in the *negotiation* treatment for additional demand, to ensure comparability with the other treatments, this price is fictitiously set to 1.30.
H_0: Identical to uniform price treatment, H_1: greater or smaller than uniform price treatment; $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.
H_0: Observed = Theoretical, H_1: Observed $\neq$ Theoretical; $^{\dagger}p < 0.1$, $^{\dagger\dagger}p < 0.05$, $^{\dagger\dagger\dagger}p < 0.01$.

significant difference in buyer profit between the uniform price and negotiation treatments ($p = 0.80$).

We therefore find partial support for hypothesis 12, which predicted no difference in buyer profits across treatments. The variances in buyer profits are different too.

Result 15. *While theory predicts the lowest variance in buyer profit, we find the variance to be the highest in the uniform price treatment. Using an F-test on independent observations, we find significantly higher variances in buyer profit in the uniform price treatment vis-à-vis the price premium treatment ($p = 0.04$). There is no significant difference in variances between the uniform price and negotiation treatments ($p = 0.29$).*

We therefore find no support for hypothesis 13.

Let us consider in more detail the additional demand. We hypothesised that trade would always take place for the additional demand in the negotiation treatment. This is not what we see. In approximately one-third of the rounds in which the buyer had additional demand, trade did not take place because the winning seller's price for additional demand exceeded the buyer's maximum willingness-to-pay.

The high rejection rate in the negotiation treatment is an indication that money is being left on the table. Indeed, had the buyer followed the theoretical prediction of setting their maximum willingness-to-pay to $v = 1.30$ instead of 0.67, ceteris paribus, their average profit would have been significantly larger at 110.45 vis-à-vis both the negotiation and price premium treatments (both $p < 0.01$). Observe that trade not taking place is a two-fold blow to the buyer: first, because it is always profitable for the buyer to trade, and second, because sellers whose price requests for the additional demand were rejected should revise their belief about the attractiveness of winning the auction downward and thus respond by bidding less aggressively in future auctions, leading to higher prices for the fixed demand.

But it is not only the buyers whose maximum willingness-to-pay is too low when it comes to the additional demand. While sellers bid most aggressively for the fixed quantity in the negotiation treatment, their bids for the additional quantity are not high enough to make up for this. To see this, note that the average seller profit for the additional demand is given by 7.6 in the uniform price treatment, 20.0 in the price premium treatment, and 8.6 in the negotiation treatment. Theoretically, seller profits should be increasing across these treatments. Another way to look at this is to compare the requested price premiums in the negotiation treatment to the price premium of 0.30 in the price premium treatment. On average, sellers in the negotiation treatment requested a price premium of 0.08, which is significantly lower than 0.30 ($p < 0.01$). However, it is close to the buyer's average accepted price premium of 0.09. In order to attain the same profit on the additional demand as in the price premium treatment—assuming the buyer always accepts the price request—sellers would have to increase their price premium in the negotiation treatment by on average 0.23. Lower bids for the fixed quantity, coupled with lower price premiums for the additional quantity, result in significantly lower sellers profits than in both the price premium treatment and uniform price treatments (both $p < 0.05$).

Lastly, efficiency is significantly lower than theoretically predicted in all treatments (all $p < 0.01$), and significantly higher in the negotiation treatment than in the price premium treatment ($p < 0.05$).

Looking holistically at the results: recall that the theoretical prediction is that, compared to the uniform price treatment, sellers stand to make more profit on the additional demand in the price premium and negotiation treatments. Our results show that sellers submit significantly more aggressive bids to win the contract for the fixed quantity in these two treatments. More than that, our results exhibit the theoretical ranking of bids and prices for the fixed demand across the treatments.

Our results for the additional demand in the negotiation treatment follows the theoretical prediction less closely. This is driven by sellers who bid too aggressively for the fixed demand, given that the buyers are only willing to accept low price premiums for the additional demand.

Exploratory analysis

In this section, we will conduct an exploratory analysis of the experimental results. We saw above that while sellers are in line with the theoretical prediction of bidding most aggressively in the negotiation treatment, their price requests for the additional demand are not high enough to make up for the low prices for the fixed demand. As such, seller profit is lowest in the negotiation treatment. A potential explanation for sellers' behaviour comes by comparing the negotiation treatment with the price premium treatment. In the price premium treatment, sellers receive a fixed per-unit premium for the additional demand. While there is demand uncertainty in the price premium treatment, there is no uncertainty over whether or not trade will take place should the additional demand realise. Rather, trade will always take place and sellers know for which price they will sell any additional units. Of course, in the negotiation treatment, this in not the case: it is possible that negotiations break down for the additional quantity. The fact that negotiations break down regularly leads sellers to request conservative price premiums for the additional quantity.

While uncertainty about trade taking place may be a plausible explanation for sellers' low price premiums, the behaviour of buyers is harder to explain. At 0.67, the buyer's average maximum willingness-to-pay is far away from the theoretical prediction. As described above, if buyers had followed their dominant strategy of setting the maximum willingness-to-pay to $v = 1.30$, ceteris paribus, they would have made significantly more profit. In ultimatum bargaining, fairness concerns are frequently cited as a explanation for deviations from the theoretical predictions (see e.g. Van Dijk et al. (2004)). The argument goes that subjects are willing to reject profitable but unbalanced splits as a result of their preference for fairness. That being said, in our specific case, we can rule out fairness concerns as a potential explanation for buyer behaviour with a high degree of certainty. This is because the results show that the total profit for the additional demand is, if anything, split in favour of the buyer. When simply accepting the seller's requested price for additional quantity, buyers would, on average, receive more than four-fifths of the total profit. Note that this share is computed using the winning seller's *actual* cost. If instead the buyer inverted eq. (4.3) to estimate the winning seller's costs, the buyer share of profits becomes even larger. Note that in 82% of cases in which negotiations broke down, buyers would have made a higher profit than the seller.

At the end of the experiment, all subjects participated in a survey intended to characterise subjects' personality traits. To this end, subjects answered a battery of questions taken from Brandstätter (1988). All questions are on a nine-point Likert scale, and we averaged across all questions for a given personality trait to determine a subject's level of Risk-aversion, Self-control, Independence, Tough-mindedness, and Extraversion. The summary statistics of the survey are reported in table 4.2, while the battery of questions can be found in section 4.A.5.

	Risk aversion	**Self control**	**Independence**	**Tough-mindedness**	**Extraversion**
Min.	2.0	3.0	3.0	1.0	2.6
1. Quartile	5.0	4.8	4.7	4.5	4.6
Median	6.0	5.7	5.3	5.5	5.6
Mean	5.9	5.8	5.4	5.5	5.6
3. Quartile	7.0	6.5	6.0	6.5	6.4
Max.	9.0	9.0	9.0	9.0	8.6

Table 4.2: Summary statistics of survey results. All 180 subjects participated in the survey.

An interesting question to ask is whether the deviations from the theoretical predictions we observe are correlated with differences in personality. Table 4.3 reports the results of a random-effects panel regression of sellers' bids for the fixed demand on their costs, the respective treatment, their personality traits, as well as the respective round number, which allows us to capture learning effects over time. We also include interactions between the treatment dummies and the personality traits, which allows the personality traits to have differing effects across treatments. Looking at the full model that includes the round number, the panel regression demonstrates that personality traits do not seem to be driving bidding behaviour for the fixed demand: the only significant effect we find is that higher extraversion is associated with lower bids in the price premium treatment. Overall though, we observe that bids are decreasing over time, that is, sellers' bids become more aggressive. Note that the coefficient on the round number remains negative and significant even when only considering the negotiation treatment ($p < 0.01$). This is a surprising finding. Recall that we observe a higher-than-expected rejection rate for the additional quantity in the negotiation treatment. This means that seller profits on the additional demand are lower than theoretically predicted, which should result in sellers bidding *higher* for the fixed demand over time to make up for this. We observe the opposite. Another way to check for learning effects is to regress (winning) sellers' profits on the round number. In both cases, a random-effects panel regression yields a coefficient on the round number that is both negative and not statistically significant ($p = 0.42$ for all sellers, and $p = 0.16$ when considering only winning sellers). As a result, we find no evidence for learning on the sellers' side.

	Dependent variable: Bid for the fixed demand			
	(1)	(2)	(3)	(4)
Constant	0.272***	0.393***	0.410***	0.101
	(0.064)	(0.096)	(0.096)	(0.113)
Cost	0.670***	0.670***	0.670***	0.670***
	(0.005)	(0.005)	(0.005)	(0.005)
Risk aversion	0.006	−0.007	−0.007	−0.008
	(0.005)	(0.008)	(0.008)	(0.008)
Self control	−0.004	−0.003	−0.003	−0.008
	(0.006)	(0.009)	(0.009)	(0.009)
Independence	0.013*	0.005	0.005	0.004
	(0.008)	(0.010)	(0.010)	(0.010)
Tough minded	−0.004	−0.013	−0.013	−0.008
	(0.006)	(0.010)	(0.010)	(0.009)
Extraversion	−0.003	0.015	0.015	0.008
	(0.006)	(0.010)	(0.010)	(0.009)
TreatmentPP		−0.024	−0.024	−0.090
		(0.147)	(0.147)	(0.136)
TreatmentNeg		−0.200	−0.200	−0.299**
		(0.151)	(0.151)	(0.141)
TreatmentPP× Risk aversion		0.017	0.017	0.013
		(0.012)	(0.012)	(0.011)
TreatmentNeg× Risk aversion		0.005	0.005	0.004
		(0.011)	(0.011)	(0.010)
TreatmentPP × Self control		−0.013	−0.013	−0.0004
		(0.015)	(0.015)	(0.014)
TreatmentNeg× Self control		0.003	0.003	0.013
		(0.014)	(0.014)	(0.013)
TreatmentPP × Independence		0.008	0.008	0.003
		(0.017)	(0.017)	(0.016)
TreatmentNeg× Independence		0.009	0.009	0.006
		(0.017)	(0.017)	(0.015)
TreatmentPP × Tough-mindedness		0.013	0.013	0.004
		(0.013)	(0.013)	(0.012)
TreatmentNeg× Tough-mindedness		0.015	0.015	0.014
		(0.014)	(0.014)	(0.013)
TreatmentPP × Extraversion		−0.031**	−0.031**	−0.015
		(0.013)	(0.013)	(0.013)
TreatmentNeg× Extraversion		−0.011	−0.011	−0.0003
		(0.014)	(0.014)	(0.013)
Round			−0.001***	−0.001***
			(0.0001)	(0.0001)
BidPreTreatment				0.005***
				(0.001)
Observations	4,800	4,800	4,800	4,800
R^2	0.804	0.805	0.807	0.807
Adjusted R^2	0.804	0.804	0.806	0.806
F Statistic	19,631.950***	19,676.510***	19,970.450***	20,002.920***
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 4.3: Random effects panel regression of the sellers' bids for the fixed demand across treatments and including personality traits. BidPreTreatment refers to a subject's average bid in the pre-treatment.

However, personality traits are correlated with sellers' price requests for the additional demand in the negotiation treatment; the results of the random-effects panel regression are reported in table 4.4. Higher levels of self control, tough-mindedness, and extraversion are associated with lower price requests, while higher independence is associated with higher price requests. It is of note that neither risk aversion nor the round number seems to influence sellers' price requests. How do these findings compare with the literature? First of all, risk aversion is often found to influence behaviour in ultimatum bargaining, see e.g. Holt and Laury (2002); Fehr and Gächter (2000), something we do not find in our context. Nevertheless, generally speaking, personality traits have been shown to influence behaviour in negotiation contexts. That being said, the direction of the effect that a given personality trait has on behaviour is anything but homogeneous. A review of this literature is provided in John et al. (1999).

	Dependent variable:	
	Price request for the additional demand	
	(1)	(2)
Constant	0.443***	0.449***
	(0.153)	(0.153)
Cost	0.692***	0.692***
	(0.011)	(0.011)
Risk aversion	0.00004	0.00004
	(0.009)	(0.009)
Self control	−0.027**	−0.027**
	(0.013)	(0.013)
Independence	0.057***	0.057***
	(0.017)	(0.017)
Tough mindedness	−0.016	−0.016
	(0.013)	(0.013)
Extraversion	−0.030**	−0.030**
	(0.012)	(0.012)
Round		−0.0003
		(0.0003)
Observations	1,600	1,600
R^2	0.710	0.710
Adjusted R^2	0.709	0.709
F Statistic	3,893.291***	3,895.743***
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 4.4: Random effects panel regression of the sellers' requests for the additional demand including personality traits.

Consider the buyers next. Table 4.5 reports the results of a random-effects panel regression of buyers' maximum willingness-to-pay for the additional demand on the auction price and their personality traits. Looking at model (1), we see that buyers' maximum willingness-to-pay is positively associated with the auction price. The theoretical prediction is that the maximum willingness-to-pay is independent of the auction price. In model (2), we include the buyer's expected payoff in the pre-treatment as a proxy for their sophistication. However, it is of note that this coefficient is not significant in any of the models. In model (3), we see that the personality do not directly influence the willingness-to-pay. However, they do correlate with how strong the influence of the auction price is. As seen in model (4), higher degrees of tough-mindedness and extraversion are associated with a smaller influence of the auction price on the buyer's maximum willingness-to-pay for additional demand. It is worth noting that the maximum willingness-to-pay does increase over time, i.e. moves closer to the theoretical prediction. However, this does not have an influence on buyer profits: the coefficient of a random-effects panel regression of the buyers' profits on the round number yields $p = 0.27$.

	Dependent variable:			
	Maximum willingness-to-pay for the additional demand			
	(1)	(2)	(3)	(4)
Constant	0.435***	0.411	0.549	1.187**
	(0.038)	(0.295)	(0.424)	(0.533)
AuctionPrice	0.469***	0.469***	0.492***	−0.777*
	(0.038)	(0.038)	(0.037)	(0.405)
PayoffPreTreatment		0.001	0.030	0.028
		(0.018)	(0.025)	(0.030)
Risk aversion			−0.046	−0.041
			(0.031)	(0.039)
Self control			−0.037	−0.065
			(0.043)	(0.056)
Independence			−0.005	0.021
			(0.050)	(0.064)
Tough-mind.			−0.067*	−0.128***
			(0.039)	(0.048)
Extraversion			0.029	−0.027
			(0.048)	(0.061)
Round			0.003***	0.003***
			(0.001)	(0.001)
AuctionPrice × Risk aversion				−0.010
				(0.031)
AuctionPrice × Self control				0.055
				(0.045)
AuctionPrice × Independence				−0.044
				(0.053)
AuctionPrice × Tough-mind.				0.131***
				(0.032)
AuctionPrice × Extraversion				0.107**
				(0.050)
Observations	800	800	800	800
R^2	0.159	0.159	0.202	0.237
Adjusted R^2	0.158	0.157	0.194	0.224
F Statistic	150.889***	151.042***	199.780***	244.124***
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 4.5: Random effects panel regression of the buyers' maximum willingness-to-pay for the additional demand including personality traits. PayoffPreTreatment refers to a subject's expected payoff in the pre-treatment.

Taken together, when looking at the additional demand, we find that sellers' price requests for the additional demand are decreasing over time (albeit not statistically significantly), while buyers' maximum willingness-to-pay increases with the rounds. Lower price requests coupled with a higher willingess-to-pay means that that trade in the negotiation treatment is more likely to take place. This should result in lower prices for the additional demand in later rounds. Table 4.6 reports the results of random-effects panel regression of the price for the additional demand in the negotiation treatment. Indeed, as expected, we find a coefficient on the round number that is both negative and significant ($p < 0.05$).

	Dependent variable:	
	Price for additional quantity	
	(1)	(2)
Constant	0.517^{***}	0.568^{***}
	(0.038)	(0.044)
WinnerCost	0.677^{***}	0.670^{***}
	(0.052)	(0.052)
Round		-0.002^{**}
		(0.001)
Observations	800	800
R^2	0.173	0.179
Adjusted R^2	0.172	0.177
F Statistic	167.242^{***}	173.635^{***}

Notes: **p<0.1;* ***p<0.05;* ****p<0.01.*

Table 4.6: Random effects panel regression of the price for additional quantity on the winner's costs and round.

To sum up: for the fixed quantity, our results are in line with theoretical predictions. Across treatments, sellers bid more aggressively the higher their expected profit for the additional demand is. As a result, we observe the theoretical ranking of prices for the fixed demand. For the additional demand, however, our results in the negotiation treatment do not follow the theoretical predictions as closely. Rather, a high rejection rate on the part of buyers means that sellers' profits for the additional demand are lower than expected. While this should result in sellers adjusting their bids upwards for the fixed demand, we find no evidence for this. Our analysis above also shows that personality traits influence negotiation dynamics for both sides, indicating that outcomes are shaped not only by economic incentives but also by behavioural factors.

4.5 Discussion of the model

In this section, we will discuss the robustness of our game theoretic model. We will demonstrate that the theoretical predictions we derive are robust to many alternative formulations of the model, but also discuss formulations for which this is not the case.

Constant marginal production costs In our model, we assume sellers have constant marginal production costs of c. This is without loss of generality as long as the equilibrium quantities traded are equivalent across the three settings. In such a case, even with convex or concave production costs, all three settings yield the same expected buyer expenses. A sufficient condition for identical equilibrium trading quantities is that it is always beneficial to trade. That is, if $\bar{c} < v$, our results hold even for alternative formulations of sellers' production costs.

Bargaining power In the negotiation treatment—mirroring the "fundamental transformation" of Williamson (1985) that we have discussed in the introduction—the winning seller has all the bargaining power in the price negotiation for the additional demand. As a result, the theoretical prediction is that they extract full surplus by requesting a unit price of v for all additional units. What would happen in a case in which sellers had less bargaining power? Consider, for example in the case where buyer and seller engage in Rubinstein bargaining (Rubinstein, 1982). Less bargaining power on the seller's part translates directly into lower expected profit for the additional demand. The seller accounts for this by bidding less aggressively in the auction. In aggregate, their expected payoff remains unchanged. However, note that there may exist a critical threshold for the buyer's bargaining power above which sellers may find it attractive to pool in the auction for the fixed demand in order to conceal their private information. Pooling may be worth it because if a seller's bid is indicative of their costs, the buyer can leverage this information in the negotiation, which, in the extreme case, means sellers would make no profit for the additional demand. An example for a case in which sellers pool in order to conceal their private information can be found in Fugger et al. (2019).

Fallback option In our model, the buyer does not have a fallback option for the additional quantity in the negotiation treatment. That is, if the winning seller's requested price for additional quantity exceeds the buyer's maximum willingness-to-pay, no trade takes place for this additional quantity, leaving both buyer and seller with zero profit for these units. Qualitatively, our results do not change even in the presence of a fallback option. In fact, the model in its current form is mathematically equivalent to a situation in which the buyer can purchase additional units at a unit price of v in case the negotiation fails. If the buyer had an outside option $\bar{c} < r < v$, then r would be the maximum the buyer is willing to pay for additional units. Note that even in this setting, the equivalence between the three settings we consider still holds. The argument is that the better the buyer's outside option is, the lower the seller's expected profit for additional demand becomes, which leads to less aggressive bidding in the auction for the fixed demand. In aggregate, seller profit is unchanged.

Delivery obligation In our model, we have an implicit assumption that the winning seller is always obligated to deliver. For the uniform cost distribution and the parameters we consider in the experiment, this assumption is unproblematic, since sellers never bid below their costs. Hence, it is always beneficial for the winning seller to deliver, since the price they receive exceeds their costs. However, for a sufficiently high price premium, sellers would bid below their costs. For example, any price premium in the experiment greater than 1 makes winning the auction so attractive that sellers would bid below their costs. The same would happen in the negotiation treatment for sufficiently high v. In both these cases, if sellers were not obligated to deliver at a loss, the equivalence we show between the three settings would break down. However, note that this is only true if sellers observe the total demand before having to deliver to the buyer. If the additional quantity were revealed only *after* having delivered the fixed quantity, the equivalency between the three settings holds even if sellers bid below their costs for the fixed quantity.

To sum up, the theoretical result on the equivalency of the three settings we consider holds more generally. To be precise, it holds for alternative assumptions on sellers' production costs, for different distributions of bargaining power between buyer and seller, as well as when sellers bid below their costs for the fixed quantity.

4.6 Conclusion

This paper set out to answer a central question in procurement under demand uncertainty: how should a buyer design an auction when the total demand is unknown? We compared three relevant contracting settings: (1) a *uniform price* setting, in which all units—fixed and additional—are procured at the same per-unit price determined in the auction; (2) a *price premium* setting, in which the buyer defines a fixed contingency in the case of additional demand; and (3) a *negotiation* setting, in which the buyer procures the fixed demand via auction and bargains ex-post with the winning seller for additional units if needed. Our theoretical framework demonstrated that—even when considering differing assumptions about sellers' costs, bargaining power, and delivery obligations—all three formats should yield identical expected profits for buyers and sellers. The intuition is straightforward: sellers adjust their auction bids to reflect the expected value of any additional demand, competing away format-specific advantages in the auction stage. Thus, even though the ex-post pricing rules differ, the ex-ante competitive process equalises outcomes.

We tested these predictions in a controlled laboratory experiment and found that actual behaviour diverges substantially from the theoretical benchmark. In line with the model, sellers did bid more aggressively in formats where they expected higher profits from additional demand, reproducing the predicted ranking of prices for the fixed demand: highest in the uniform price treatment, lower in the price premium treatment, and lowest in the negotiation treatment. However, the negotiation format revealed significant inefficiencies. Buyers often set their maximum willingness-to-pay for additional units far below the theoretical optimum, resulting in failed negotiations in roughly one-third of cases where extra units were needed. This behaviour meant passing up profitable trades and reduced the potential cost advantage of negotiation. The behaviour of sellers in the negotiation treatment also does not follow the theoretical predictions. Rather than requesting the high premiums predicted by theory, they demanded amounts far lower than the fixed premium offered in the price premium treatment. Consequently, their profits from additional demand were modest, and total seller profits in the negotiation treatment fell below those in both the price premium and uniform price settings. Personality traits such as independence, self-control, and extraversion correlated with differences in negotiation behaviour, suggesting that individual characteristics can influence outcomes in subtle but important ways.

The broader implication is that behavioural frictions—aggressive bargaining by buyers, cautious premium-setting by sellers, and possibly personality-driven decision patterns—can erode the theoretical equivalence of procurement formats under demand uncertainty. While the uniform price format performed relatively well for buyers in our setting, this advantage likely arose from the absence of negotiation breakdowns. With improved negotiation strategies—such as buyers consistently setting willingness-to-pay closer to their valuation and sellers recognising their bargaining position—the flexible negotiation format could match or outperform the other settings in terms of buyer surplus and allocative efficiency.

For practitioners, these results carry two key messages. First, simpler procurement formats may be more reliable in delivering predictable outcomes. Second, where negotiation is feasible and potentially beneficial, targeted interventions—such as clear decision rules, transparency about counterparties' incentives, or structured bargaining protocols—could help unlock latent gains and reduce the frequency of failed trades. Future research could explore how changes in bargaining power, the presence of fallback supply options, or alternative cost structures affect both theoretical predictions and experimental behaviour. Extending the model to richer strategic environments and testing interventions in the laboratory would help bridge the gap between normative auction theory and the behavioural realities of procurement under uncertainty. In doing so, we can better understand not only which formats are theoretically equivalent, but also which are most likely to deliver desirable outcomes in practice.

Appendix of Chapter 4

4.A Experiment: screenshots

4.A.1 Pre-treatment

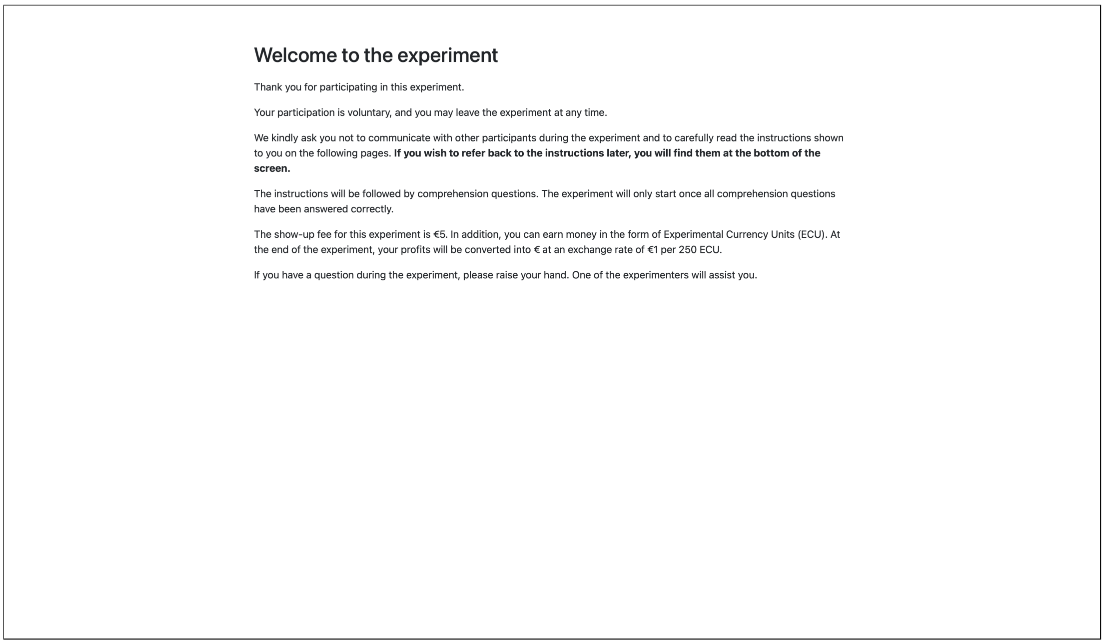

Welcome to the experiment

Thank you for participating in this experiment.

Your participation is voluntary, and you may leave the experiment at any time.

We kindly ask you not to communicate with other participants during the experiment and to carefully read the instructions shown to you on the following pages. **If you wish to refer back to the instructions later, you will find them at the bottom of the screen.**

The instructions will be followed by comprehension questions. The experiment will only start once all comprehension questions have been answered correctly.

The show-up fee for this experiment is €5. In addition, you can earn money in the form of Experimental Currency Units (ECU). At the end of the experiment, your profits will be converted into € at an exchange rate of €1 per 250 ECU.

If you have a question during the experiment, please raise your hand. One of the experimenters will assist you.

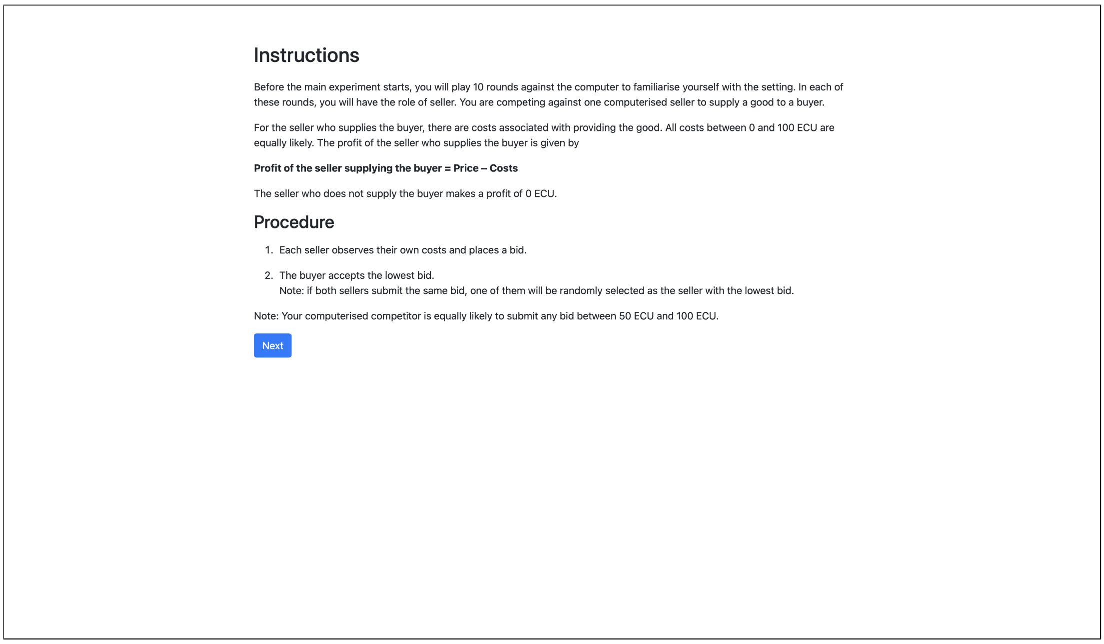

Instructions

Before the main experiment starts, you will play 10 rounds against the computer to familiarise yourself with the setting. In each of these rounds, you will have the role of seller. You are competing against one computerised seller to supply a good to a buyer.

For the seller who supplies the buyer, there are costs associated with providing the good. All costs between 0 and 100 ECU are equally likely. The profit of the seller who supplies the buyer is given by

Profit of the seller supplying the buyer = Price – Costs

The seller who does not supply the buyer makes a profit of 0 ECU.

Procedure

1. Each seller observes their own costs and places a bid.
2. The buyer accepts the lowest bid.
   Note: if both sellers submit the same bid, one of them will be randomly selected as the seller with the lowest bid.

Note: Your computerised competitor is equally likely to submit any bid between 50 ECU and 100 ECU.

Next

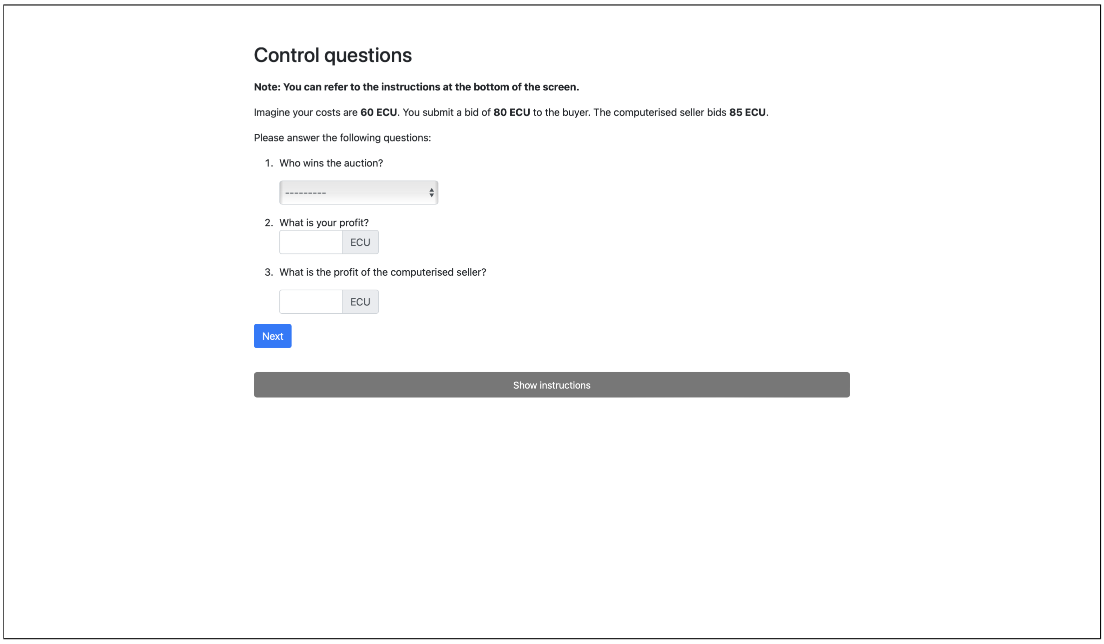
Control questions
Note: You can refer to the instructions at the bottom of the screen.
Imagine your costs are 60 ECU. You submit a bid of 80 ECU to the buyer. The computerised seller bids 85 ECU.
Please answer the following questions:
1. Who wins the auction?
----
2. What is your profit?
ECU
3. What is the profit of the computerised seller?
ECU
Next
Show instructions

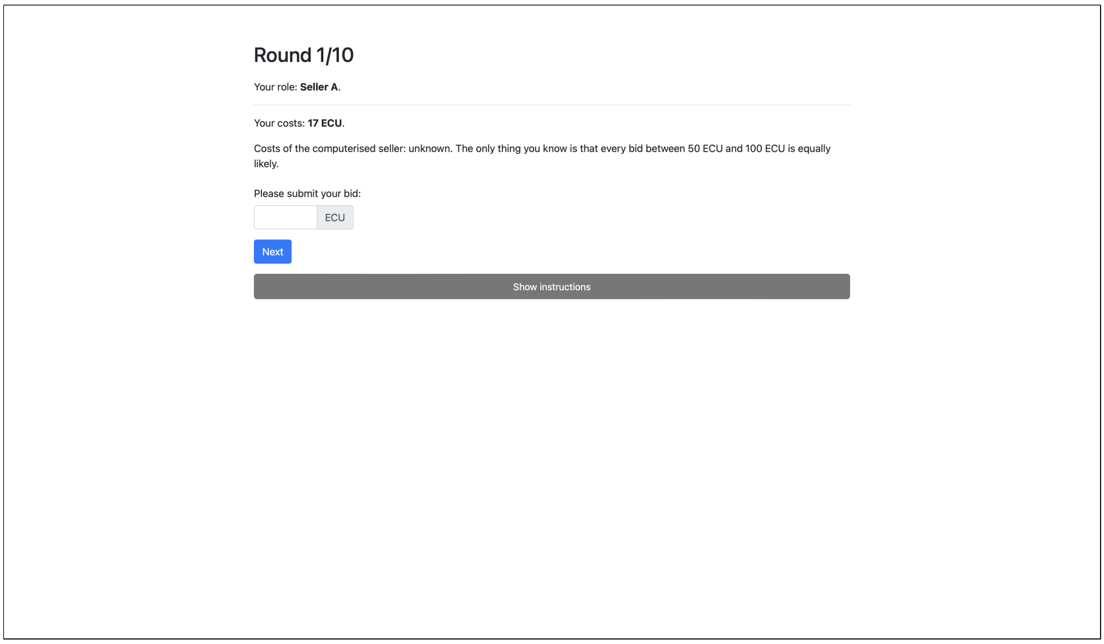
Round 1/10
Your role: Seller A.
Your costs: 17 ECU.
Costs of the computerised seller: unknown. The only thing you know is that every bid between 50 ECU and 100 ECU is equally likely.
Please submit your bid:
ECU
Next
Show instructions

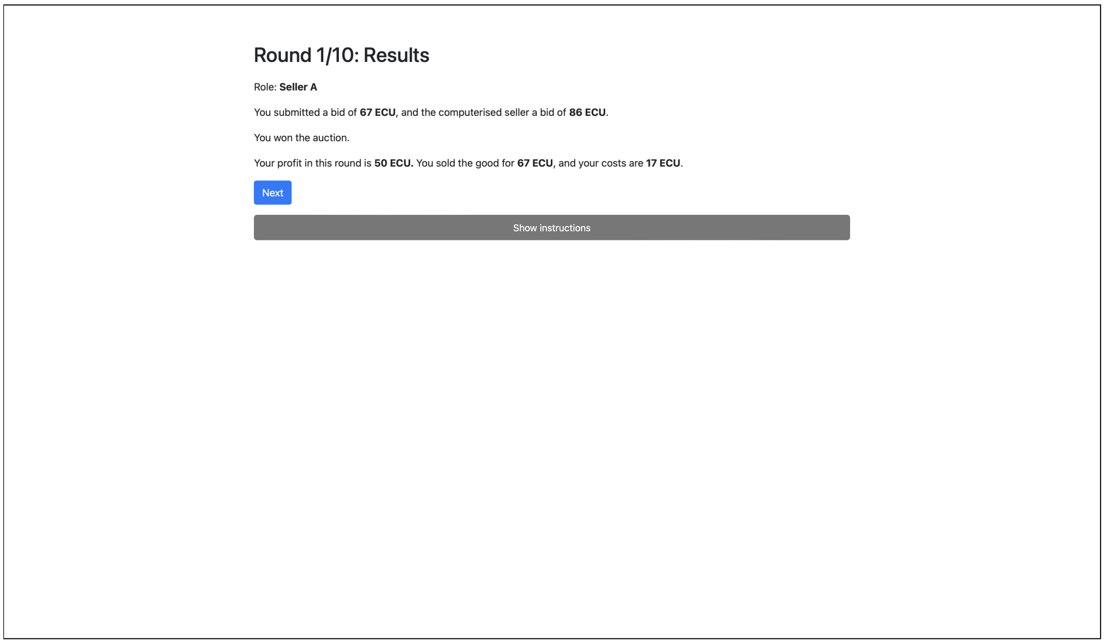
Round 1/10: Results
Role: Seller A
You submitted a bid of 67 ECU, and the computerised seller a bid of 86 ECU.
You won the auction.
Your profit in this round is 50 ECU. You sold the good for 67 ECU, and your costs are 17 ECU.
Next
Show instructions

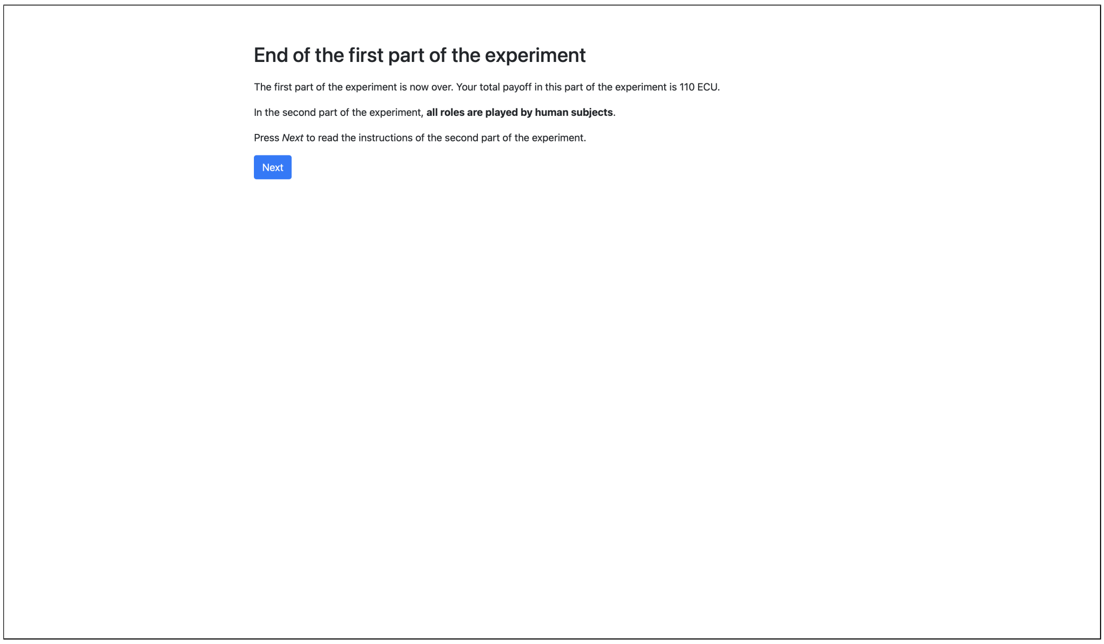
End of the first part of the experiment
The first part of the experiment is now over. Your total payoff in this part of the experiment is 110 ECU.
In the second part of the experiment, all roles are played by human subjects.
Press Next to read the instructions of the second part of the experiment.
Next

4.A.2 Uniform price

Instructions

The experiment consists of 40 rounds, and in each round, you will have the same role. The possible roles are buyer or seller. At the beginning of each round, one buyer and two sellers will be randomly matched.

In each round, the buyer will conduct an auction to purchase a certain number of units of a good that they sell for 1.30 ECU per unit.

In each round, the buyer faces a certain demand of 100 units. With a probability of 50%, the buyer faces an additional demand of another 100 units.

The buyer's profit is equal to the total demand multiplied with the difference between their per-unit selling price of 1.30 ECU and the per-unit price they pay to the selected seller.

Buyer's Profit = Total Demand x (1.30 ECU – Per-unit Price)

Sellers' profit

In each round, all sellers receive a fixed payment of 30 ECU. The seller who supplies the buyer makes an additional profit.

For the seller who supplies the buyer, there are per-unit costs associated with providing the good. At the beginning of each round, these costs are randomly and independently determined for each seller. All costs between 0 and 1 ECU are equally likely. The profit of the seller who supplies the buyer is given by

Profit of the seller supplying the buyer = Total Demand x (Per-unit Price – Per-unit Costs) + 30 ECU

The seller who does not supply the buyer makes a profit of 30 ECU.

Procedure

1. Each seller observes their own per-unit costs and places a per-unit bid.
2. The buyer accepts the lowest bid.
3. The total demand realises.
4. Subjects are informed of their profits in that round.

Next

Control questions

Note: You can refer to the instructions at the bottom of the screen.

Please answer the following questions:

1. In each round, what is the buyer's total demand?

   ----

2. Imagine you are in the role of the seller. Your competitor submits the lowest bid. What is your profit?

   ECU

Next

Show instructions

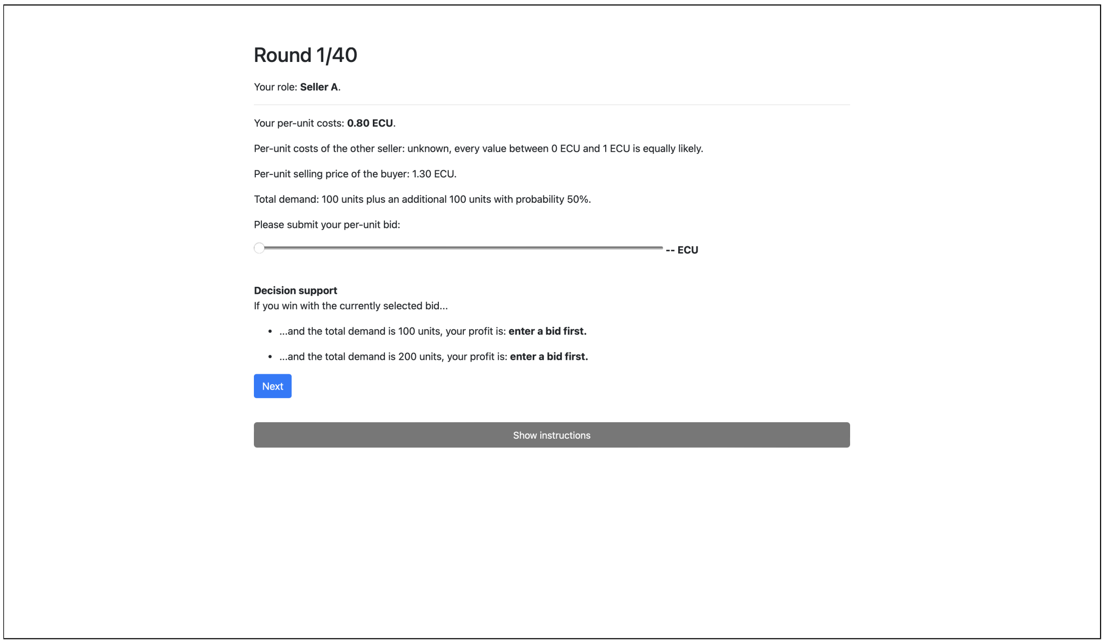
Round 1/40
Your role: Seller A.
Your per-unit costs: 0.80 ECU.
Per-unit costs of the other seller: unknown, every value between 0 ECU and 1 ECU is equally likely.
Per-unit selling price of the buyer: 1.30 ECU.
Total demand: 100 units plus an additional 100 units with probability 50%.
Please submit your per-unit bid:
-- ECU
Decision support
If you win with the currently selected bid...
...and the total demand is 100 units, your profit is: enter a bid first.
...and the total demand is 200 units, your profit is: enter a bid first.
Next
Show instructions

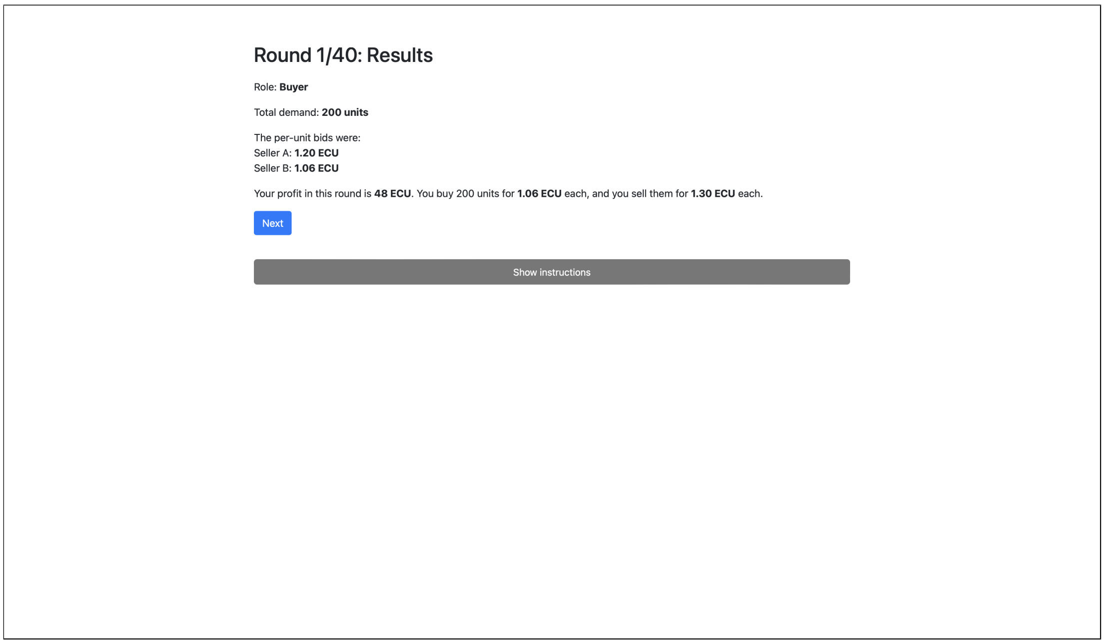
Round 1/40: Results
Role: Buyer
Total demand: 200 units
The per-unit bids were:
Seller A: 1.20 ECU
Seller B: 1.06 ECU
Your profit in this round is 48 ECU. You buy 200 units for 1.06 ECU each, and you sell them for 1.30 ECU each.
Next
Show instructions

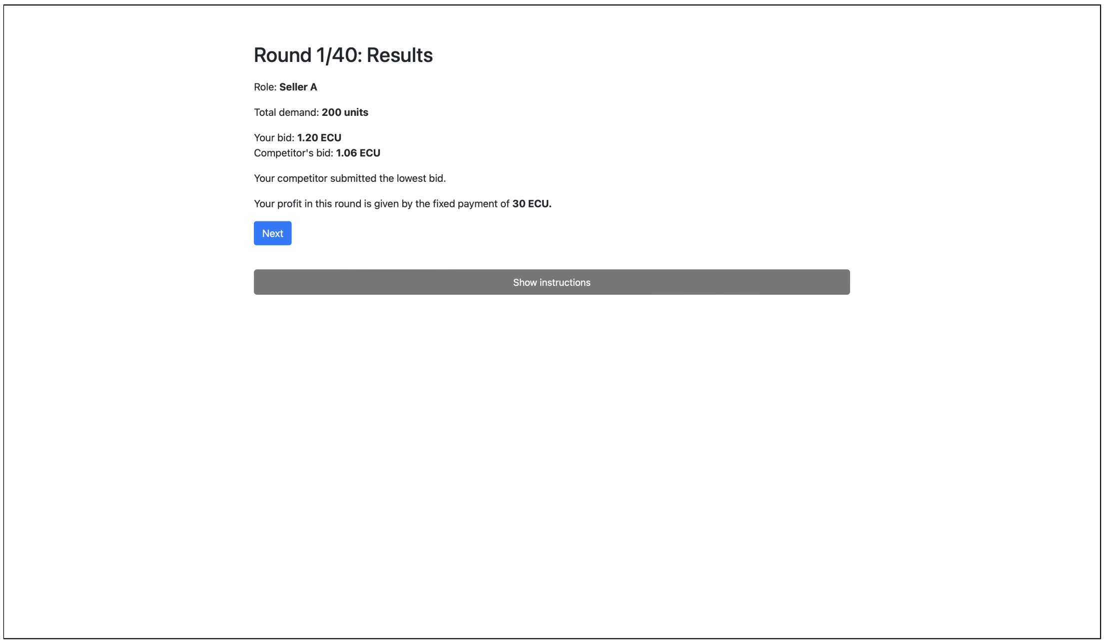
Round 1/40: Results
Role: Seller A
Total demand: 200 units
Your bid: 1.20 ECU
Competitor's bid: 1.06 ECU
Your competitor submitted the lowest bid.
Your profit in this round is given by the fixed payment of 30 ECU.
Next
Show instructions

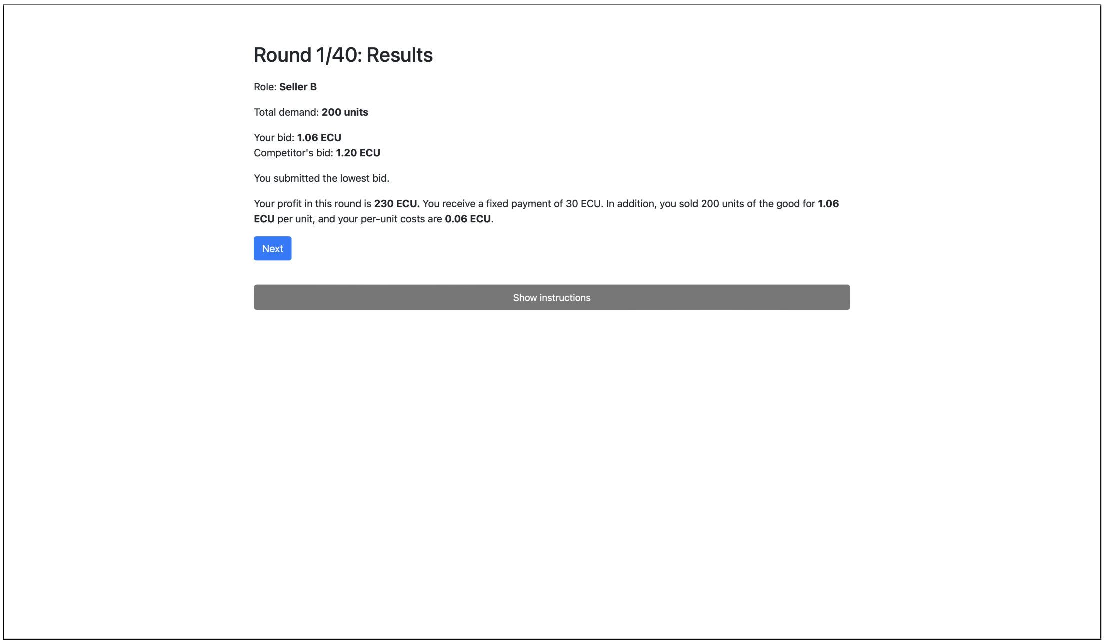
Round 1/40: Results
Role: Seller B
Total demand: 200 units
Your bid: 1.06 ECU
Competitor's bid: 1.20 ECU
You submitted the lowest bid.
Your profit in this round is 230 ECU. You receive a fixed payment of 30 ECU. In addition, you sold 200 units of the good for 1.06 ECU per unit, and your per-unit costs are 0.06 ECU.
Next
Show instructions

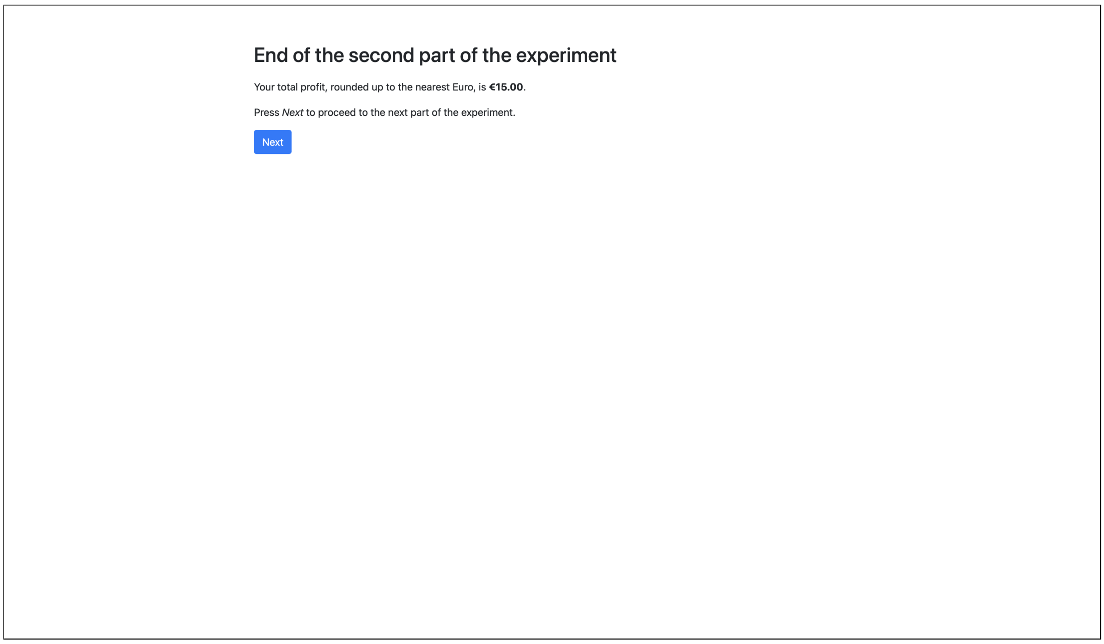
End of the second part of the experiment
Your total profit, rounded up to the nearest Euro, is €15.00.
Press Next to proceed to the next part of the experiment.
Next

4.A.3 Price premium

Instructions

The experiment consists of 40 rounds, and in each round, you will have the same role. The possible roles are buyer or seller. At the beginning of each round, one buyer and two sellers will be randomly matched.

In each round, the buyer will conduct an auction to purchase a certain number of units of a good that they sell for 1.30 ECU per unit.

In each round, the buyer faces a certain demand of 100 units. With a probability of 50%, the buyer faces an additional demand of another 100 units.

If the additional demand realises, the buyer will pay the winning seller a premium of 0.30 ECU on top of the auction price for every unit of additional demand.

Buyer's profit

- If the total demand is **100 units**, the buyer's profit is equal to the total demand multiplied with the difference between their per-unit selling price of 1.30 ECU and the per-unit auction price they pay:
  100 x (1.30 ECU – Per-unit Price)
- If the total demand is **200 units**, the buyer pays the per-unit auction price for the first 100 units. For the additional demand of 100 units, the buyer pays the per-unit auction price plus a per-unit premium of 0.30 ECU. The buyer's profit is therefore equal to:
  100 x (1.30 ECU – Per-unit Price) + 100 x (1.30 ECU – [Per-unit Price + 0.30 ECU])

Sellers' profit

In each round, all sellers receive a fixed payment of 30 ECU. The seller who supplies the buyer makes an additional profit.

Profit of the seller who supplies the buyer

For the seller who supplies the buyer, there are per-unit costs associated with providing the good. At the beginning of each round, these costs are randomly and independently determined for each seller. All costs between 0 and 1 ECU are equally likely.

- If the total demand is **100 units**, the profit of the seller who supplies the buyer is given by:
  100 x (Per-unit Price – Per-unit Costs) + 30 ECU
- If the total demand is **200 units**, the profit of the seller who supplies the buyer is given by:
  100 x (Per-unit Price – Per-unit Costs) + 100 x ([Per-unit Price + 0.30 ECU] – Per-unit Costs) + 30 ECU

Profit of the seller who does not supply the buyer

The seller who does not supply the buyer makes a profit of 30 ECU.

Procedure

1. Each seller observes their own per-unit costs and places a per-unit bid.
2. The buyer accepts the lowest bid.
3. The total demand realises.
4. Subjects are informed of their profits in that round.

Next

Control questions

Note: You can refer to the instructions at the bottom of the screen.

Please answer the following questions:

1. In each round, what is the buyer's total demand?

   ----

2. Imagine you are in the role of the seller. Your competitor submits the lowest bid. What is your profit?

   ECU

3. Imagine the buyer's demand in a round is 200 units. The winning seller receives a per-unit premium of 0.30 ECU on which units?

Next

Show instructions

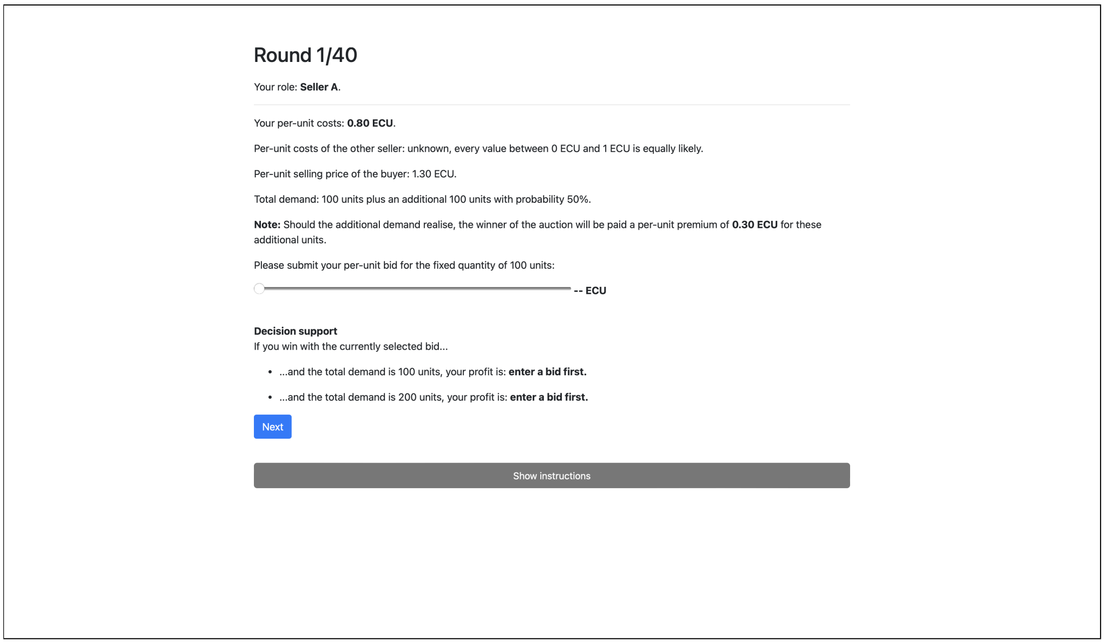

Round 1/40

Your role: **Seller A**.

Your per-unit costs: **0.80 ECU**.

Per-unit costs of the other seller: unknown, every value between 0 ECU and 1 ECU is equally likely.

Per-unit selling price of the buyer: 1.30 ECU.

Total demand: 100 units plus an additional 100 units with probability 50%.

Note: Should the additional demand realise, the winner of the auction will be paid a per-unit premium of **0.30 ECU** for these additional units.

Please submit your per-unit bid for the fixed quantity of 100 units:

-- ECU

Decision support
If you win with the currently selected bid...

- ...and the total demand is 100 units, your profit is: **enter a bid first.**
- ...and the total demand is 200 units, your profit is: **enter a bid first.**

Next

Show instructions

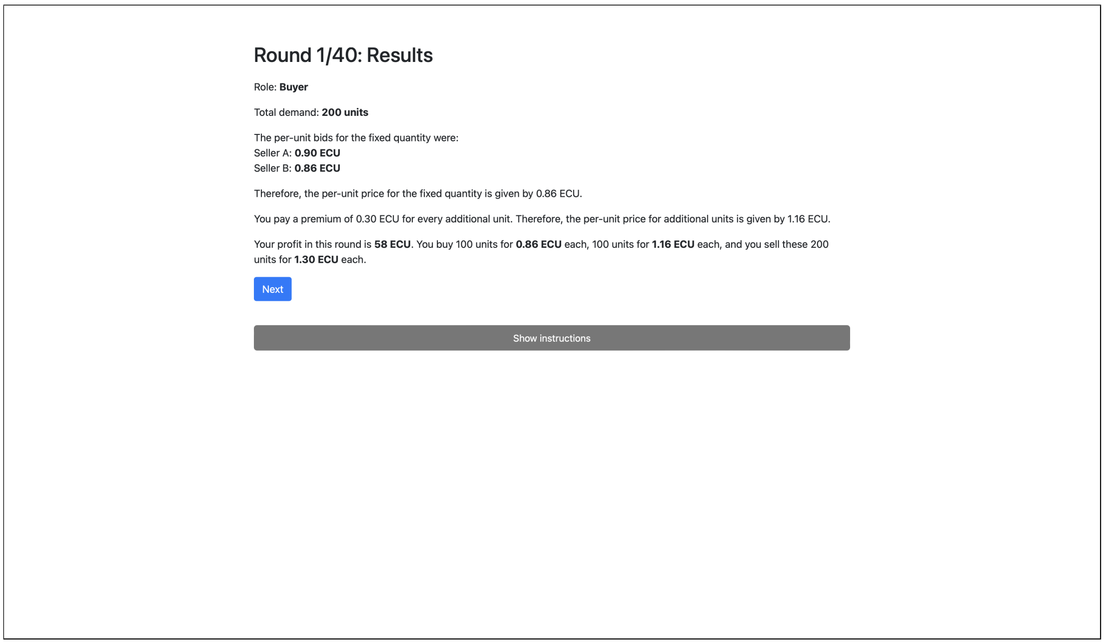

Round 1/40: Results

Role: **Buyer**

Total demand: **200 units**

The per-unit bids for the fixed quantity were:
Seller A: **0.90 ECU**
Seller B: **0.86 ECU**

Therefore, the per-unit price for the fixed quantity is given by 0.86 ECU.

You pay a premium of 0.30 ECU for every additional unit. Therefore, the per-unit price for additional units is given by 1.16 ECU.

Your profit in this round is **58 ECU**. You buy 100 units for **0.86 ECU** each, 100 units for **1.16 ECU** each, and you sell these 200 units for **1.30 ECU** each.

Next

Show instructions

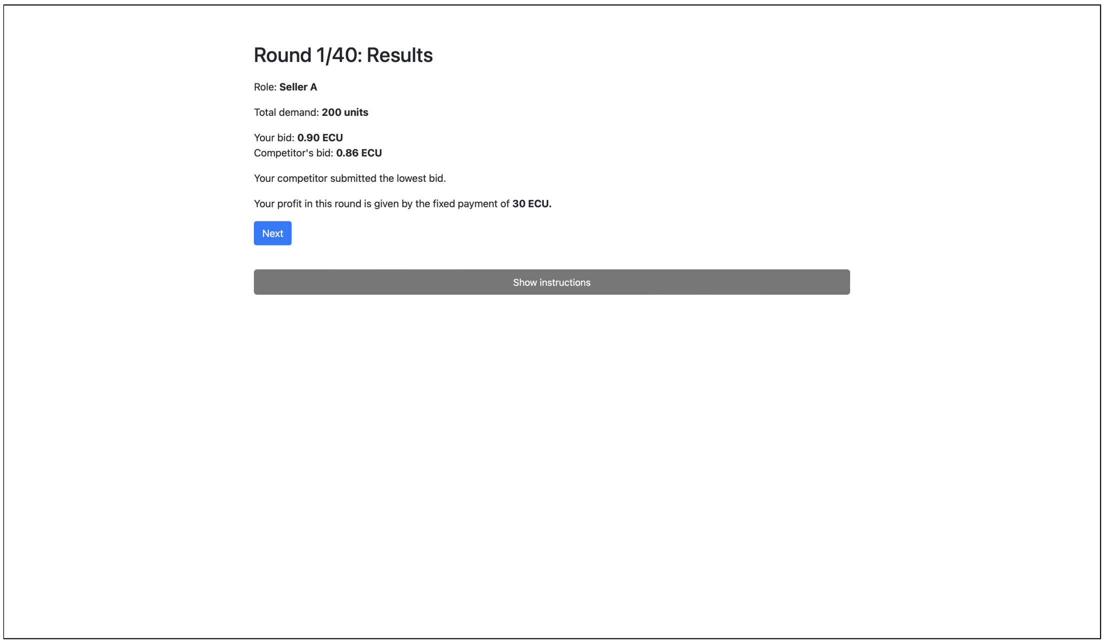
Round 1/40: Results
Role: Seller A
Total demand: 200 units
Your bid: 0.90 ECU
Competitor's bid: 0.86 ECU
Your competitor submitted the lowest bid.
Your profit in this round is given by the fixed payment of 30 ECU.
Next
Show instructions

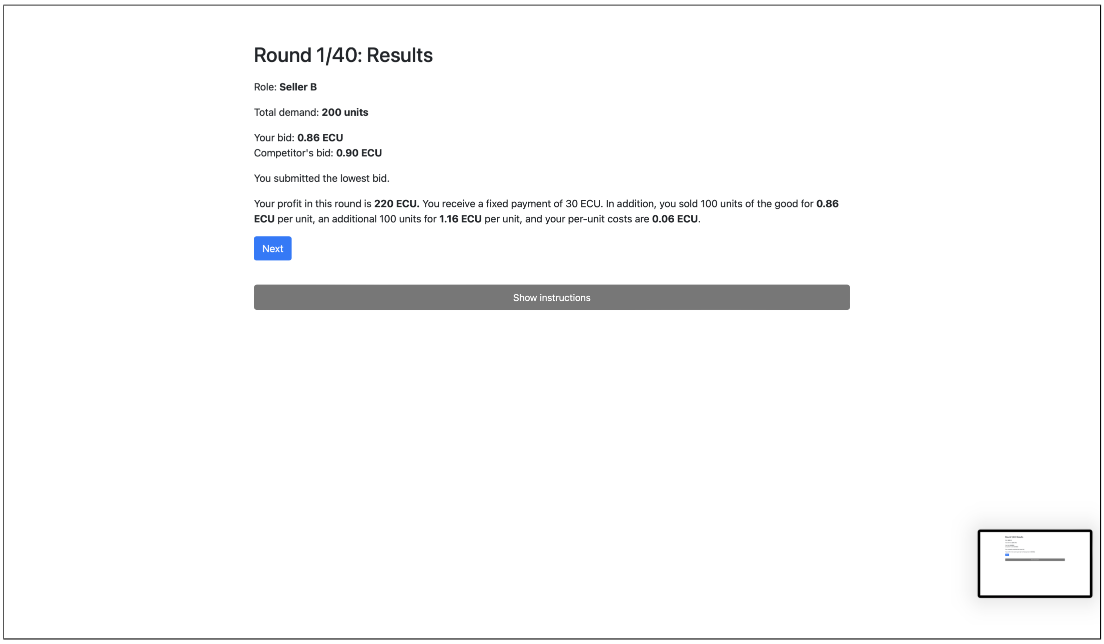
Round 1/40: Results
Role: Seller B
Total demand: 200 units
Your bid: 0.86 ECU
Competitor's bid: 0.90 ECU
You submitted the lowest bid.
Your profit in this round is 220 ECU. You receive a fixed payment of 30 ECU. In addition, you sold 100 units of the good for 0.86 ECU per unit, an additional 100 units for 1.16 ECU per unit, and your per-unit costs are 0.06 ECU.
Next
Show instructions

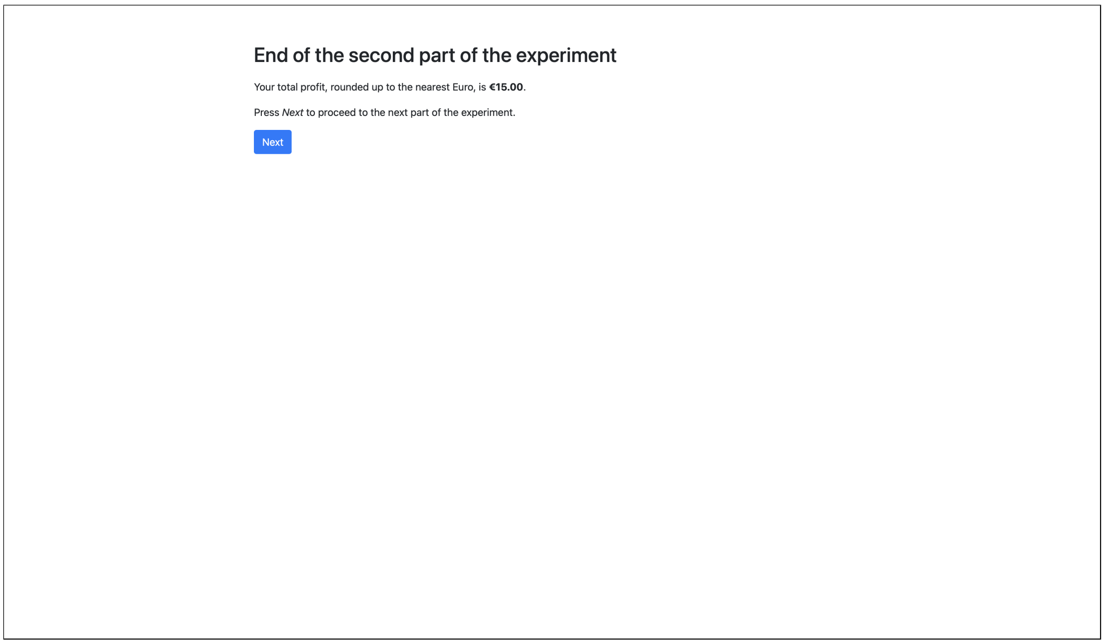
End of the second part of the experiment
Your total profit, rounded up to the nearest Euro, is €15.00.
Press Next to proceed to the next part of the experiment.
Next

4.A.4 Negotiation

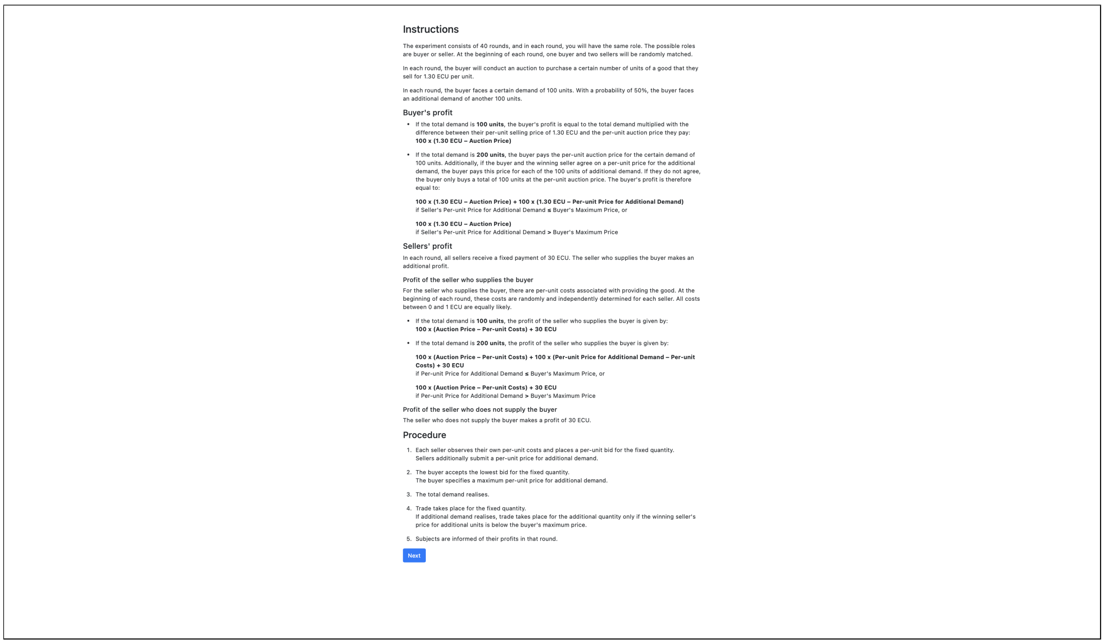

Instructions

The experiment consists of 40 rounds, and in each round, you will have the same role. The possible roles are buyer or seller. At the beginning of each round, one buyer and two sellers will be randomly matched.

In each round, the buyer will conduct an auction to purchase a certain number of units of a good that they sell for 1.30 ECU per unit.

In each round, the buyer faces a certain demand of 100 units. With a probability of 50%, the buyer faces an additional demand of another 100 units.

Buyer's profit

- If the total demand is **100 units**, the buyer's profit is equal to the total demand multiplied with the difference between their per-unit selling price of 1.30 ECU and the per-unit auction price they pay:
  100 x (1.30 ECU – Auction Price)
- If the total demand is **200 units**, the buyer pays the per-unit auction price for the certain demand of 100 units. Additionally, if the buyer and the winning seller agree on a per-unit price for the additional demand, the buyer pays this price for each of the 100 units of additional demand. If they do not agree, the buyer only buys a total of 100 units at the per-unit auction price. The buyer's profit is therefore equal to:

  100 x (1.30 ECU – Auction Price) + 100 x (1.30 ECU – Per-unit Price for Additional Demand)
  if Seller's Per-unit Price for Additional Demand ≤ Buyer's Maximum Price, or

  100 x (1.30 ECU – Auction Price)
  if Seller's Per-unit Price for Additional Demand > Buyer's Maximum Price

Sellers' profit

In each round, all sellers receive a fixed payment of 30 ECU. The seller who supplies the buyer makes an additional profit.

Profit of the seller who supplies the buyer

For the seller who supplies the buyer, there are per-unit costs associated with providing the good. At the beginning of each round, these costs are randomly and independently determined for each seller. All costs between 0 and 1 ECU are equally likely.

- If the total demand is **100 units**, the profit of the seller who supplies the buyer is given by:
  100 x (Auction Price – Per-unit Costs) + 30 ECU
- If the total demand is **200 units**, the profit of the seller who supplies the buyer is given by:

  100 x (Auction Price – Per-unit Costs) + 100 x (Per-unit Price for Additional Demand – Per-unit Costs) + 30 ECU
  if Per-unit Price for Additional Demand ≤ Buyer's Maximum Price, or

  100 x (Auction Price – Per-unit Costs) + 30 ECU
  if Per-unit Price for Additional Demand > Buyer's Maximum Price

Profit of the seller who does not supply the buyer

The seller who does not supply the buyer makes a profit of 30 ECU.

Procedure

1. Each seller observes their own per-unit costs and places a per-unit bid for the fixed quantity.
   Sellers additionally submit a per-unit price for additional demand.
2. The buyer accepts the lowest bid for the fixed quantity.
   The buyer specifies a maximum per-unit price for additional demand.
3. The total demand realises.
4. Trade takes place for the fixed quantity.
   If additional demand realises, trade takes place for the additional quantity only if the winning seller's price for additional units is below the buyer's maximum price.
5. Subjects are informed of their profits in that round.

Next

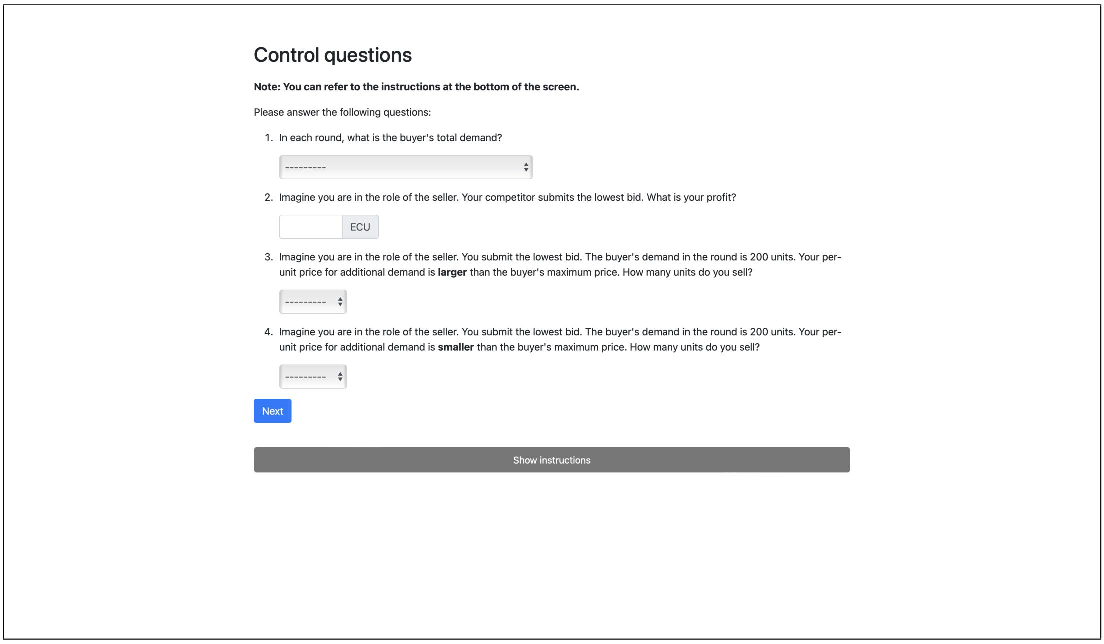

Control questions

Note: You can refer to the instructions at the bottom of the screen.

Please answer the following questions:

1. In each round, what is the buyer's total demand?
   ----
2. Imagine you are in the role of the seller. Your competitor submits the lowest bid. What is your profit?
   ECU
3. Imagine you are in the role of the seller. You submit the lowest bid. The buyer's demand in the round is 200 units. Your per-unit price for additional demand is **larger** than the buyer's maximum price. How many units do you sell?
   ----
4. Imagine you are in the role of the seller. You submit the lowest bid. The buyer's demand in the round is 200 units. Your per-unit price for additional demand is **smaller** than the buyer's maximum price. How many units do you sell?

Next

Show instructions

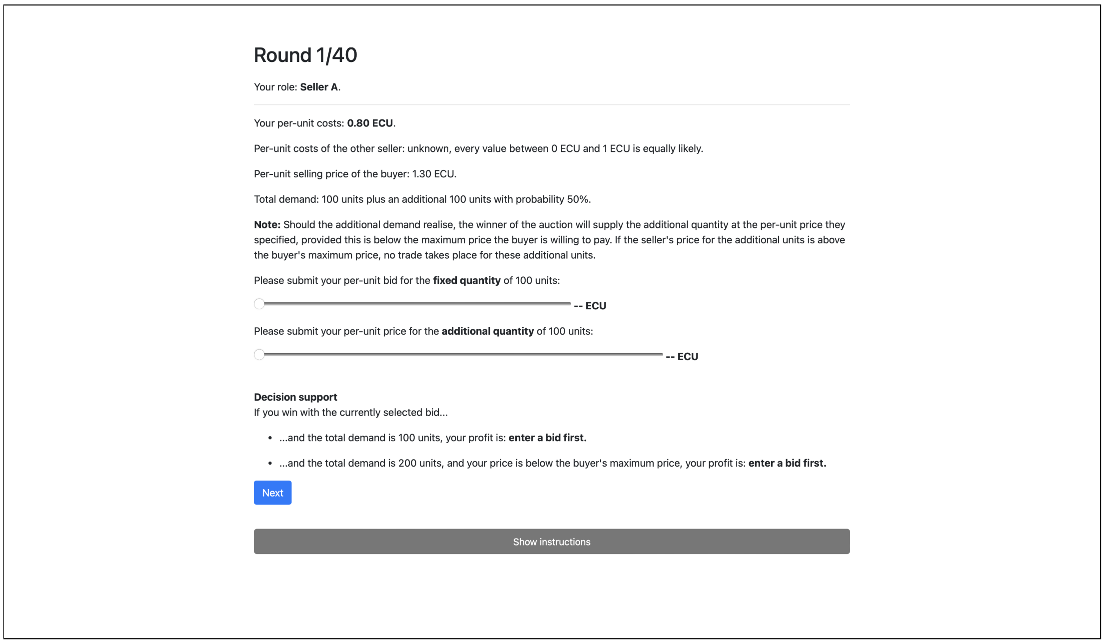
Round 1/40
Your role: Seller A.
Your per-unit costs: 0.80 ECU.
Per-unit costs of the other seller: unknown, every value between 0 ECU and 1 ECU is equally likely.
Per-unit selling price of the buyer: 1.30 ECU.
Total demand: 100 units plus an additional 100 units with probability 50%.
Note: Should the additional demand realise, the winner of the auction will supply the additional quantity at the per-unit price they specified, provided this is below the maximum price the buyer is willing to pay. If the seller's price for the additional units is above the buyer's maximum price, no trade takes place for these additional units.
Please submit your per-unit bid for the fixed quantity of 100 units:
-- ECU
Please submit your per-unit price for the additional quantity of 100 units:
-- ECU
Decision support
If you win with the currently selected bid...
...and the total demand is 100 units, your profit is: enter a bid first.
...and the total demand is 200 units, and your price is below the buyer's maximum price, your profit is: enter a bid first.
Next
Show instructions

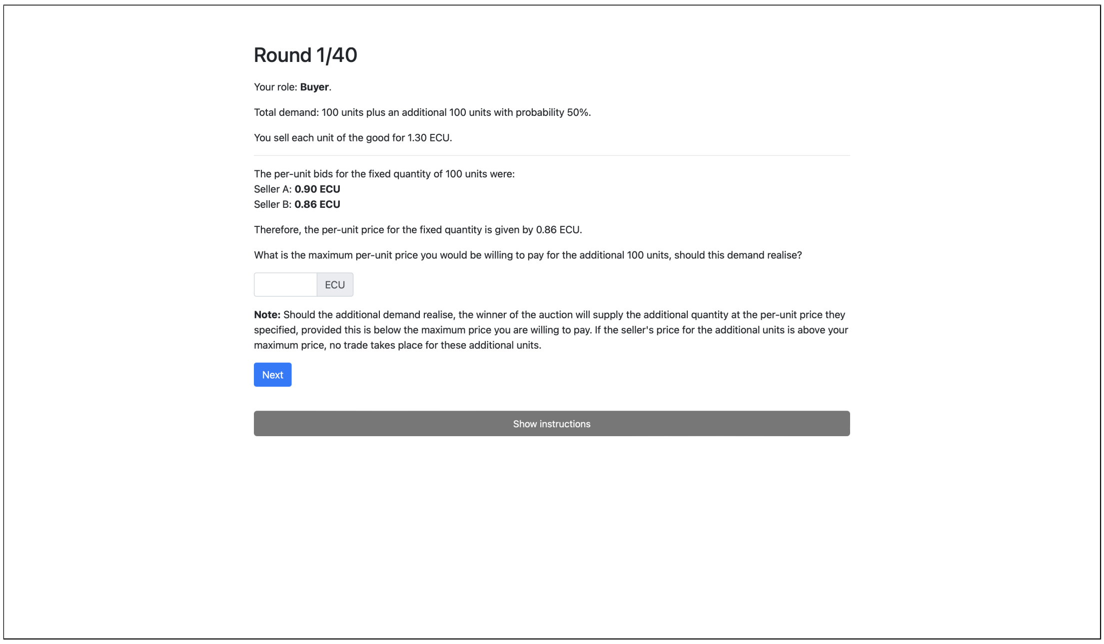
Round 1/40
Your role: Buyer.
Total demand: 100 units plus an additional 100 units with probability 50%.
You sell each unit of the good for 1.30 ECU.
The per-unit bids for the fixed quantity of 100 units were:
Seller A: 0.90 ECU
Seller B: 0.86 ECU
Therefore, the per-unit price for the fixed quantity is given by 0.86 ECU.
What is the maximum per-unit price you would be willing to pay for the additional 100 units, should this demand realise?
ECU
Note: Should the additional demand realise, the winner of the auction will supply the additional quantity at the per-unit price they specified, provided this is below the maximum price you are willing to pay. If the seller's price for the additional units is above your maximum price, no trade takes place for these additional units.
Next
Show instructions

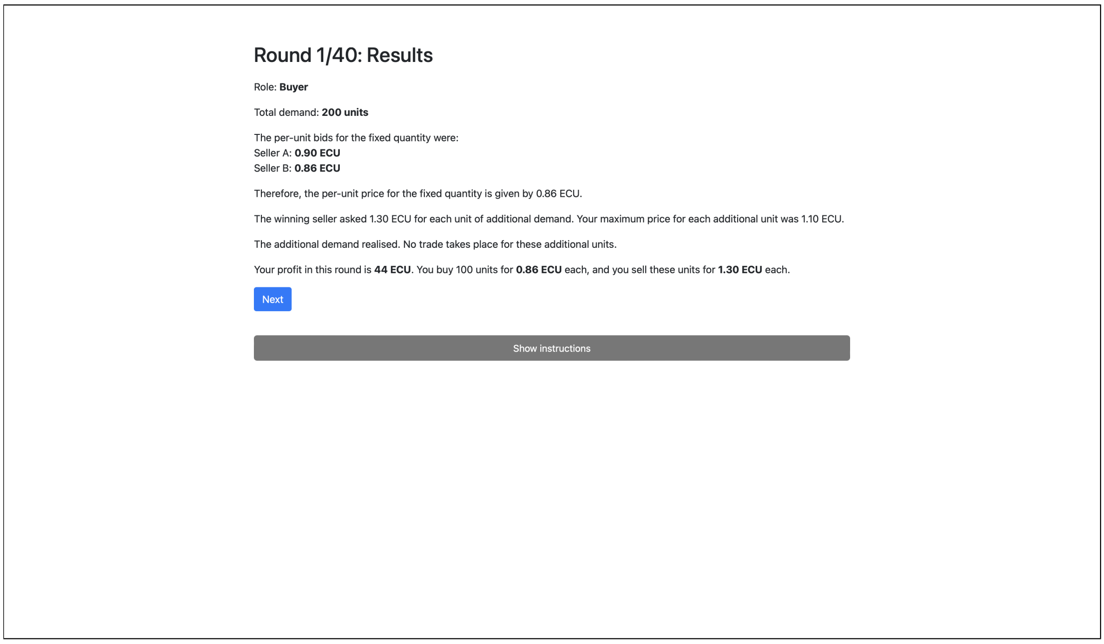
Round 1/40: Results
Role: Buyer
Total demand: 200 units
The per-unit bids for the fixed quantity were:
Seller A: 0.90 ECU
Seller B: 0.86 ECU
Therefore, the per-unit price for the fixed quantity is given by 0.86 ECU.
The winning seller asked 1.30 ECU for each unit of additional demand. Your maximum price for each additional unit was 1.10 ECU.
The additional demand realised. No trade takes place for these additional units.
Your profit in this round is 44 ECU. You buy 100 units for 0.86 ECU each, and you sell these units for 1.30 ECU each.
Next
Show instructions

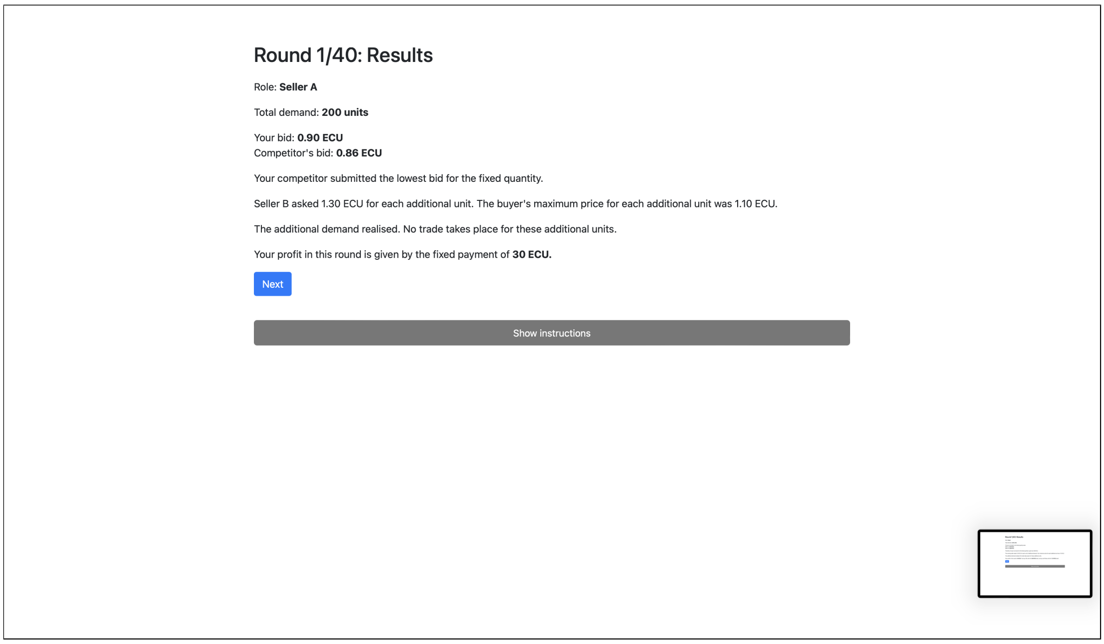
Round 1/40: Results
Role: Seller A
Total demand: 200 units
Your bid: 0.90 ECU
Competitor's bid: 0.86 ECU
Your competitor submitted the lowest bid for the fixed quantity.
Seller B asked 1.30 ECU for each additional unit. The buyer's maximum price for each additional unit was 1.10 ECU.
The additional demand realised. No trade takes place for these additional units.
Your profit in this round is given by the fixed payment of 30 ECU.
Next
Show instructions

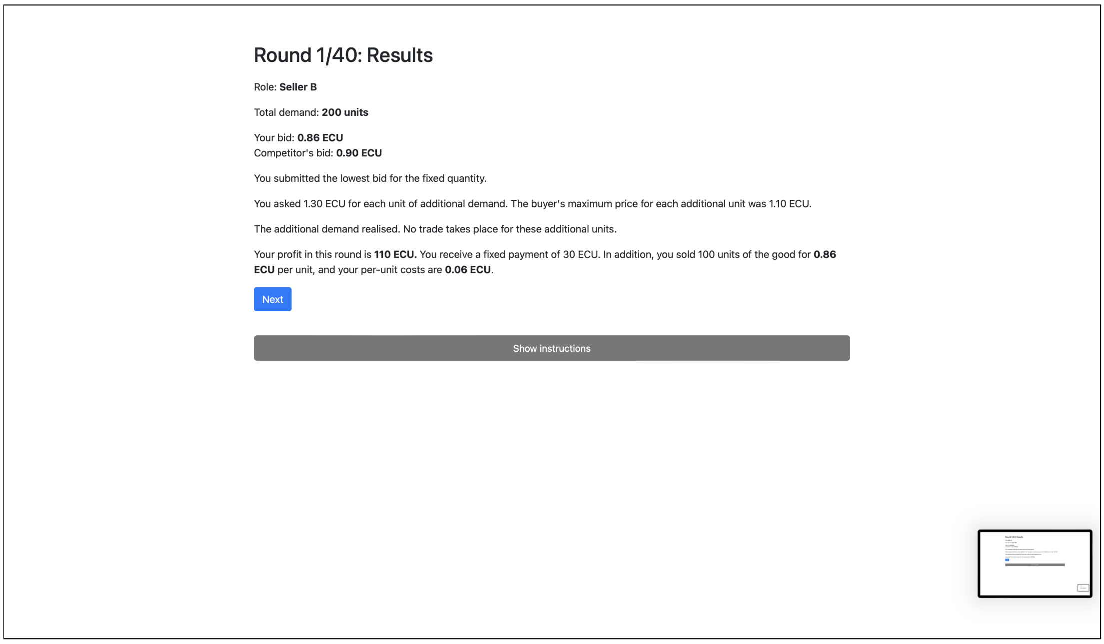
Round 1/40: Results
Role: Seller B
Total demand: 200 units
Your bid: 0.86 ECU
Competitor's bid: 0.90 ECU
You submitted the lowest bid for the fixed quantity.
You asked 1.30 ECU for each unit of additional demand. The buyer's maximum price for each additional unit was 1.10 ECU.
The additional demand realised. No trade takes place for these additional units.
Your profit in this round is 110 ECU. You receive a fixed payment of 30 ECU. In addition, you sold 100 units of the good for 0.86 ECU per unit, and your per-unit costs are 0.06 ECU.
Next
Show instructions

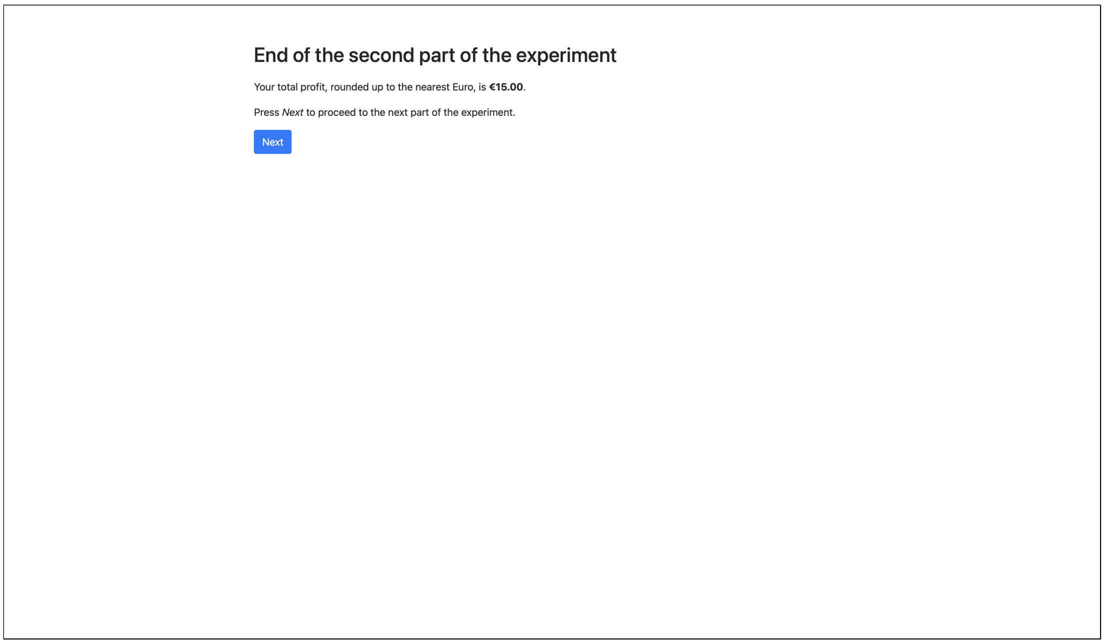
End of the second part of the experiment
Your total profit, rounded up to the nearest Euro, is €15.00.
Press Next to proceed to the next part of the experiment.
Next

4.A.5 Survey

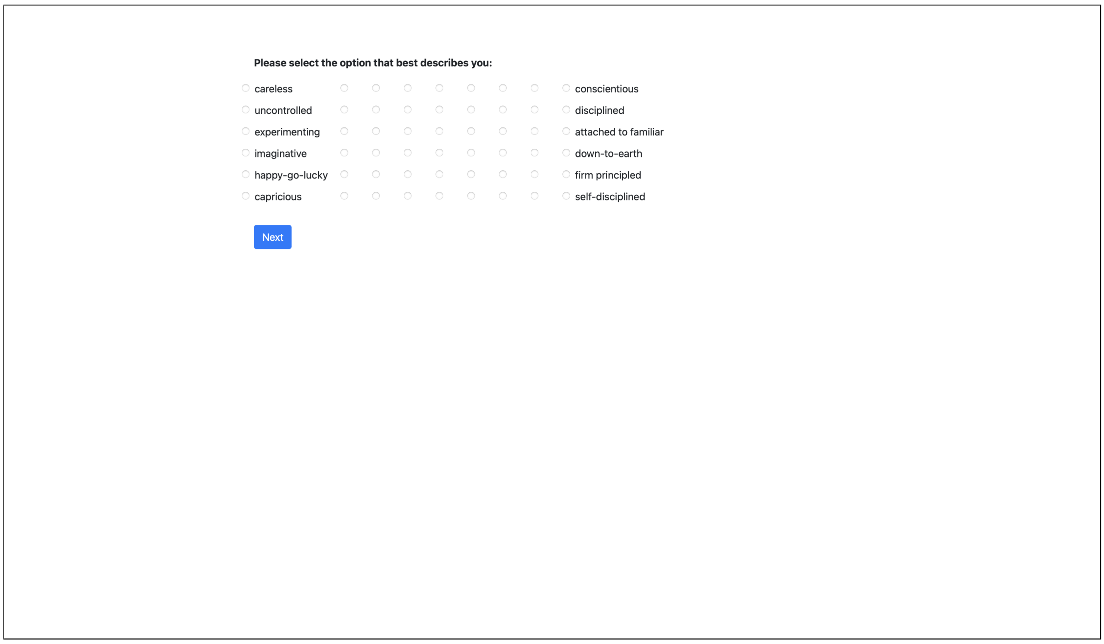

Please select the option that best describes you:

careless	conscientious
uncontrolled	disciplined
experimenting	attached to familiar
imaginative	down-to-earth
happy-go-lucky	firm principled
capricious	self-disciplined

Next

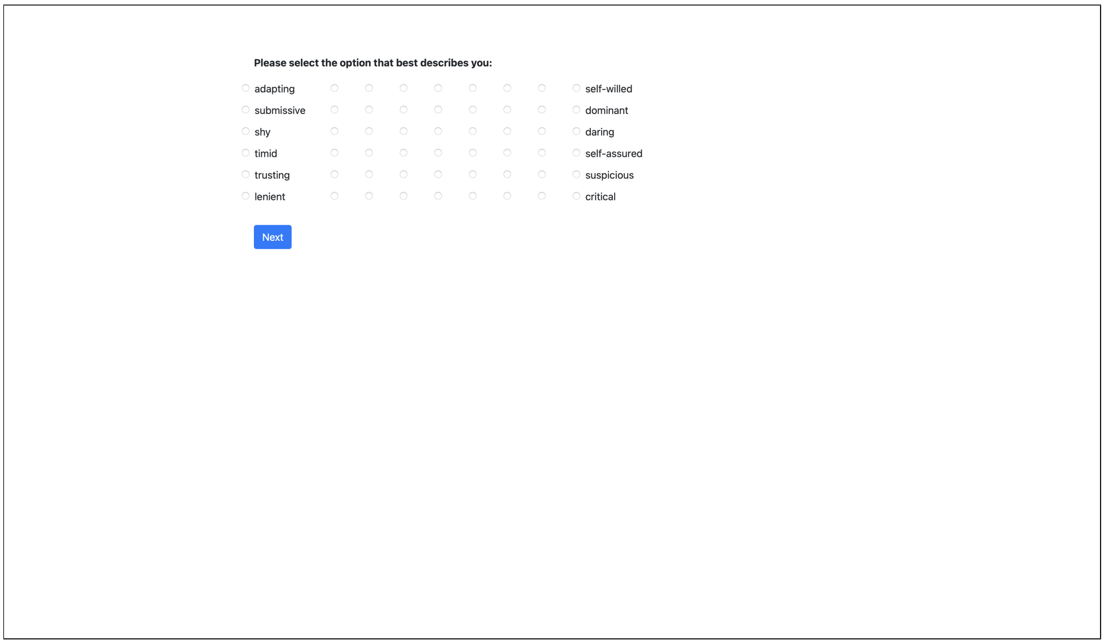

Please select the option that best describes you:

adapting	self-willed
submissive	dominant
shy	daring
timid	self-assured
trusting	suspicious
lenient	critical

Next

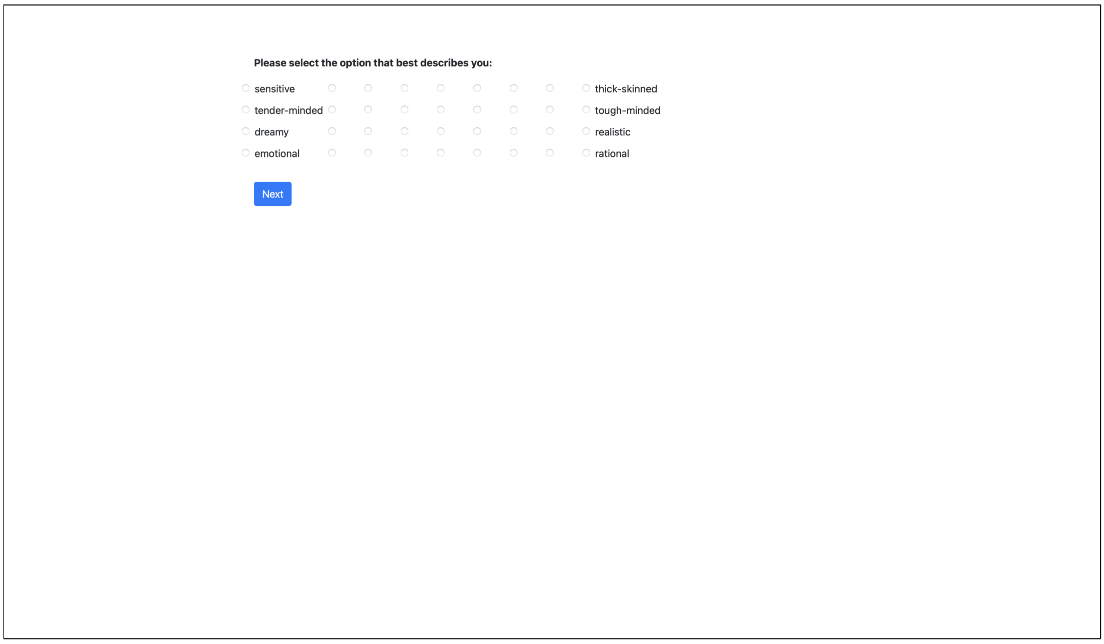
Please select the option that best describes you:
sensitive
thick-skinned
tender-minded
tough-minded
dreamy
realistic
emotional
rational
Next

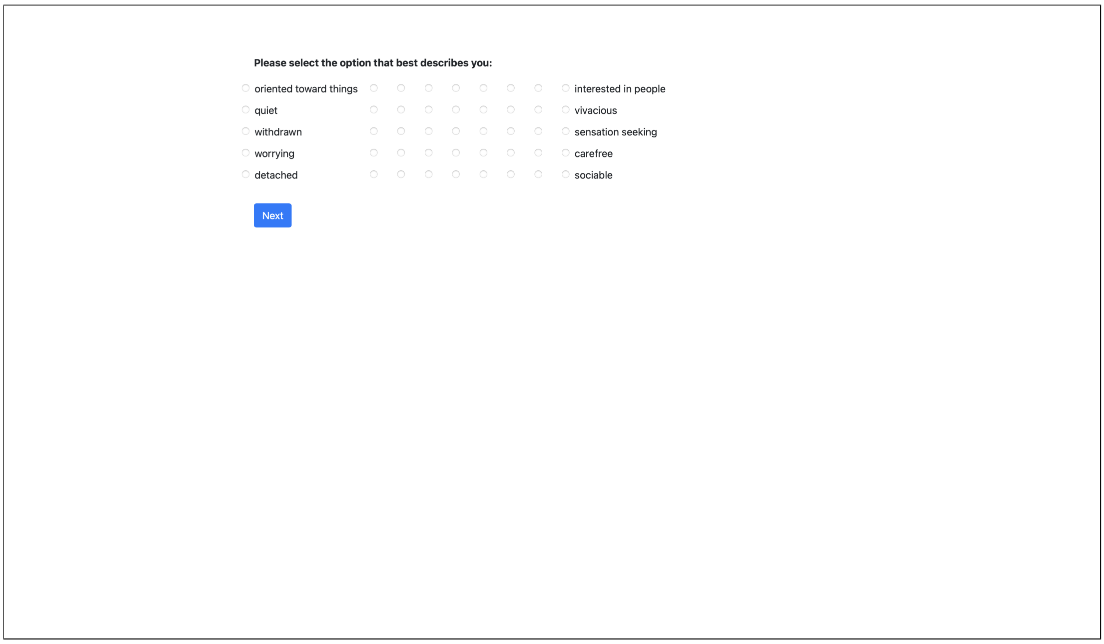
Please select the option that best describes you:
oriented toward things
interested in people
quiet
vivacious
withdrawn
sensation seeking
worrying
carefree
detached
sociable
Next

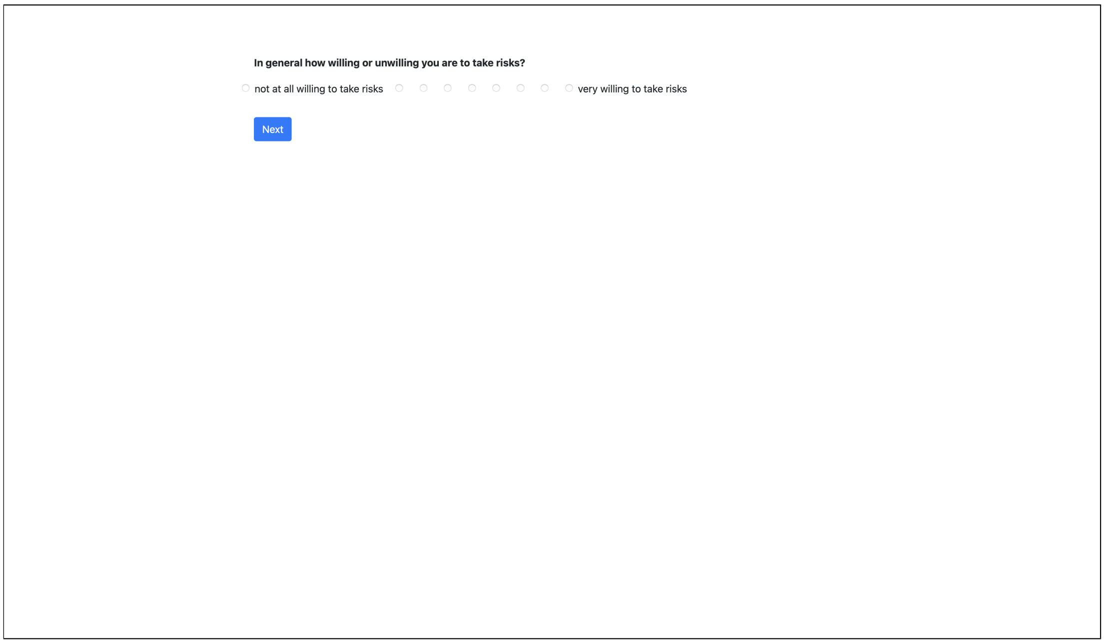
In general how willing or unwilling you are to take risks?
not at all willing to take risks
very willing to take risks
Next

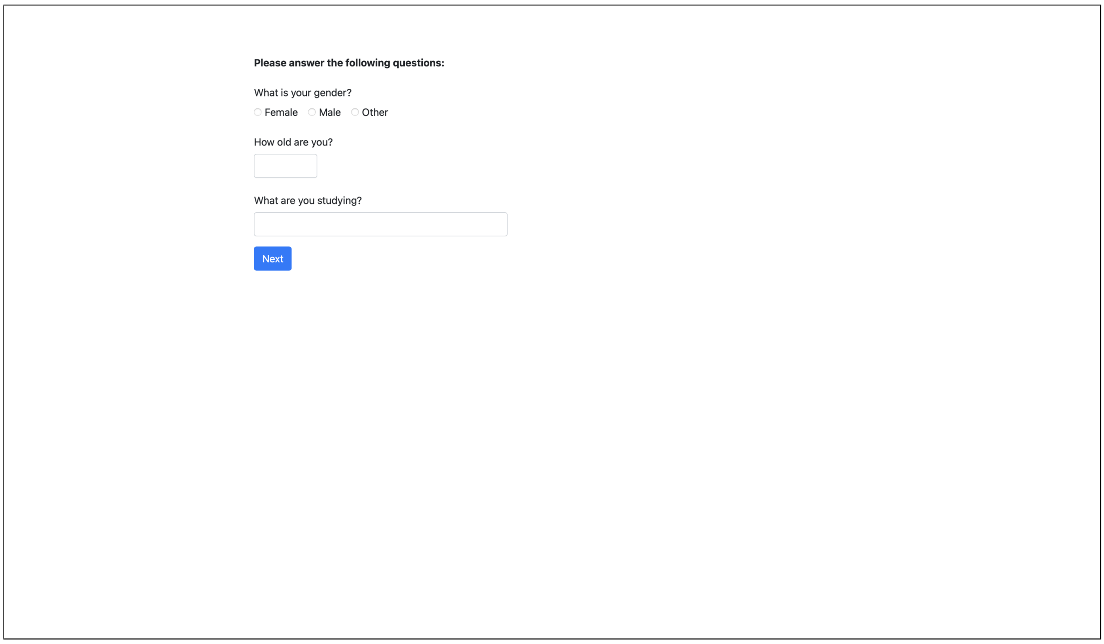
Please answer the following questions:
What is your gender?
Female
Male
Other
How old are you?
What are you studying?
Next

4.A.6 Final screen

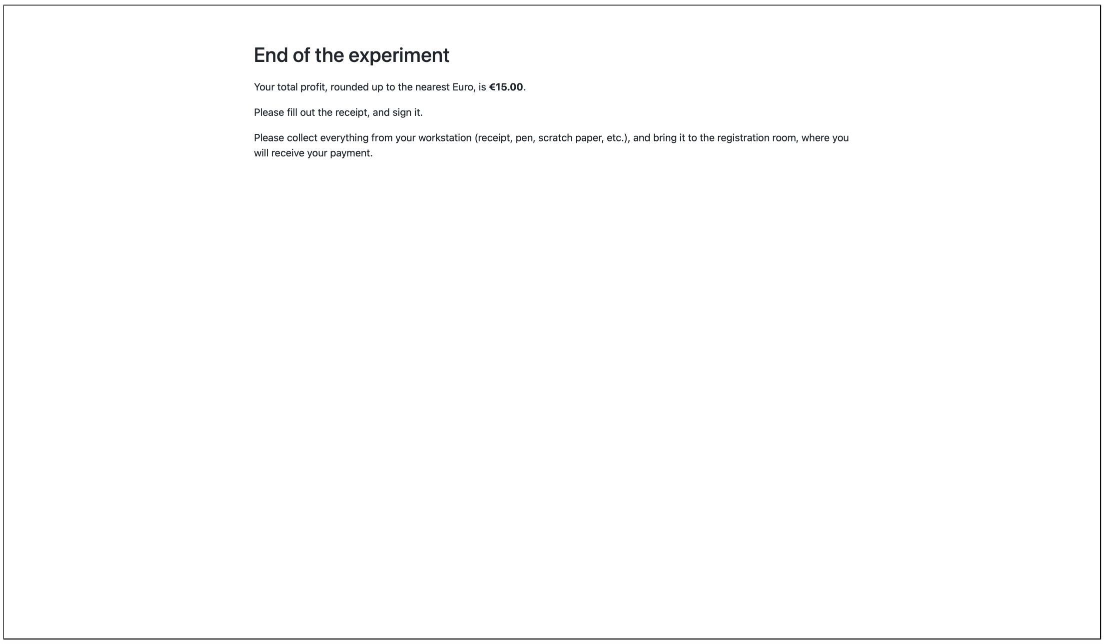

Chapter 5

Improving Incentive Schemes for Procurement Managers

There is anecdotal evidence that procurement managers who are rewarded for both low contract prices and continual price reductions may face a trade-off between receiving a bonus now and receiving a bonus later. We introduce a game-theoretic model to formally study the conflict of interest faced by a procurement manager with savings targets. We consider a dynamic game of incomplete information in which a firm delegates costly qualification and investment decisions to a procurement manager. Qualification increases competition, while investment refers to the effort exerted to identify and resolve potential inefficiencies with the contract supplier. The firm implements an incentive scheme that rewards low prices for the initial awarding as well as price reductions in the second period. We use this model to characterize the procurement manager's key trade-off: qualification increases competition, which leads to lower initial prices, but makes it harder for the procurement manager to secure a future bonus, since the value of investing decreases with the number of qualified bidders. Our main result is that even under an optimal incentive scheme, a procurement manager makes inefficient qualification and/or investment decisions, which directly translates into higher procurement expenses for the firm. We extend our analysis to show that the firm can minimize these inefficiencies by committing to an incentive scheme across both periods upfront.

5.1 Introduction and motivation

In 2021, procurement departments across industries managed an average of 78% of a firm's sourceable spend (CAPS, 2022). Given this substantial influence on financial resources, firms must ensure the efficient utilization of their funds by incentivizing procurement managers to act in alignment with the firm's long-term goals. Recently, incentive schemes based on annual savings targets have gained prominence (McKinsey & Company, 2014). In practice, this means that firms set annual savings targets for procurement departments and offer performance-based bonuses to managers who achieve these targets. Although performance-based bonuses are typically associated with explicit monetary payments, they may also be complemented or even substituted by more implicit forms of reward, such as (faster) promotions, preferential project assignments, or increased visibility to senior management. Depending on the design of the incentive scheme, target fulfillment can be measured, for example, as a reduction relative to the previous year's prices or relative to a cost calculation. To meet these targets, managers negotiate price reductions with their suppliers. This approach aims to foster a culture of cost discipline by encouraging consistent savings that benefit the firm over time.

In practice, measuring performance and target fulfillment is not straightforward. A central difficulty lies in the fact that observable outcomes, such as the negotiated price, only reveal full value creation relative to a counterfactual that is typically not observable. Without knowing what the price would have been without the manager's effort, performance cannot be reliably inferred from the observed outcome alone. For example, in a market with rising resource costs, being able to negotiate the same price for the coming period may be considered a success. In addition, by shifting the reference point for the savings calculation, managerial efforts undertaken early in a project to improve the competitive environment may make it harder to demonstrate savings in subsequent annual negotiations. To see this, consider two procurement managers starting in equally challenging market conditions with limited competition. The first procurement manager engages primarily in 'firefighting' during negotiations, stressing the difficulty of the situation while investing little to address its root causes. As a result of the state of emergency, this procurement manager may still receive substantial recognition for achieving modest savings. In contrast, the second procurement manager invests significant effort throughout the year to screen and qualify new suppliers. These efforts strengthen the firm's negotiation position and allow for lower prices in the long run. However, because the final negotiation appears comparatively easy, the contribution of this procurement manager is less visible and often underappreciated.

In general, challenges related to the measurement of value interact with the dynamic nature of savings generation. As a result, at first glance, continuous price reductions driven by savings targets seem beneficial. However, in such incentive schemes, procurement managers may face a trade-off: meeting savings targets in one period makes it more difficult to achieve savings targets in future periods. As such, the timing of savings becomes relevant to the procurement manager. Industry consultants underline this, noting that "many buyers are inherently incentivized not to harvest the full savings potential in year one [...] because that would mean they are not having much to offer in the following one(s)" (Scharlach, 2024). As pointed out by Johnson and Leenders (2010), this strategic timing of savings can be harmful to the firm. Ellram et al. (2020) provide a concrete example of this. Using case studies, they highlight that incentive schemes that reward ongoing price reductions can lead to inflated initial prices, with procurement managers continually negotiating price reductions to secure bonuses. At the expense of the firm, this strategic behavior can benefit both suppliers, who secure higher prices initially, and procurement managers, who meet their savings targets. Indeed, procurement professionals surveyed by Ellram et al. acknowledge that strategically delaying savings occurs in practice. The severity of the problem is difficult to overstate; increased procurement costs directly influence a firm's market competitiveness.

Thus far, to the best of our knowledge, the trade-offs a procurement manager faces as a result of such incentive schemes have only been studied anecdotally or using case studies, but not formally. This study aims to fill this gap. We formally characterize the trade-off a procurement manager faces under an incentive scheme that encourages cost savings and show that this trade-off results in higher expenses for the firm, even under optimal incentivization. Moreover, we use our model to demonstrate how a firm can use incentives to improve the overall outcome despite the procurement manager's trade-off.

To this end, we introduce a game-theoretic model to study the conflict of interest faced by a procurement manager with savings targets. We consider a dynamic two-stage game of incomplete information, in which the firm delegates costly qualification and investment decisions to a procurement manager. Qualification increases competition, whereas investment refers to the effort exerted to identify and resolve potential inefficiencies with the contract supplier. Both qualification and investment reduce the firm's total expenses in expectation. The firm implements an incentive scheme that rewards low prices in the first period, and price reductions in the second period. We use this model to characterize the procurement manager's key trade-off: qualification increases competition—which leads to lower initial prices—but makes it harder for the procurement manager to secure a future bonus, since the value of investing is decreasing in the number of qualified bidders. Our main result is that even under an optimal two-period incentive scheme, the total procurement expenses remain inefficiently high for the firm. Under such an incentive scheme, a procurement manager makes inefficient qualification and investment decisions, which directly translates into higher procurement expenses for the firm. In our initial analysis, we consider a firm that sets optimal incentives at the beginning of each period. We extend our analysis to include a firm that can credibly commit to an incentive scheme across both periods upfront. We show that this commitment is beneficial: it allows the firm to reduce the incentive necessary for qualification in the first period at the expense of less investment in the second period. In aggregate, the former effect outweighs the latter, resulting in lower total procurement costs. Nevertheless, even under an optimal two-period incentive scheme, the total procurement expenses remain inefficiently high for the firm.

The remainder of this paper is structured as follows. Section 5.2 connects our work with the related literature. The model is formally introduced in section 5.3 and analyzed in section 5.4. In sections 5.4.1 and 5.4.2, we solve the model with and without delegation, respectively. In section 5.5, we discuss the assumptions of our game-theoretic model, as well as various attempts to mitigate the trade-off between qualification and investment. Finally, section 5.6 discusses the managerial implications and concludes the paper.

5.2 Related Literature

Our research relates to the literature on procurement incentives and dynamic contracts, bridging the insights from business and economics research. We contribute by providing a theoretical framework that characterizes a key trade-off procurement managers face, namely, between creating competition on the one hand and the ability to negotiate price reductions on the other.

Our research builds on and extends the existing literature on dynamic incentive contracts, a field that is mainly concerned with the regulatory perspective. Much of the early work in this area focused on public goods provision (e.g., Laffont and Tirole (1986, 1987, 1988)), central planning (e.g., Freixas et al. (1985); Holstrom (1982)), and monopoly regulation (for example, Baron and Myerson (1982)). A typical model considers a two-period principal-agent setting with asymmetric information. The agent's efficiency type and effort level are unknown to the principal, who observes only the resulting costs after the first period. A challenge of commitment arises because the principal cannot credibly commit to a second-period incentive scheme at the outset of the interaction. Instead, after observing the first-period performance, the principal updates its beliefs about the agent's type and selects an optimal incentive structure for the second period (Laffont and Tirole, 1988). Such dynamic adjustments to incentive schemes can lead to a "ratchet effect", whereby an agent initially deliberately underperforms to make future incentives more attainable (Freixas et al., 1985; Laffont and Tirole, 1988). Various strategies have been proposed to address these problems. For instance, Baron and Besanko (1987) extend the model to allow for a principal who can guarantee the agent a non-negative profit if the agent forfeits the right to refuse participation in the second period. They characterize cases in which both the principal and agent are better off under such a scheme. Alternatively, Holstrom (1982) theoretically demonstrate that involving the agent when setting their targets and reward structure can reduce the commitment challenges faced by the principal.

We contribute to the literature by considering a principal-agent model commonly used in procurement contexts. In the setting we consider, the agent makes a qualification and an investment decision on the principal's behalf. In our model, we also find a "ratchet effect" that arises from price reduction targets: the agent is rewarded for achieving a certain price in the first period and a price reduction in the second period. However, this incentive scheme causes a conflict of interest for the agent: lower prices in the first period make it more difficult for the agent to meet the targets in subsequent periods. We demonstrate that this leads to inefficiently low supplier qualification in the first period and/or inefficiently low investment in the second period. As a result, the total procurement expenses for the firm are inefficiently high.

Our work also ties into the management literature. Conflicts of interest faced by agents have been studied, for example, in the context of physicians (Dai et al., 2022). Dai et al. empirically study the case of cardiologists to whom a test becomes available that informs their decision on whether or not a patient should undergo surgery. The authors demonstrated that overtreatment in cardiology cannot be avoided simply by the availability of better testing methodologies. Rather, because the testing decision is endogenous, a change in physicians' financial incentives is required. This is closely related to what we will demonstrate: while both qualification and investment are beneficial to the principal, the financial incentives of the agent can get in the way of their efficient use.

Looking specifically at literature concerned with the field of procurement, related work highlights that procurement managers may delay or forego potential savings owing to inadequate performance metrics, such as the inconsistent treatment of cost avoidance versus cost savings (Ellram and Tate, 2021), or as a result of goal misalignment between different actors within the firm (Ellram et al., 2020). For example, when the responsibility for initial procurement and recurring purchases is split between different departments, managers may lack the incentive to optimize total costs. We formalize the incentives a procurement manager faces and show that conflicts of interest can arise even with adequate performance tracking and when only considering a single principal and a single agent. Specifically, our model shows that having a single procurement manager responsible for both initial and recurring purchases does not necessarily ensure goal alignment. Rather, the trade-off between competitive prices upfront and the ability to negotiate price reductions results in inefficiently high total procurement costs.

Moving from private to public procurement, Coviello et al. (2018) leverage data on public procurement tenders in Italy to provide empirical evidence that procurement managers limit upfront competition. For tenders below a certain threshold, Italian legislation gives buyers more discretion on whom (not) to invite to bid. By studying tenders in the neighborhood of this threshold, they identify that this buyer discretion directly leads to an increased likelihood that the same firm will repeatedly win tenders.

The studies listed above have in common that they rely on empirical case studies. In contrast, we introduce a theoretical model in this study. Using this model, we show that even under optimal incentivization, a firm's delegation of qualification and investment decisions to a procurement manager can lead to inefficiencies. In the following section, we formally introduce our model.

5.3 Model

In this section, we begin by formally introducing the model. Next, we solve the model assuming that the firm can make all decisions. This serves as a benchmark for a setting in which the firm must delegate its decisions to a procurement manager.

5.3.1 General Framework

We model a two-period setting in which a firm awards a contract to a single supplier. The firm initially has access to a supplier base consisting of two prequalified suppliers. In each period, the firm faces a binary decision: (i) in the first period ($t = 1$), it chooses whether to qualify an additional supplier to expand the supplier base; and (ii) in the second period ($t = 2$), it decides whether to invest in improving the efficiency of the selected contract supplier. Such an investment aims to reduce the price paid for the good by addressing potential inefficiencies in the supplier's operations. These inefficiencies may stem from outdated production processes, coordination and scheduling problems, quality-related rework, or equipment constraints. An example of such investments comes from the automotive industry: Toyota continually invests to increase the efficiency of its suppliers' production processes and thus reduce procurement costs (Toyota, 2024).

Importantly, the firm does not observe suppliers' efficiency levels ex ante; thus, investment outcomes are uncertain—the supplier may already be operating efficiently, in which case the investment yields no benefit. In both periods, the firm may alternatively opt to take no action. The firm's objective is to minimize the total procurement cost.

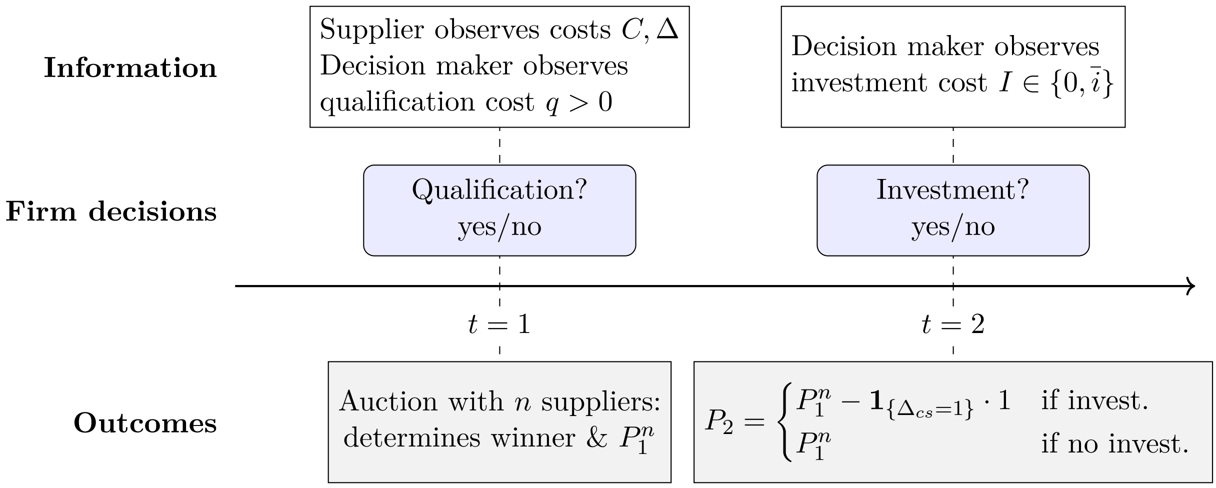

Figure 5.1: Timeline of information, decisions, and outcomes

At the beginning of $t = 1$, the firm selects a supplier through a second-price auction. We denote by n the number of suppliers that participate in the auction. The winning supplier becomes the contract supplier (indexed as “cs” in the analysis), and the auction determines the price for the good in the first period, denoted P_1^n. Without further action by the firm, the price remains unchanged in the second period, that is, $P_2^n = P_1^n$.

The total cost of a supplier consists of two independent components: (i) production costs $C \in \{0, 1\}$, each with equal probability, and (ii) inefficiency costs $\Delta \in \{0, 1\}$, again with equal probability. A supplier is considered inefficient if its inefficiency cost is 1. Supplier costs are private information, and we abstract from potential quality differences.[1] Observe that in the setting we consider, it is a weakly dominant strategy for suppliers to bid their total costs in the auction, see for example the proof in Krishna (2010).

In $t = 1$, the firm can qualify an additional supplier, which increases competition and—on average—lowers procurement costs. Additionally, it lowers the probability of contracting an inefficient supplier. The qualification cost is $q > 0$.

In $t = 2$, the firm can invest in potentially improving the contracted supplier's efficiency. The investment cost is denoted by $I \in \{0, \bar{i}\}$, with probabilities $\frac{1}{4}$ and $\frac{3}{4}$, respectively, and with $\bar{i} > 0$. The realization of I occurs at the beginning of $t = 2$, after the supplier is selected. In our model, investment cost uncertainty reflects the interaction between the firm and the contracted supplier; it captures how difficult it is for the firm to assess and resolve potential inefficiencies in the supplier's operations. The value of I captures whether the interaction is feasible at low or high cost. Alternatively, I can also be read as a good or poor fit between the firm and supplier; in the following, we will use this interpretation. If the firm invests and the contract supplier is inefficient ($\Delta = 1$), the inefficiency is resolved, reducing the second-period price to $P_2 = P_1 - \Delta$. If the contract supplier

[1] Throughout the analysis, prices are considered to be quality-adjusted. Without loss of generality, any quality disparities can be incorporated through a bonus-penalty system without affecting the general structure of the model.

is not inefficient, the price remains unchanged, but the investment cost applies nonetheless.

5.4 Analysis

In this section, we start by deriving the first-best solution, which serves as a benchmark for the other settings.

5.4.1 First-best

In the first-best setting, the firm makes all decisions, bears all resulting costs, and reaps the resulting benefits. We solve the model using backward induction.

Period t = 2

Table 5.1 shows, for both $n = 2$ and $n = 3$, the respective probability of a given first-period price occurring, as well as the probability of the contract supplier being inefficient for this price. For example, consider the case where $n = 2$, that is, the case in which no qualification takes place. With a probability of $\frac{8}{16}$, the auction price is 1. For this price, the probability of the contract supplier being inefficient is $\frac{1}{4}$. If $n = 3$, that is, if qualification takes place, the auction price is 1 with a probability of $\frac{22}{32}$, and the probability of the contract supplier being inefficient is $\frac{5}{22}$.

	$n=2$		$n=3$	
P_1	$\mathbb{P}[P_1]$	$\mathbb{P}[\Delta=1\|P_1]$	$\mathbb{P}[P_1]$	$\mathbb{P}[\Delta_{cs}=1\|P_1]$
0	$\frac{1}{16}$	0	$\frac{5}{32}$	0
1	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{22}{32}$	$\frac{5}{22}$
2	$\frac{7}{16}$	$\frac{3}{7}$	$\frac{5}{32}$	$\frac{2}{5}$
$\mathbb{E}[P_1]$	$\frac{11}{8}$	—	1	—

Table 5.1: First-period prices and inefficiency probabilities for $n=2$ and $n=3$ suppliers.

Investment is beneficial for a firm whenever the expected savings exceed the investment costs. That is, whenever

$$\mathbb{P}[\Delta_{cs}=1|p]\cdot\Delta\geq I. \tag{5.1}$$

For a firm with a good firm-supplier fit ($I=0$), this always holds, and the firm always invests. For a firm with a poor firm-supplier fit ($I=\bar{i}$), the investment decision depends on the realization of p and $\bar{i}$.

Due to the model design, three price realizations are possible: $0,1,2$. Denote by $\hat{P}^n$ the cut-off price from which, if realized, a firm with high investment costs and n qualified suppliers invests. The cut-off price $\hat{P}^n$ is endogenous to the realization of $\bar{i}$: higher investment costs increase the cut-off price (less price realizations for which investments take place) and lower investment costs lower the cut-off price (see Figure 5.2).

Figure 5.2 shows the first-period prices at which a firm with a poor firm-supplier fit invests, depending on the number of qualified suppliers and the investment costs. For example, for $\bar{i}\in(\frac{1}{4},\frac{2}{5})$, the high-cost firm invests only when observing $P_1=2$ with both $n=2$ and $n=3$: $\hat{P}^{n=2}=\hat{P}^{n=3}=2$.

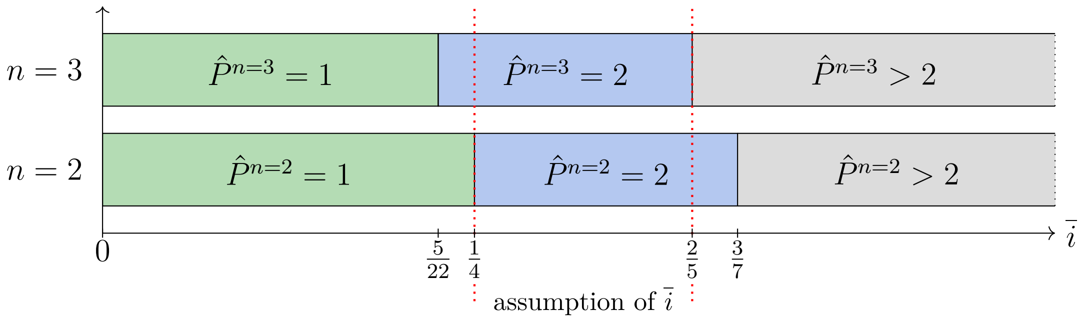

Figure 5.2: Investment decision of high-cost firm depending on i and n

Period t = 1

In the first period, the firm must weigh the costs and benefits of qualifying an additional supplier. An additional supplier promotes competition in the auction, lowers expected procurement costs, and reduces the probability of selecting an inefficient supplier. However, this comes at a fixed qualification cost $q > 0$, which must be justified by the potential downstream benefits. The firm's total cost is given by

$$\begin{aligned} TC^n = & \sum_p \mathbb{P}^n[p] \cdot p + \sum_{p<\hat{P}^n} \mathbb{P}^n[p] \cdot \left(p - \frac{1}{4} \cdot \mathbb{P}^n[\Delta_{cs} = 1 \mid p]\right) \\ & + \sum_{p \geq \hat{P}^n} \mathbb{P}^n[p] \cdot \left(p - \mathbb{P}_n[\Delta_{cs} = 1 \mid p] + \frac{3}{4} \cdot \bar{i}\right) + \mathbf{1}^{[n=3]} \cdot q. \end{aligned} \tag{5.2}$$

The first term in eq. (5.2) represents the expected price in $t = 1$. In $t = 2$, for $p < \hat{P}^n$, only firms with a good firm-supplier fit invest. The expected price reduction in this case is the product of the probability that the firm has a good firm-supplier fit and the probability that the contract supplier is inefficient. For $p \geq \hat{P}^n$, all firms invest; however, firms with a poor firm-supplier fit incur $\bar{i}$.

If the firm decides to qualify the additional supplier, the computation of its total cost is analogous, except that it also bears the qualification cost in $t = 1$ (see the indicator function in eq. (5.2)).

It is beneficial for a firm to qualify whenever $TC^{n=3} \leq TC^{n=2}$. Table 5.2 lists the qualification thresholds for each interval of $\bar{i}$. If $q \leq q(\bar{i})$ in an interval, a firm qualifies an additional supplier.

	$TC^{n=2}$	$TC^{n=3}$	$q(\bar{i})$
$0 < \bar{i} \leq \frac{5}{22}$	$\frac{39}{16} + \frac{45}{64} \cdot \bar{i}$	$\frac{57}{32} + \frac{81}{128} \cdot \bar{i} + q$	$\frac{21}{32} + \frac{9}{128} \cdot \bar{i}$
$\frac{5}{22} < \bar{i} \leq \frac{1}{4}$	$\frac{39}{16} + \frac{45}{64} \cdot \bar{i}$	$\frac{243}{128} + \frac{15}{128} \cdot \bar{i} + q$	$\frac{69}{128} + \frac{79}{128} \bar{i}$
$\frac{1}{4} < \bar{i} \leq \frac{2}{5}$	$\frac{81}{32} + \frac{21}{64} \cdot \bar{i}$	$\frac{243}{128} + \frac{15}{128} \cdot \bar{i} + q$	$\frac{81}{128} + \frac{27}{128} \bar{i}$
$\frac{2}{5} < \bar{i} \leq \frac{3}{7}$	$\frac{81}{32} + \frac{21}{64} \cdot \bar{i}$	$\frac{249}{128} + q$	$\frac{75}{128} + \frac{21}{64} \bar{i}$
$\frac{3}{7} < \bar{i}$	$\frac{171}{64}$	$\frac{249}{128} + q$	$\frac{93}{128}$

Table 5.2: First-best qualification thresholds and total cost depending on $\bar{i}$.

Figure 5.3 summarizes the firm's qualification and investment decisions across both periods. The qualification threshold curve shows that qualification becomes more attractive when investment costs increase; that is, the firm is willing to incur higher qualification costs when investing in the second period becomes more expensive. The colored regions indicate the first-period cut-off prices above which the firm chooses to invest. As investment costs increase, it becomes less attractive for the firm to invest, and the cut-off price increases accordingly. For certain realizations of $\bar{i}$, the cut-off price also depends on the qualification decision: when no additional supplier is qualified, the firm is willing to invest for a wider range of first-period prices, as the potential gains from investment become more valuable as a result of the higher probability of the contract supplier being inefficient. Outside these ranges, the investment regions are identical for $n = 2$ and $n = 3$.

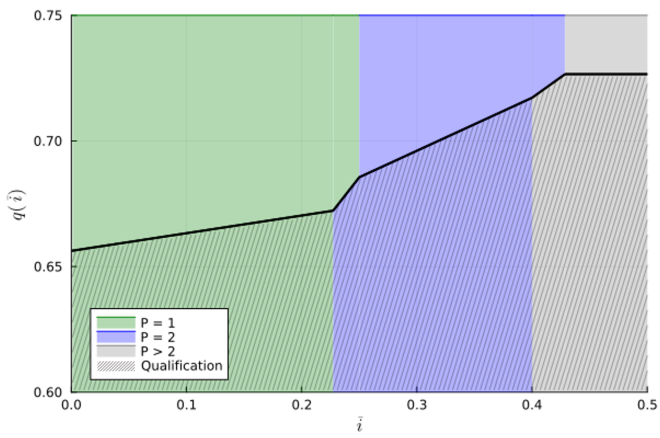

Figure 5.3: Qualification and investment decision dependent on q and $\bar{i}$

As shown in the figure above, multiple cases must be considered. For example, for very low $\bar{i}$, the firm will invest for all $P_1 > 0$, while for very high $\bar{i}$, the firm never invests. We will consider the intermediate case of $\bar{i} \in (\frac{1}{4}, \frac{2}{5})$ (see the red interval in Figure 5.2). In this case, a firm with a poor firm-supplier fit invests only when observing $P_1 = 2$ with both $n = 2$ and $n = 3$: $\hat{P}^{n=2} = \hat{P}^{n=3} = 2$, whereas a firm with a good firm-supplier fit always invests. We consider this case because, in extreme cases, studying the effects of delegation to a procurement manager is less interesting, since the decisions do not change. For example, for large $\bar{i}$, neither the firm nor the procurement manager would invest.

The corresponding qualification threshold (for both types of firms) for the case we consider is given by:

$$q_{FB} \leq \frac{81}{128} + \frac{27}{128}\bar{i}. \tag{5.3}$$

In this section, we determine the optimal qualification and investment levels. In the following section, we analyze a situation in which the firm delegates qualification and investment decisions to a procurement manager.

5.4.2 Delegation to a procurement manager

In large firms, the owner(s) cannot make all operational decisions themselves and, therefore, delegate the responsibility to functional managers within the firm. Incentive schemes are implemented to ensure that delegated decisions serve the firm's objectives. However, incentive schemes can lead managers to take actions that increase their personal gain at the expense of the firm's cost efficiency. We extend our model by assuming that procurement decisions are delegated to a procurement manager. The manager seeks to maximize personal profit, which may—and, as we will see below, does—conflict with the firm's cost-minimization objective.

What happens if the firm must delegate the qualification and investment decisions to an agent? As we show below, there exists a trade-off between qualification and investment: qualifying the additional supplier increases the agent's first-period payoff but decreases the likelihood of the contract supplier being inefficient and thus of receiving a bonus in the second period. Consequently, the first-best solution cannot be obtained through delegation. With delegation, a procurement manager makes inefficiently low qualification and/or investment decisions, which directly translates into higher procurement expenses for the firm.

Under delegation, the procurement manager bears the qualification cost q and the investment cost I, whereas the firm benefits from any resulting price reduction. Without additional incentives, the procurement manager will not qualify an additional supplier and, when incurring a poor firm-supplier fit, will not invest in resolving a potential inefficiency. Therefore, the firm implements a bonus scheme. This scheme encompasses both periods and is communicated to the procurement manager at the beginning of $t = 1$ before any decisions are made.

Consider $t = 1$. If the firm wants the manager to qualify the additional supplier, it will set the first-period bonus such that the manager is exactly indifferent between qualifying and not qualifying. In other words, the expected first-period bonus is the lowest amount the firm can pay the manager and still have them qualify. We show that it suffices to consider optimal bonus schemes that reward the manager only for achieving the lowest possible price in $t = 1$. The following lemma formalizes this observation.

Lemma 3. *An optimal incentive scheme can be implemented in* $t = 1$ *by offering the procurement manager a bonus* $B_1 \geq 0$ *only if the best possible price,* $P_1 = 0$ *is attained.*

Proof. First, observe that $P_1 \in \{0, 1, 2\}$. Next, note that we do not have to consider a case in which the manager receives a bonus $B_1', B_1'', B_1''' > 0$ if $P_1 = 0$, $P_1 = 1$, and $P_1 = 2$ respectively, since such a bonus scheme rewards the manager for attaining the worst possible price, which cannot be optimal.

We show that any incentive scheme that rewards multiple price realizations and under which the manager qualifies can also be represented by an incentive scheme in which the procurement manager receives a bonus only if $P_1 = 0$.

Suppose there is a bonus scheme under which the manager qualifies that pays a bonus $B_1', B_1'' > 0$ when attaining prices $P_1 = 0$ and $P_1 = 1$, respectively. The expected payment to the manager is given by

$$\mathbb{P}[P = 0] \cdot B_1' + \mathbb{P}[P = 1] \cdot B_1''.$$

Then, there is always an incentive scheme under which the manager qualifies that only pays a bonus for attaining $P_1 = 0$ and yields the same or lower total cost to the firm. To see this, note that the expected payment to the manager under our proposed incentive scheme is given by

$$\mathbb{P}[P = 0] \cdot B_1.$$

Observe that setting

$$B_1 := B_1' + \frac{\mathbb{P}[P = 1]}{\mathbb{P}[P = 0]} \cdot B_1''$$

yields the same expected payment to the manager. This holds whenever $\mathbb{P}[P = 0] > 0$.

Now, observe that in any optimal $t = 1$ incentive scheme, the manager must be exactly indifferent between qualifying and not qualifying. In other words, it is optimal for the firm to pay the manager the lowest possible bonus B_1^* such that they still qualify. Clearly, $B_1^* \leq B_1$. This concludes the proof.

□

In $t = 2$, the bonus scheme entails a bonus $B_2 > 0$ for the procurement manager if a price reduction is achieved, that is, if the contract supplier is inefficient and the inefficiency is resolved. No bonus is awarded if an investment is undertaken, but no price reduction occurs.

We consider two cases: *no-commitment* and *commitment*.

In the no-commitment case, the firm sets bonuses sequentially: it chooses B_1 at the beginning of $t = 1$, and B_2 after observing the realized outcomes of $t = 1$, at the beginning of $t = 2$. Consequently, the firm selects B_2 to minimize its second-period cost, given the observed first-period outcome.

In the commitment case, the firm decides on both bonuses ex ante: it sets B_1 and B_2 simultaneously at the beginning of $t = 1$, and commits to not revise them after observing the outcomes of $t = 1$. Consequently, it sets both bonuses to minimize expected total cost across both periods, based on information available at $t = 1$. The commitment bonus scheme can be seen as more complex for the firm because the bonuses are functions of the not-yet-realized first-period price.

We will show two key results. First, the ranking of the three settings from most to least beneficial for the firm in terms of total cost is: first-best, commitment, no-commitment. Second, delegation to an agent leads to inefficiently low qualification and investment in both settings. At this point, we want to stress that these results are not a relic of the specific intermediate case we are considering. Rather, they hold generally for all cut-off prices smaller than 2 (so all situations where prices can realize that lead to investment). The proofs have been relegated to section 5.A (see Propositions 20 and 21).

We will consider both settings in turn, starting with the no-commitment setting.

5.4.3 No-commitment

Period t = 2:

Recall that the firm does not know whether it is in a situation with a good or poor firm-supplier fit. In general, procurement managers invest if their expected bonus is at least as high as their investment costs. This is always fulfilled when the firm-supplier fit is good. The expected payoff to the manager from investing in a situation where the firm-supplier fit is poor is given by

$$\mathbb{P}(\Delta_{cs} = 1 | P_1) \cdot B_2 - \bar{i}. \tag{5.4}$$

Rearranging yields the optimal bonus in $t = 2$

$$B_2^* = \frac{\bar{i}}{\mathbb{P}(\Delta_{cs} = 1 | P_1)}. \tag{5.5}$$

This is the optimal bonus because it minimizes the firm's expected cost in $t = 2$ while ensuring that managers choose to invest when the firm-supplier fit is poor. Specifically, the firm sets B_2 such that in a situation with a poor firm-supplier fit, managers are indifferent between investing and not investing.

Note that managers with a good firm-supplier fit make a surplus under this bonus scheme. This surplus is an information rent and derives from the fact that the firm offers B_2^*—which incorporates the high investment costs of $\bar{i}$—but investment costs are 0 when the firm-supplier fit is good.

Proposition 14. *Relative to first-best, investment is inefficiently low in the no-commitment setting.*

Proof. Consider a price $p \geq \hat{P}^n$. In the first-best case, the firm will always invest for this p. In the no-commitment case, whether the firm wants to incentivize investment depends on $\bar{i}$.

If the firm does not incentivize investment, their cost in $t = 2$ is given by

$$p - \mathbb{P}(\Delta_{cs} = 1|p) \cdot \frac{1}{4}.$$

In this case, investment will only take place when the firm-supplier fit is good, which is the case in a quarter of the time. The price reduction only occurs when investment takes place and the contract supplier is inefficient.

On the other hand, if the firm decides to incentivize the manager to invest, firm cost in $t = 2$ is given by

$$p - \mathbb{P}(\Delta_{cs} = 1|p) \cdot (1 - B_2^*).$$

All managers invest independently of the firm-supplier fit, and the price is reduced whenever the contract supplier is inefficient. However, the firm must pay B_2^* to achieve this.

We find that the firm does not want the manager to invest whenever

$$\bar{i} \geq \bar{i}_{crit} := \frac{3}{4} \cdot \mathbb{P}(\Delta_{cs} = 1|p). \tag{5.6}$$

We have shown that for certain investment costs, the firm will not incentivize investment, even though this is optimal in first-best. This concludes the proof.

□

For the analysis that follows, denote by $\hat{P}_{NC} \geq \hat{P}$ the cut-off price above which the firm wants to incentivize the manager to invest in the non-commitment setting.

We show that investment is inefficiently low in the no-commitment setting. Next, we consider the manager's qualification decision.

Period t = 1:

In the first period, the firm must decide whether to incentivize the manager to qualify the additional supplier.

Considering only optimal bonus schemes, the firm has two options: (i) pay no bonus in $t = 1$, which means the manager will not qualify the additional supplier, or (ii) pay a B_1 such that, in expectation, the manager makes at least as much profit when they qualify as when they do not qualify. Comparing the firm's total costs reveals when the firm wants to incentivize the manager to qualify.

Note that the manager faces a trade-off. Qualifying *increases* the chances of receiving B_1, but—because more suppliers in the auction lead to lower probabilities of the contract supplier being inefficient— *decreases* the chances of receiving B_2.

The firm must account for this when setting B_1.

When will the manager qualify? Whenever their total profit with qualification is weakly greater than their total profit when they do not qualify.

The manager's total profit is given by

$$\mathbb{P}[P_1^n = 0] \cdot B_1 + \frac{1}{4} \sum_{p \geq \hat{P}_{NC}} \mathbb{P}^n[p] \cdot \mathbb{P}^n[\Delta_{cs} = 1|p] \cdot B_2^* - \begin{cases} 0, \text{if } n = 2 \\ q, \text{if } n = 3 \end{cases}$$
$$= \mathbb{P}[P_1^n = 0] \cdot B_1 + \frac{1}{4} \sum_{p \geq \hat{P}_{NC}} \mathbb{P}^n[p] \cdot \bar{i} - \begin{cases} 0, \text{if } n = 2 \\ q, \text{if } n = 3 \end{cases} \tag{5.7}$$

Whether the firm is willing to incentivize the manager depends on the investment cost $\bar{i}$ and qualification cost q. Based on Proposition 14, three relevant intervals can be distinguished:

$\bar{i}$	$t = 2$	$t = 1$
$[\frac{1}{4}, \frac{3}{10})$	$\hat{P}_{NC} = \hat{P} = 2$	$q \leq q_{NC}^{(1)} := \frac{3}{5} \cdot (\frac{81}{128} + \frac{21}{128}\bar{i})$
$[\frac{3}{10}, \frac{9}{28})$	$\hat{P}_{NC}^{n=2} = \hat{P} = 2$ and $\hat{P}_{NC}^{n=3} > 2$	$q \leq q_{NC}^{(2)} := \frac{3}{5} \cdot (\frac{75}{128} + \frac{49}{192}\bar{i})$
$[\frac{9}{28}, \frac{2}{5}]$	$\hat{P}_{NC} = \hat{P} > 2$	$q \leq q_{NC}^{(3)} := \frac{3}{5} \cdot \frac{93}{128}$

Table 5.3: Qualification incentives for different $\bar{i}$ and q

Table 5.3 summarizes the firm's investment and qualification incentives for the manager as a function of $\bar{i}$ and q. Column $t = 2$ shows the cut-off price above which the firm would invest in the first-best setting, as well as the cut-off price above which the firm chooses to incentivize investment in the no-commitment setting. Column $t = 1$ shows the maximum qualification cost q for which the firm finds it optimal to incentivize the manager to qualify an additional supplier in the no-commitment setting. If q exceeds this threshold, the cost of providing incentives becomes too high, and the firm refrains from incentivizing qualification, resulting in no qualification. These results highlight that under no-commitment, qualification occurs less frequently than in the first-best scenario. The derivation of the maximum qualification cost per interval is relegated to section 5.A.

Proposition 15. *Qualification is inefficiently low in the no-commitment setting.*

Proof. Denote by q_{FB} the threshold below which the additional supplier is qualified in the first-best setting. Denote by $q_{NC}^{(1)}, q_{NC}^{(2)}, q_{NC}^{(3)}$ the thresholds below which the additional supplier is qualified in each of the three intervals in the no-commitment setting.

Comparing the different thresholds, we find that $q_{NC}^{(1)} < q_{NC}^{(2)} < q_{NC}^{(3)} < q_{FB}$.

This means that for a given interval k in the no-commitment setting, $\forall q \in (q_{NC}^{(k)}, q_{FB})$, no qualification takes place in the no-commitment setting even though it would in first-best.

This establishes that the qualification cost threshold is inefficiently low in the no-commitment setting. As a result, there are parameter ranges in which qualification would be efficient, but no qualification takes place as it is not incentivized by the firm. This completes the proof. □

We have shown that the no-commitment setting leads to inefficiently low investment and qualification. These inefficiencies translate directly into higher total costs for the firm.

Proposition 16. *The total costs to the firm are higher in the no-commitment setting than in the first-best setting.*

Proof. Depending on the values of $\bar{i}$ and q, there are multiple avenues through which firm costs can be higher in the no-commitment setting.

By definition, inefficient investment and qualification translate directly into higher costs for the firm. However, even in the case of efficient investment and qualification, firm costs are higher than in first-best. In the former case, this is a result of the firm not knowing which type of manager they are facing. This information asymmetry means that low-cost managers make a positive profit in $t = 2$. In the latter case, the firm needs to pay $B_1^* > q$ to incentivize qualification because the procurement manager internalizes the impact of qualification on future bonus prospects.

In all possible cases, costs to the firm are higher than in first-best, which concludes the proof. □

We show that delegation in the no-commitment setting leads to inefficiently low investment and qualification, which translates into higher total costs for the firm. In the following section, we show that commitment is beneficial to the firm.

5.4.4 Commitment

In this section, we examine the case in which the firm simultaneously sets both bonuses at the beginning of $t = 1$ and credibly commits to not revising them thereafter. We show that such a commitment is beneficial to the firm. In the no-commitment setting, the firm considers each period separately when setting incentives. Credible commitment allows the firm to jointly consider both periods when setting the incentive scheme. As we will show below, under joint optimization, there are cases in which it is worth it for the firm to reduce the second-period bonus B_2 (to $B_2 = 0$).[2]

While this weakens investment incentives, it also lowers the bonus B_1 necessary to incentivize qualification. In aggregate, the cost savings from the lower B_1 outweigh the efficiency loss from reduced investment. Through decreased investment and increased qualification, the commitment setting leads to a lower total cost for the firm compared to the no-commitment setting. However, overall, the firm's total cost in the commitment setting remains above the first-best benchmark. To understand the trade-off between qualification and investment, we revisit the three relevant investment cost intervals defined earlier and explore how the firm adjusts its strategy under commitment compared to the no-commitment setting. The detailed derivation is relegated to section 5.A.

- **Interval 1: $\bar{i} \in [\frac{1}{4}, \frac{3}{10}]$:** In the no-commitment setting discussed above, we saw that in $t = 2$, for all n, the firm incentivizes the manager to invest for $P_1 = 2$ and incentivizes to qualify when $q \leq q_{NC}^{(1)}$. In the commitment setting, firms deviate from this strategy. It is optimal for the firm not to incentivize investment at all and to incentivize qualification when $q_{NC}^{(1)} < q \leq q_C$, where:

$$q_C := \frac{3}{5} \cdot (\frac{75}{128} + \frac{7}{16}\bar{i}). \tag{5.8}$$

  In summary, for Interval 1, we find that through decreased investment and increased qualification, the commitment setting leads to lower total cost to the firm compared to the no-commitment setting.

- **Interval 2: $\bar{i} \in [\frac{3}{10}, \frac{9}{28})$:** In the no-commitment setting, the firm incentivizes investment for $P_1 = 2$ only when no qualification took place ($n = 2$) and incentivizes to qualify when $q \leq q_{NC}^{(2)}$. In the commitment setting, the firm

[2] First, note that offering $B_2 > B_2^*$ is not beneficial to the firm, as all managers already invest for B_2^* independent of the firm-supplier fit. Second, offering $0 < B_2 < B_2^*$ is not beneficial either, since under such a bonus, investment only takes place when there is a good firm-supplier fit—but in these cases investment would also take place for $B_2 = 0$.

deviates from this strategy for a certain range of qualification costs q:

  - *For* $q \leq q_{NC}^{(2)}$, the firm sets the same incentives as under no-commitment.
  - *For* $q_{NC}^{(2)} < q < q_C$, the firm incentivizes qualification, but does not incentivize investment.

  In summary, for Interval 2, we find that through decreased investment and increased qualification, the commitment setting leads to lower total cost for the firm compared to the no-commitment setting.

- **Interval 3:** $\bar{i} \geq \frac{9}{28}$**:** In the no-commitment setting, the firm never incentivizes the manager to invest, and only incentivizes the manager to qualify when $q \leq q_{NC}^{(3)}$. In this interval, the firm sets the same incentives in the commitment setting as it does in the no-commitment setting. In summary, for Interval 3, the firm's incentives for the manager are identical in both the no-commitment and commitment settings.

The above analysis leads to three key results.

Proposition 17. *Investment is lower in the commitment setting than in the no-commitment and first-best settings.*

Proof. Denote by $\hat{P}_{FB}$, $\hat{P}_{NC}$, and $\hat{P}_C$ the cut-off prices above which a firm (or manager) with a poor firm-supplier fit invests in the first-best, no-commitment, and commitment settings, respectively.

First, in Interval 3, the investment strategies do not differ between the commitment and no-commitment settings. That is,

$$\hat{P}_C = \hat{P}_{NC}.$$

Next, observe that in the no-commitment setting, in Intervals 1 and 2, the manager always invests when the firm would in first-best. That is,

$$\hat{P}_{NC} = \hat{P}_{FB}.$$

However, we have shown that there exist values of $\bar{i}$ and q for which the total cost to the firm is lower in the commitment setting than in the no-commitment setting. Put differently, in these cases, it is beneficial for the firm to deviate from the first-best investment strategy by setting

$$\hat{P}_C > \hat{P}_{FB}.$$

Therefore, in Intervals 1 and 2, only managers in situations with a good firm-supplier fit invest in the commitment setting, while all managers, independent of the firm-supplier fit, invest in the no-commitment and first-best settings. This concludes the proof. □

The intuition behind this result is that by committing to an ex-post sub-optimal bonus in $t = 2$, the firm is able to to reduce its total cost by incentivizing qualification at the expense of investment. This means that we have more qualification in the commitment setting compared with no-commitment, which gives us another proposition.

Proposition 18. *Qualification is inefficiently low in the commitment setting compared to first-best. However, qualification is higher than in the no-commitment setting.*

Proof. Denote by q_{FB}, q_{NC} and q_C the thresholds below which the additional supplier is qualified in the first-best, no-commitment, and commitment settings, respectively.

Interval 1 and 2: We find that $q_{NC} < q_C < q_{FB}$ (see eqs. (5.3), (5.9) and (5.30) in sections 5.A, 5.4.1 and 5.4.2).

This means that $\forall q \in (q_{NC}, q_C)$, no qualification takes place in the no-commitment setting, even though it would in commitment. Similarly, $\forall q \in (q_C, q_{FB})$, no qualification takes place in the commitment setting, even though it would in first-best.

Interval 3: We find that $q_{NC} = q_C < q_{FB}$ (see eqs. (5.3) and (5.29) in sections 5.A and 5.4.1).

This means that $\forall q \in (q_{NC}, q_{FB})$, no qualification takes place in the no-commitment and commitment settings even though it would in first-best.

This concludes the proof. □

One way to think about the commitment setting is that it places fewer restrictions on the firm's second-period bonus scheme. Under no-commitment, the firm is forced to implement an optimal second-period incentive scheme. On the other hand, under commitment, by jointly optimizing the first- and second-period incentive schemes, the firm can credibly commit even to an ex-post sub-optimal second period incentive scheme. As described above, by incentivizing qualification instead of investment, the firm can leverage this additional freedom and attain lower total costs. This leads to our final proposition.

Proposition 19. *The total cost to the firm is lower in the commitment setting than in the no-commitment setting but higher than in the first-best setting.*

Proof. This follows immediately from the two previous propositions. Investment and qualification are inefficient in the commitment setting. Thus, total costs are higher than in the first-best setting. However, we have demonstrated cases in which total costs are lower for the firm in the commitment setting than in the no-commitment setting. This concludes the proof. □

This follows directly from the previous two propositions. Although commitment leads to inefficient qualification and investment, the total cost is lower in the commitment setting than in the no-commitment setting.

The commitment setting shows how increased flexibility in bonus choice can be strategically valuable for the firm. By committing to a suboptimal bonus scheme in the second period, the firm credibly shifts incentives toward qualification and away from costly investment. This is not efficient, but it is cost-minimizing relative to the no-commitment setting. In summary, commitment allows the firm to restructure incentives in a way that outperforms a simple period-by-period optimization.

5.5 Discussion

In this section, we discuss the assumptions of our game-theoretic model and examine the efficacy of various attempts to mitigate the trade-off between qualification and investment.

5.5.1 Investment costs of the procurement manager

We start with the assumption of the investment costs of the procurement manager. In our model, high (low) investment costs (i.e., good (poor) firm-supplier fit) occur with a probability of $\frac{3}{4}$ ($\frac{1}{4}$). These values were chosen for illustrative purposes; our model generally holds for all probabilities of low investment costs below a critical threshold. To see this, first observe that the higher the probability of low investment costs, the less attractive it is for the firm to offer a bonus for investment in the second period. Specifically, rearranging eq. (5.6) from proposition 14, we find that firm will not incentivize investment in $t = 2$ at all whenever the probability of low investment costs exceeds

$$1 - \frac{\bar{i}}{\mathbb{P}(\Delta_{cs} = 1|p)}.$$

For probabilities of low investment costs above this threshold, there is no longer a trade-off for the procurement manager: both the low and high investment cost managers make zero profit in $t = 2$. As such, only the first-period bonus is relevant, resulting in optimal qualification on the manager's part. Looking at this through a practitioner's lens, this means that funding programs to reduce the share of procurement managers with high investment costs (or improving the firm's capacity to assess and resolve potential inefficiencies in the supplier's operations) for example, through training, may well be worth it for the firm.

5.5.2 Level of inefficiency

Next, let us consider the assumption on the level of inefficiency Δ. For ease of exposition, we assume $\Delta = 1$. Doing so reduces the number of possible first-period prices from four to three, which considerably reduces the number of cases we need to consider when delegating qualification and investment decisions to the procurement manager.[3] However, note that our results hold independently of this assumption on Δ. This is because the dynamics of our model do not change

[3] For general Δ, we have $P_1 \in \{0, \Delta, 1, 1 + \Delta\}$.

even with the addition of a further possible price: the probability of a given price occurring decreases in n, while the probability of the contract supplier being inefficient increases with price. As such, our results remain qualitatively unchanged.

5.5.3 Firm as sole beneficiary of a resolved inefficiency

Another assumption of our model is that a resolved inefficiency in the second period exclusively benefits the firm and not the contract supplier. By relaxing this assumption and sharing the cost reduction created by resolving inefficiencies, the supplier's bidding strategies in the first period are altered. In our current setting, bidding one's true cost is a dominant strategy, which simplifies the analysis and avoids the need for distributional assumptions regarding supplier types. However, if suppliers anticipate capturing a portion of future gains, they may choose to bid below their costs to increase their probability of winning the contract. This behavior reflects a willingness to incur short-term losses in exchange for expected future gains. However, because these gains are uncertain and depend on both the probability of inefficiency and the firm's investment decision, bidding below cost is not ex-post optimal. Consequently, dominant strategies no longer exist, and optimal supplier offers become belief-dependent, requiring additional assumptions regarding their type distributions for analysis. We expect the degree of distortion in bidding behavior to increase with the size of the potential future gains. As such, when the prospective gains are small, the supplier's incentives remain closely aligned with the baseline case analyzed in this study, and the qualitative nature of our results is preserved.

5.5.4 Assigning decisions to different managers

When it comes to attempts to mitigate the trade-off procurement managers face in our model, two prominent possibilities come to mind. First, that of assigning the qualification and investment decisions to different procurement managers. This would remove the dynamic interdependence between bonus periods, ensuring that qualification and investment decisions are not distorted by concerns about future cost-saving targets. However, this approach might only work in markets for commoditized goods, where deep knowledge of the product is not necessary for a successful awarding and no significant synergy effects exist between initial awarding and future price negotiations. Indeed, as illustrated by Cöster et al. (2023), domain-specific expertise often prevents such a division of qualification and investment decisions across more than one procurement manager.

5.5.5 Incentivizing action instead of outcome

Another possibility would be to make the bonus in $t = 1$ dependent not on the price that was attained, but rather to explicitly incentivize managers to build competition and strengthen the firm's negotiation position. This fosters supplier development and the qualification of new suppliers. Observe that even when directly rewarding the manager for qualification, the fundamental trade-off of the procurement manager remains unchanged: qualifying an additional supplier increases the first-period bonus at the expense of the second-period bonus. Moreover, there are also risks associated with such an approach, namely that such incentives could lead to over-qualification or the inclusion of non-competitive suppliers. As a result, such an incentive scheme need not necessarily reduce total costs to the firm. In fact, if the procurement manager qualifies a high-cost inefficient supplier, they are actually better off than under this incentive scheme than under the one we consider in our model. This is because such a qualification allows the procurement manager to receive the first-period bonus without reducing their second-period bonus. Therefore, incentive schemes that encourage the creation of competition require, among other safeguards, a performance measurement system that looks at more than just the number of qualified suppliers.

5.6 Conclusion

In this paper, we formally examine the trade-offs faced by a procurement manager with cost-saving incentives. In our model, the procurement manager makes qualification and investment decisions on the firm's behalf. While both qualification and investment reduce procurement costs, we show that there is an inherent trade-off in any incentive scheme that solely rewards cost savings: achieving low costs upfront makes future cost savings harder to attain.

In our model, this means that there is a conflict between qualification and investment for the procurement manager. Specifically, qualifying an additional supplier enhances competition and increases the manager's first-period bonus, but simultaneously reduces the scope for achieving savings in subsequent negotiations, thereby lowering the second-period bonus. As a result, we show that—even under optimal incentivization—delegation of qualification and investment decisions to procurement managers directly translates into higher procurement expenses for the firm. This result formalizes the anecdotal evidence we have previously discussed regarding procurement managers not acting in the firm's best interest whenever they are offered cost-saving incentives.

The managerial implications of this are clear: firms should holistically consider the incentives their procurement managers face. What appears straightforward at first—rewarding managers for driving down costs—proves counterproductive when considered in a larger context. Managers are forced to trade-off between short-term and long-term bonuses, resulting in decisions that undermine the firm's broader sourcing strategy. While our model outlines that a firm's ability to commit to an incentive scheme upfront can mitigate these trade-offs, we also show that outcomes remain inefficient. This highlights the limitations of purely savings-based incentive schemes. Moreover, as discussed above, overcoming this problem in practice may prove difficult. Domain-specific expertise often prohibits the simple splitting of qualification and investment decisions across different procurement managers. Similarly, incentivizing the creation of competition instead of rewarding low prices may result in over-qualification or even the qualification of non-competitive suppliers.

All in all, our analysis shows that designing effective incentive schemes in the field of procurement is a problem of dynamic optimization rather than static measurement. Future research could contribute by developing incentive schemes along the proposed lines that are applicable in practice.

Appendix of Chapter 5

5.A Appendix

5.A.1 Appendix: First-best

Proposition 20. *In all cases in which investment takes place at at least one price realization (at least one cut-off price is below 2), there exists $\bar{i}$ for which investment is lower under delegation than under first-best.*

Proof. Observe that proposition 14 holds for all $\bar{i}$. The claim immediately follows. □

Proposition 21. *In all cases in which investment takes place at at least one price realization (at least one cut-off price is below 2), there exists $\bar{i}$ for which qualification is lower under delegation than under first-best.*

Proof. Observe first that, by definition, qualification can never be higher under delegation. For certain values of $\bar{i}$, the qualification decision does not differ across delegation and first-best. However, in the delegation setting, the firm must account for the procurement manager's conflict between qualification and investment. This results in an expected payment to the manager in $t = 1$ that is greater than q. In other words, to get the manager to qualify, the firm must compensate the manager for more than the qualification costs. This implies that there exist values for $\bar{i}$ for which qualification occurs under first-best, but not under delegation.

To see this, consider the case of $q = q_{FB}$. That is, q is equal to the highest qualification cost for which the firm will qualify the additional supplier under first-best. Above, we have just argued above that under delegation, managers need be over-compensated for qualification. As such, for the $\bar{i}$ for which $q = q_{FB}$, the firm will not incentivize the manager to qualify. □

5.A.2 Appendix: No-commitment

Derivation of the maximum qualification cost per interval for which the firm is willing to incentivize the manager:

Lemma 4. *When* $\bar{i} \in [\frac{1}{4}, \frac{3}{10})$*, the firm wants the manager to invest for* $P_1 = 2$ *for both* $n = 2$ *and* $n = 3$ *and incentivizes the manager to qualify whenever*

$$q \leq q_{NC}^{(1)} := \frac{3}{5} \cdot (\frac{81}{128} + \frac{21}{128}\bar{i}) \tag{5.9}$$

Proof. We know that for $\bar{i} \in [\frac{1}{4}, \frac{3}{10})$, the firm wants the manager to invest for $P_1 = 2$ for both $n = 2$ and $n = 3$. That is, we have $\hat{P}_{NC} = \hat{P} = 2$. In this case The manager's total profit is then

$$\frac{1}{16}B_1 + \frac{7}{64}\bar{i} \tag{5.10}$$

for $n = 2$, and

$$\frac{5}{32}B_1 + \frac{5}{128}\bar{i} - q \tag{5.11}$$

for $n = 3$.

Using these two equations, we find that the manager will qualify the additional supplier whenever the first-period bonus is at least

$$B_1^* = \frac{32}{3}q + \frac{3}{4}\bar{i} \tag{5.12}$$

Does the firm incentivize the manager to qualify? This depends on the firm's total costs.

If the firm offers B_1^*, their total cost is given by

$$TC^n = \begin{cases} \mathbb{P}^n[P_1 = 0] \cdot B_1, \text{if } B_1 = B_1^* \\ 0, \text{if } B_1 = 0 \end{cases}$$

$$+ \sum_p \mathbb{P}^n[p] \cdot p + \sum_{p<\hat{P}_{NC}} \mathbb{P}^n[p] \cdot (p - \frac{1}{4} \cdot \mathbb{P}^n[\Delta_{cs} = 1|p]) + \sum_{p \geq \hat{P}_{NC}} \mathbb{P}^n[p] \cdot (p - \mathbb{P}^n[\Delta_{cs} = 1|p] + \bar{i}) \tag{5.13}$$

This gives us:

$$TC^{n=2} = \frac{81}{32} + \frac{7}{16}\bar{i} \tag{5.14}$$

and

$$TC^{n=3} = \frac{243}{128} + \frac{35}{128}\bar{i} + \frac{5}{3}q \tag{5.15}$$

□

Lemma 5. *When $\bar{i} \in [\frac{3}{10}, \frac{9}{28})$, the firm only incentivizes investment for $P_1 = 2$ if $n = 2$ and incentivizes the manager to qualify whenever*

$$q \leq q_{NC}^{(2)} := \frac{3}{5} \cdot (\frac{75}{128} + \frac{49}{192}\bar{i}) \tag{5.16}$$

Proof. We know that for $\bar{i} \in [\frac{3}{10}, \frac{9}{28})$, the firm only incentivizes investment for $P_1 = 2$ if $n = 2$. In the case of $n = 3$, the firm does not incentivize investment; formally, $\hat{P}_{NC}^{n=2} = \hat{P} = 2$ and $\hat{P}_{NC}^{n=3} > 2$.The manager's total profit is then

$$\frac{1}{16}B_1 + \frac{7}{64}\bar{i} \tag{5.17}$$

for $n = 2$, and

$$\frac{5}{32}B_1 - q \tag{5.18}$$

for $n = 3$.

Using these two equations, we find that the manager will qualify the additional supplier whenever the first-period bonus is at least

$$B_1^* = \frac{32}{3}q + \frac{7}{6}\bar{i} \tag{5.19}$$

Does the firm incentivize the manager to qualify? This depends on the firm's total cost.

Using eq. (5.13), firm total cost is

$$TC_{n=2} = \frac{81}{32} + \frac{7}{16}\bar{i} \tag{5.20}$$

and

$$TC_{n=3} = \frac{249}{128} + \frac{35}{192}\bar{i} + \frac{5}{3}q \tag{5.21}$$

We find that the firm incentivizes the manager to qualify whenever

$$q \leq q_{NC}^{(2)} := \frac{3}{5} \cdot \left(\frac{75}{128} + \frac{49}{192}\bar{i}\right) \tag{5.22}$$

□

Lemma 6. *When $\bar{i} \in [\frac{9}{28}, \frac{2}{5}]$, the firm never wants to incentivize the manager to invest and but incentivizes the manager to qualify whenever*

$$q \leq q_{NC}^{(3)} := \frac{3}{5} \cdot \frac{93}{128} \tag{5.23}$$

Proof. When $\bar{i} \in [\frac{9}{28}, \frac{2}{5}]$, the firm never wants to incentivize the manager to invest. That is, $\hat{P}_{NC} > 2$. In other words, the firm offers no bonus in $t = 2$, and thus, all manager types make zero profit in the second stage.

The manager's total profit is then

$$\frac{1}{16}B_1 \tag{5.24}$$

for $n = 2$, and

$$\frac{5}{32}B_1 - q \tag{5.25}$$

for $n = 3$.

Using these two equations, we find that the manager will qualify the additional supplier whenever the first-period bonus is at least

$$B_1^* = \frac{32}{3}q \tag{5.26}$$

Does the firm incentivize the manager to qualify? This depends on the firm's total cost.

Using eq. (5.13), firm total cost is

$$TC_{n=2} = \frac{171}{64} \tag{5.27}$$

and

$$TC_{n=3} = \frac{249}{128} + \frac{5}{3}q \tag{5.28}$$

We therefore find that the firm incentivizes the manager to qualify whenever

$$q \leq q_{NC}^{(3)} := \frac{3}{5} \cdot \frac{93}{128} \tag{5.29}$$

This concludes the proof. □

5.A.3 Appendix: Commitment

Interval 1: $\bar{i} \in [\frac{1}{4}, \frac{3}{10})$ In the no-commitment setting discussed above, we saw that in $t = 2$, for all n, the firm incentivizes the manager to invest for $P_1 = 2$. Given this price realization, this is the firm's optimal strategy. However, what if, at the beginning of $t = 1$, the firm could credibly commit to a different second-period bonus scheme?

First, note that offering $B_2 > B_2^*$ is not beneficial to the firm, as all manager types already invest for B_2^*. Second, offering $0 < B_2 < B_2^*$ is not beneficial either, since under such a bonus, only low-cost managers invest, but these managers would also do so for $B_2 = 0$. Therefore, the only potentially beneficial alternative is to commit setting $B_2 = 0$.

We must analyze two subcases within Interval 1.

Sub-case (i): $q \leq q_{NC}^{(1)}$ In the no-commitment setting, the firm incentivizes the manager to qualify and invest for all n. In the commitment setting, the firm has two additional options available: (a) incentivize investment only for $n = 2$, or (b) refrain from incentivizing investment altogether.

Comparing eq. (5.15) with eq. (5.21), we find that deviating to option (a) is only worthwhile for values of $\bar{i}$ larger than the maximum value of $\bar{i}$ in Case 1, which is a contradiction.

By comparing eq. (5.15) with eq. (5.28), we find that option (b) yields a lower total cost to the firm for all values of $\bar{i}$ permitted in Case 1. In other words, the ability to commit to an ex-post suboptimal bonus scheme in $t = 2$ is beneficial for the firm.

Sub-case (ii): $q > q_{NC}^{(1)}$ In the no-commitment setting, the firm does not incentivize qualification, but it incentivizes investment for all n. In the commitment setting, the firm has three other options available: (a) do not incentivize investment at all, (b) incentivize qualification but investment only for $n = 2$, and (c) incentivize qualification but no investment.

For option (a), comparing eq. (5.14) with eq. (5.27), we find a contradiction for the investment cost $\bar{i}$.

Similarly, for option (b), comparing eq. (5.14) with eq. (5.21), we find a contradiction for the qualification cost q.

However, comparing eq. (5.14) with eq. (5.28), we find that option (c) is worth it $\forall q \in (q_{NC}^{(1)}, q_C)$, with

$$q_C := \frac{3}{5} \cdot (\frac{75}{128} + \frac{7}{16}\bar{i}) \tag{5.30}$$

In summary, for Interval 1, we find that through decreased investment and increased qualification, the commitment setting leads to lower total cost to the firm vis-à-vis the no-commitment setting.

Interval 2: $\bar{i} \in [\frac{3}{10}, \frac{9}{28})$ In the no-commitment setting, in $t = 2$, the firm incentivizes the manager to invest only for $n = 2$.

Similar to the above, we must analyze two sub-cases within Interval 2.

Sub-case (i): $q \leq q_{NC}^{(2)}$ In the no-commitment setting, the firm incentivizes qualification but not investment. Since no investment takes place in the no-commitment setting, for this sub-case, there is no difference between commitment and no-commitment.

Sub-case (ii): $q > q_{NC}^{(2)}$ In the no-commitment setting, the firm does not incentivize qualification, but it incentivizes investment for $n = 2$. In the commitment setting, the firm has two other options available: (a) do not incentivize investment at all, and (b) incentivize qualification but no investment.

For option (a), comparing eq. (5.20) with eq. (5.27), we find a contradiction for the investment cost $\bar{i}$.

However, comparing eq. (5.20) with eq. (5.28), we find that option (b) is beneficial $\forall q \in (q_{NC}^{(2)}, q_C)$, with q_C as defined above.

In summary, for Interval 2, we find that through decreased investment and increased qualification, the commitment setting leads to lower total costs to the firm vis-à-vis the no-commitment setting.

References

Arrow, K. J., Harris, T. and Marschak, J. (1951), ‘Optimal inventory policy’, *Econometrica: Journal of the Econometric Society* pp. 250–272.

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

Baron, D. P. and Besanko, D. (1987), ‘Commitment and fairness in a dynamic regulatory relationship’, *Review of Economic Studies* **54**(3), 413–436.
URL: *https://www.doi.org/10.2307/2297567*

Baron, D. P. and Myerson, R. B. (1982), ‘Regulating a monopolist with unknown costs’, *Econometrica* **50**(4), 911–930.
URL: *https://www.jstor.org/stable/1912769*

Brandstätter, H. (1988), ‘Sechzehn persönlichkeits-adjektivskalen (16 pa) als forschungsinstrument anstelle des 16 pf.’, *Zeitschrift für experimentelle und angewandte Psychologie* .

Bulow, J. and Klemperer, P. (1996), ‘Auctions versus negotiations’, *The American Economic Review* **86**(1), 180–194.

CAPS, R. (2022), ‘Managed spend growing influence’, *The Metrics of Supply Management: A Cross-Industry Report 2021* . last accessed: 24.07.2023.
URL: *https://www.capsresearch.org/blog/posts/2022/january/caps-stats-managed-spend-growing-influence/*

Chen, D. L., Schonger, M. and Wickens, C. (2016), ‘oTree—an open-source platform for laboratory, online, and field experiments’, *Journal of Behavioral and Experimental Finance* **9**, 88–97.

Chen, F. (2007), ‘Auctioning supply contracts’, *Management Science* **53**(10), 1562–1576.

Chwe, M. S.-Y. (1989), 'The discreet bid first-price auction', *Economic Letters* **31**(4), 303–306.

Cöster, M., Danielson, M., Ekenberg, L., Gullberg, C., Titlestad, G., Westelius, A. and Wettergren, G. (2023), 'Procurement competence'.

Coviello, D., Guglielmo, A. and Spagnolo, G. (2018), 'The effect of discretion on procurement performance', *Management Science* **64**(2), 715–738.

Cox, J. C., Roberson, B. and Smith, V. L. (1982), 'Theory and behavior of single object auctions', *Research in Experimental Economics* **2**(1), 1–43.

Crocker, K. J. and Reynolds, K. J. (1993), 'The efficiency of incomplete contracts: An empirical analysis of air force engine procurement', *The RAND Journal of Economics* pp. 126–146.

Cui, R., Li, M. and Zhang, S. (2022), 'AI and procurement', *Manufacturing & Service Operations Management* **24**(2), 691–706.

Dai, T., Wang, X. and Hwang, C.-W. (2022), 'Clinical ambiguity and conflicts of interest in interventional cardiology decision making', *Manufacturing & Service Operations Management* **24**(2), 864–882.

Duenyas, I., Hu, B. and Beil, D. R. (2013), 'Simple auctions for supply contracts', *Management Science* **59**(10), 2332–2342.

Dufwenberg, M. and Gneezy, U. (2000), 'Price competition and market concentration: an experimental study', *International Journal of Industrial Organization* **18**(1), 7–22.

Ellram, L. M. and Tate, W. L. (2021), 'Cost avoidance: Not everything that counts is counted', *Journal of Business Logistics* **42**(4), 406–427.
URL: *https://doi.org/10.1111/jbl.12262*

Ellram, L. M., Tate, W. L. and Choi, T. Y. (2020), 'The conflicted role of purchasing in new product development costing', *Journal of Supply Chain Management* **56**(1).
URL: *https://doi.org/10.1111/jscm.12217*

Elmaghraby, W. J. (2000), 'Supply contract competition and sourcing policies', *Manufacturing & Service Operations Management* **2**(4), 350–371.

Engelbrecht-Wiggans, R. and Katok, E. (2007), 'Regret in auctions: Theory and evidence', *Economic Theory* **33**(1), 81–101.

European Union (2022), EU COVID-19 vaccine procurement, Technical report, European Union.

Fehr, E. and Gächter, S. (2000), 'Fairness and retaliation: The economics of reciprocity', *Journal of Economic Perspectives* **14**(3), 159–182.

Freixas, X., Guesnerie, R. and Tirole, J. (1985), 'Planning under incomplete information and the ratchet effect', *Review of Economic Studies* **52**(2), 173–191.
URL: *https://www.doi.org/10.2307/2297615*

Fugger, N., Gretschko, V. and Pollrich, M. (2019), Sequential procurement with limited commitment, Discussion Paper 19–030, ZEW.

Fugger, N., Katok, E. and Wambach, A. (2016), 'Collusion in dynamic buyer-determined reverse auctions', *Management Science* **62**(2), 518–533.

Gulati, R. (1995), 'Does familiarity breed trust? The implications of repeated ties for contractual choice in alliances', *Academy of Management Journal* **38**(1), 85–112.

Gur, Y., Lu, L. and Weintraub, G. Y. (2017), 'Framework agreements in procurement: An auction model and design recommendations', *Manufacturing & Service Operations Management* **19**(4), 586–603.

Hart, O. and Moore, J. (2008), 'Contracts as reference points', *Quarterly Journal of Economics* **123**(1), 1–48.

Helper, S. and Sako, M. (2010), 'Management innovation in supply chain: Appreciating Chandler in the twenty-first century', *Industrial and Corporate Change* **19**(2), 399–429.

Holstrom, B. (1982), 'Design of incentive schemes and the new soviet incentive model', *European Economic Review* **17**(2), 127–148.
URL: *https://doi.org/10.1016/S0014-2921(82)80010-X*

Holt, C. A. and Laury, S. K. (2002), 'Risk aversion and incentive effects', *American Economic Review* **92**(5), 1644–1655.

Inverto (2025*a*), 'Combining e-auctions & game theory for better results', *Inverto Insights* .
URL: *https://inverto.com/en/insights/why-combining-e-auctions-and-game-theory-can-deliver-better-results*

Inverto (2025*b*), 'Next level procurement negotiations–using game theory and psychology', *Inverto Insights* .
URL: *https://inverto.com/en/expertise/negotiation-excellence*

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

John, O. P., Srivastava, S. et al. (1999), 'The big-five trait taxonomy: History, measurement, and theoretical perspectives'.

Johnson, P. F. and Leenders, M. R. (2010), 'Minding the supply savings gaps', *MITSloan Management Review* **51**(2), 25–31.

Kagel, J. H. and Levin, D. (1993), 'Independent private value auctions: Bidder behaviour in first-, second-and third-price auctions with varying numbers of bidders', *The Economic Journal* **103**(419), 868–879.

Katok, E. and Tan, L. (2025), 'Renegotiations in the presence of supply disruptions', *Production and Operations Management* **34**(2), 153–167.

Khouja, M. (1999), 'The single-period (news-vendor) problem: Literature review and suggestions for future research', *Omega* **27**(5), 537–553.

Klemperer, P. (1999), 'Auction theory: A guide to the literature', *Journal of Economic Surveys* **13**(3), 227–286.

Krishna, V. (2010), *Auction Theory*, 2. edn, Elsevier.

Laffont, J.-J. and Tirole, J. (1986), 'Using cost observations to regulate firms', *Journal of Political Economy* **94**(3), 614–641.
URL: *https://www.jstor.org/stable/1833051*

Laffont, J.-J. and Tirole, J. (1987), 'Comparative statics of the optimal dynamic incentive contract', *European Economic Review* **31**(4), 901–926.
URL: *https://doi.org/10.1016/0014-2921(87)90006-7*

Laffont, J.-J. and Tirole, J. (1988), 'The dynamics of incentive contracts', *Econometrica* **56**(5), 1153–1175.
URL: *https://www.jstor.org/stable/1911362*

Li, C. and Scheller-Wolf, A. (2011), 'Push or pull? Auctioning supply contracts', *Production and Operations Management* **20**(2), 198–213.

Liu, Q., Mierendorff, K., Shi, X. and Zhong, W. (2019), 'Auctions with limited commitment', *American Economic Review* **109**(3), 876–910.

Lucking-Reiley, D. (2000), 'Vickrey auctions in practice: From nineteenth-century philately to twenty-first-century e-commerce', *Journal of Economic Perspectives* **14**(3), 183–192.

Macaulay, S. (2018), Non-contractual relations in business: A preliminary study, *in* 'The sociology of economic life', Routledge, pp. 198–212.

Malhotra, D. and Lumineau, F. (2011), 'Trust and collaboration in the aftermath of conflict: The effects of contract structure', *Academy of Management Journal* **54**(5), 981–998.

McAfee, R. P. and McMillan, J. (1987), 'Auctions and bidding', *Journal of Economic Literature* **25**(2), 699–738.

McKinsey & Company (2014), 'A better procurement incentive model', *McKinsey Insights–Operations* .
URL: *https://www.mckinsey.com/capabilities/operations/our-insights/a-better-procurement-incentive-model*

Mogge, F. (2023), 'Automotive suppliers steer through turbulent waters'.
URL: *https://www.rolandberger.com/en/Insights/Publications/Automotive-suppliers-steer-through-turbulent-waters.html*

Poppo, L. and Zenger, T. (2002), 'Do formal contracts and relational governance function as substitutes or complements?', *Strategic Management Journal* **23**(8), 707–725.

PwC (2025), 'Optimise purchasing and procurement using game theory', *PwC Insights–Game Theory in Negotiations and Auctions* .
URL: *https://www.pwc.de/en/strategy-organisation-processes-systems/operations/procurement/game-theory-in-negotiations-and-auctions.html*

Ring, P. S. and Van de Ven, A. H. (1994), 'Developmental processes of cooperative interorganizational relationships', *Academy of Management Review* **19**(1), 90–118.

Rubinstein, A. (1982), 'Perfect equilibrium in a bargaining model', *Econometrica* pp. 97–109.

Scharlach, A. (2024), 'Wrong incentives (almost) everywhere?'. last accessed: 20.09.2024.
URL: $https://www.linkedin.com/posts/arminscharlach_wrong - incentives - almost - everywhere - i - activity - 7164211837436129281 - NiUN/$

Schelling, T. (1960), 'The strategy of conflict', *Harvard University* .

Schummer, J. and Vohra, R. V. (2003), 'Auctions for procuring options', *Operations Research* **51**(1), 41–51.

Shachat, J. and Swarthout, J. T. (2013), 'Auctioning the right to play ultimatum games and the impact on equilibrium selection', *Games* **4**(4), 738–753.

Shachat, J. and Tan, L. (2015), 'An experimental investigation of auctions and bargaining in procurement', *Management Science* **61**(5), 1036–1051.

Taylor, T. A. and Plambeck, E. L. (2007), 'Simple relational contracts to motivate capacity investment: Price only vs. price and quantity', *Manufacturing & Service Operations Management* **9**(1), 94–113.

Thaler, R. H. (1988), 'Anomalies: The winner's curse', *Journal of Economic Perspectives* **2**(1), 191–202.

Toyota (2024), 'A better procurement incentive model', *Toyota Times* .
URL: $https://toyotatimes.jp/en/toyota_news/together_with_our_suppliers/001.html$

Tunca, T. I. and Zenios, S. A. (2006), 'Supply auctions and relational contracts for procurement', *Manufacturing & Service Operations Management* **8**(1), 43–67.

Tversky, A. and Kahneman, D. (1974), 'Judgment under uncertainty: Heuristics and biases: Biases in judgments reveal some heuristics of thinking under uncertainty.', *Science* **185**(4157), 1124–1131.

Van Dijk, E., De Cremer, D. and Handgraaf, M. J. (2004), 'Social value orientations and the strategic use of fairness in ultimatum bargaining', *Journal of Experimental Social Psychology* **40**(6), 697–707.

Vickrey, W. (1961), 'Counterspeculation, auctions, and competitive sealed tenders', *Journal of Finance* **16**(1), 8–37.

Williamson, O. E. (1985), The economic institutions of capitalism. Firms, markets, relational contracting, *in* 'Das Summa Summarum des Management: Die 25 wichtigsten Werke für Strategie, Führung und Veränderung', Springer, pp. 61–75.

Zhang, F. (2010), 'Procurement mechanism design in a two-echelon inventory system with price-sensitive demand', *Manufacturing & Service Operations Management* **12**(4), 608–626.