

Experimental Evidence on Covert Bargaining Markets*

Margherita Comola[§] and Marcel Fafchamps[¶]

December 26, 2025

Abstract

We conduct a laboratory experiment to study a decentralized market where players engage in multiple (bilateral and multilateral) transactions. We propose a novel class of bargaining protocols that allow players to keep bid amounts shrouded from each other (covert bargaining). We show that these bargaining protocols double ex-post efficiency relative to a mechanism without bargaining, mainly to the benefit of players (particularly buyers) rather than the silent auctioneer. Aggregate efficiency nonetheless suffers from the fact that buyers bargain harder than sellers and that some players over-bargain to appropriate a larger share of the unknown surplus.

JEL codes: C92; C78; D47; D82

Keywords: decentralized market; bargaining; private information; experiment

*We are grateful to Roland Benabou, Francis Bloch, Olivier Bos, Yann Bramoullé, Olivier Compte, Aureo de Paula, Nicolas Jacquemet, Philippe Jehiel, Luca Merlino, Andras Niedermayer, Arno Riedl, Jean-Christophe Vergnaud, attendees to the 6th BEET Workshop, and numerous seminar participants for helpful comments and discussions. We thank Maxim Frolov and Shaden Shabayek for outstanding research assistance. Margherita Comola acknowledges the support of the grants ANR-17-EURE-0001, ANR-21-CE26-0002-01.

[§]Université Paris-Saclay (RITM) and Paris School of Economics. margherita.comola@psemail.eu

[¶]Stanford University, NBER, CEPR. fafchamp@stanford.edu

1 Introduction

Economists have held a longstanding interest in decentralized markets, where participants are free to interact within the bounds of a protocol that regulates transactions (e.g., Ochs and Roth 1989; Kriss et al. 2013; Milgrom and Segal 2020, Gaineddenova 2022). Our paper belongs to this literature and focuses on the design of bargaining markets. We propose a novel bargaining mechanism, that we name ‘covert bargaining’, where agents can keep bid amounts shrouded from each other. The purpose of this innovation is to facilitate discovery by reducing the disclosure risk of placing a high bid to identify beneficial trades. We show how this mechanism, which is simple and intuitive, can regulate a market of bilateral and multilateral transactions (i.e., among 3 or more parties). We study its performance via a laboratory experiment, showing that it fosters beneficial exchanges and allows players to extract most of the surplus from trade.

Our experimental design consists in an interactive trading floor populated by multiple agents. During the game’s unfolding, buyers and sellers have the opportunity to place and revise bids that could be positive (i.e., offers) or negative (i.e., requests), according to a bargaining protocol that we describe below. If the final sum of bids is strictly positive (i.e., the amount offered is strictly above the amount requested), the transaction occurs and each agent pays a price corresponding to her final bid. The leftover worth generated by the difference in prices goes to the auctioneer. For example, say a transaction has value -5 for the seller and $+10$ for the buyer, meaning that, if the transaction occurs, the seller incurs a cost of 5 while the buyer derives a benefit of 10 . If the final bids are -6 and $+8$ respectively (i.e., the seller requests 6 while the buyer offers 8), the transaction takes place since the sum of bids $8 - 6 = 2 > 0$. Each side pays a price equal to her final bid and gains the difference between her value and her bid. Thus, the seller gains $-5 + 6 = 1$, and the buyer gains $10 - 8 = 2$.¹ The sum of bids (2 units), which represents the value generated by the price disparity and not appropriated by the players, is collected by the auctioneer.²

This design has two novel features. First, the mechanism we propose allows to keep both values *and* bids private throughout the negotiation process (covert bargaining). This is implemented through a silent and automated auctioneer who soaks up the difference between the buying and selling price, as in the double auction of McAfee (1992). While private in-

¹This payoff rule is also used in the Nash demand game (Nash 1953) where players demand a share of a given amount of resources, and receive the amount requested as long as the sum of demands does not exceed the amount of resources available.

²Since we impose that the sum of bids must be strictly positive for a transaction to take place, in our setting there is at least a minimal fee of 1 unit going to the auctioneer when a transaction occurs. This minimal fee could be adjusted as desired (Section 6). In this example, an additional unit is appropriated by the auctioneer on top of his minimal fee.

formation about values is a common assumption, secret reserve prices have received limited theoretical attention (Andreyanov and Caoui, 2022). Extending secrecy to the bids of all agents allows to eschew the thorny issue of preference revelation,³ and to abstract from concerns related to fairness.⁴ Second, this mechanism can regulate a bargaining market with a certain degree of complexity. In our experiment, players bargain over multiple transactions simultaneously, and some of these transactions are among three or more agents. This differentiate us from the standard setting where a single pair of agents negotiate over the terms of a trade (Rubinstein 1982; Chatterjee and Samuelson, 1983), and it mimics the complexity of real world markets where multiple purchases among multiple parties are negotiated at the same time.⁵ In the current paper, different transactions are assumed independent. The application of the covert bargaining mechanism to complement or substitute goods is nonetheless possible, as discussed in Section 6.

Our benchmark is a mechanism where all players place take-it-or-leave-it bids simultaneously and no bargaining is allowed (e.g., Bloch and Jackson 2007; Haerlinger and Wooders 2011). As above, bids placed stay secret and the auctioneer soaks up the difference between the buying and selling price. We characterize this setting theoretically (Appendix B), and show that it is generally quite inefficient. We then introduce three novel protocols of covert bargaining, where different transactions are negotiated sequentially or simultaneously. These protocols impose minimal constraints on the bargaining process, i.e., just enough to be implementable in an orderly way while preserving bids privacy. Also, they allow for a small window of time to retract or renegotiate (i.e., bargaining without commitment as in Muthoo 1990) – a feature that is included in many real-life transactions, such as digital purchases and banking services.⁶ In these covert bargaining environments, ex-post efficiency and the

³In a standard bilateral bargaining game, the party who discloses their true reservation value is placed at a considerable disadvantage, as the other party can appropriate more of the surplus. As a result, both sides hesitate to reveal their reservation values, which hinders discovery and can prevent beneficial exchange from taking place. Several markets have similarly been shown to be vulnerable to truthful preference revelation (e.g., Roth and Peranson 1999; Lee 2017; Bodoh-Creed 2020).

⁴Fairness concerns appear to be less relevant when agents are less informed (Huang et al. 2022). In our game players ignore the amount of surplus generated in each transaction and who appropriates it. In our earlier example, players are not aware that the transaction would have taken place even if the seller raised her final request (-7 rather than -6 , increasing her gain by 1 unit) or if the buyer lowered her final offer (7 rather than 8 , increasing her gain by 1 unit).

⁵Flexible buyer-seller roles have been explored in the context of asset markets or markets for environmental permits (e.g., Cason et Gangadharan 2006, Sherstyuk et al. 2021).

⁶Most digital trading platform allow for full-refund order cancellation within a short time span (e.g., 30 minutes for Amazon, 24 hours for Expedia, 48 hours for AirBnb, 7 days for Microsoft Business software subscriptions) or free return (e.g., within 30 days from reception for many clothes retailers). Some legal systems give buyers a time window to withdraw from real-estate transactions without justification or penalty (e.g., 10 days in France). This feature is also present in the experiment by Camerer et al. (2019), where bargaining parties have 1.5 seconds to retract.

division of gains from trade cannot be determined theoretically, which motivates our empirical investigation. Whether such environments can sustain efficient exchange is the central focus of the paper.

We study the performance of the covert bargaining mechanism along three dimensions: ex-post efficiency, division of surplus, and bidding dynamics. Our experimental design provides a clear-cut criterion to measure efficiency: all transactions for which the sum of values is strictly positive (i.e., profitable transactions) should occur. However, our results indicate that many profitable transactions never take place. This matches theoretical predictions about the incomplete information penalty inducing delays and inefficiencies in trade (e.g., Myerson and Satterthwaite 1983; Sobel and Takahashi 1983; Williams 1987; Vincent 1989; Ausubel et al. 2002). We nonetheless find that covert bargaining markedly improves ex-post efficiency, doubling the proportion of deals on profitable transactions with respect to sealed bids (from 27% to over 60% in all three bargaining protocols). When we look at the division of gains from trade, we find that the surplus is mostly appropriated by buyers. The share of the surplus going to sellers remains low and stable in all protocols. This results reconfirms a previously documented earning gap in favor of buyers (Smith and Williams 1982; Sherstyuk et al. 2021).⁷ The same is true for the auctioneer, who only appropriates a small fraction of the surplus on average. Thus, when covert bargaining is possible, players are surprisingly good at not leaving money on the table, and the fact that bids remain secret does not result in an advantage for the auctioneer as one could have expected.

We exploit the richness of our data to investigate how these results relate to the behavioral strategies adopted by participants. We show that most players gradually increase their bids – mostly by one unit at a time – until a transaction is formed. This cautious bidding-up strategy, which overcomes the lack of information, rewards unconstrained bargaining environments. We identify two main types of frictions explaining why many profitable transactions do not occur, even when covert bargaining is possible. First, we find that buyers bargain harder than sellers: the implicit profit margin asked by buyers is 70% compared to 33% for sellers. This complements the well-known experimental result that selling prices are at least twice as large as buying prices – the so-called willingness to accept/willingness to pay ratio. This disparity has either been attributed to an endowment effect (e.g., Kahneman and Tversky 1979) or seen as an anomaly driven by inexperience, psychological traits, or poor experimental practice (e.g., Coursey et al. 1987; List 2003; Plott and Zeiler 2005; Georgantzís and Navarro-Martínez 2010; Isoni et al. 2011; Cason and Plott 2014). Our results show that this empirical pattern is associated with low offers rather than high requests, and

⁷In an experimental setting where trader roles are flexible as in ours (i.e. any trader can both buy and sell), Sherstyuk et al. (2021) find that the average earnings of buyers are 2.5 times the ones of sellers.

that it is not mitigated by bargaining dynamics. We interpret this findings in light of the market experience hypothesis, as subjects are more accustomed to bargain as buyers than as sellers (an argument already put forward by Smith and Williams 1982).⁸ Secondly, we observe that certain deals are delayed or prevented by a minority of subjects who deviate from the prevalent bidding-up strategy – presumably in an attempt to appropriate a larger share of the unknown surplus.

Our paper advances the knowledge of unstructured bargaining games, a topic of major relevance that has received surprisingly little attention by economists.⁹ Many off-line markets for differentiated goods and services already allow for unregulated negotiations between parties: e.g., collectibles, art work, second-hand cars, real estate and construction contracts, to name but a few. Furthermore, many types of differentiated goods are increasingly traded online, and we observe an exponential growth of digital marketplaces for which covert bargaining could actually be implemented and our proposed algorithm used to initiate trade while keeping offers hidden. While unstructured bargaining is ubiquitous in real-life situations, its theoretical indeterminacy has hindered economic research. This is particularly true for bargaining games involving three or more players, as we have in our experiment.¹⁰ As a result, applied economists have focused most of their attention either on auctions where bargaining is one-sided,¹¹ or on highly-structured bilateral bargaining protocols where theory gives clear and testable predictions.¹² Evidence on bargaining behavior with other protocols is limited, and the literature has focused on specific dimensions such as: fairness (e.g., Forsythe et al. 1991; Kroll et al. 2014; Galeotti et al. 2019; Luhan et al. 2019; Navarro and Veszteg 2020, Keniston et al. 2021); information (e.g., Kirchsteiger et al. 2005; Shupp et al. 2013; Agranov and Tergiman 2014; Backus et al. 2019; Camerer et al. 2019; Goeree and Lindsay 2020); and congestion (e.g., Kagel and Roth 2000; Che and Koh 2016; Ab-

⁸Smith and Williams (1982) show that in double auction markets subjects are better at being buyers, in that buyers tend to earn a larger share of the surplus. They argue that this could be driven by the fact that most subjects have more experience with the role of a buyer than with the role of a seller.

⁹There is a sizable literature in social psychology on bargaining abilities in real-life situations – see, for instance, Rubin and Brown (1975), Pruitt (2013), and Morley and Stephenson (2015).

¹⁰Most models of bargaining among multiple individuals have a unanimity or majority closing rule (e.g., Baron and Ferejohn 1989; Ali 2006; Tremewan and Vanberg 2016; Agranov et al. 2020). In our setting, partial agreement is possible, making it harder to characterize theoretically (e.g., Bennet 1997; Ambrus and Lu 2015).

¹¹The prominent theoretical literature on auctions (e.g., Milgrom and Weber 1982; Thaler 1988; Klemperer 1996; Jackson and Kremer 2006; Milgrom and Segal 2020) has recently been complemented by observational studies on online auctions (e.g., Horton et al. 2017; Bodoh-Creed et al. 2021).

¹²Notable examples are the infinite-horizon game of alternating offers by Rubinstein (1982), the double auction by Chatterjee and Samuelson (1983), and the exit game by Krishna and Serrano (1996). Experimental studies on structured bargaining protocols include Ochs and Roth (1989), Mitzkewitz and Nagel (1993), Burrows and Loomes (1994), Güth et al. (1996), Kagel and Wolfe (2001), Srivastava (2001), Johnson et al. (2002), Valley et al. (2002), Croson et al. (2003), and Kriss et al. (2013).

dulkadiroglu et al. 2017). These are all dimensions that we purposefully set aside in our experimental design. Some recent papers study bargaining games different from ours, but they reach the same conclusions: efficient exchange under private information is possible (Bochet et al., 2023; Jackson et al., 2023) and adding bargaining features increases market efficiency (Haruvy et al., 2020; Gaineddenova, 2022; Jackson et al, 2023).¹³

Our paper is closely related to a few recent studies on bargaining patterns based on observational data. To the best of our knowledge, Larsen (2021) provides the first empirical assessment on the efficiency of bargaining outcomes with two-sided private values. By estimating implicit value bounds for sequential auctions of used cars, he concludes that 17–24% of profitable negotiations fail – a result that compares well to our experimental findings on the magnitude of efficiency losses. Similarly, Freyberger and Larsen (2025) find that the fraction of failed profitable negotiations on Ebay is 37%, a number which also recalls our findings closely. This result is complemented by Larsen and Zhang (2021), who explore how the total surplus is split among different parties and show that there are efficiency gains in private-information bargaining relative to a sealed bid – two issues that are at the heart of our paper as well. In a related study, Backus et al. (2020) use data on Ebay’s Best Offer platform for collectibles (e.g. coins, antiques, toys, memorabilia, stamps, art) to study the dynamics of bilateral bargaining. They observe that subjects adopt a gradual incremental strategy that aligns with our results on bidding dynamics in the lab. By using an experimental design to overcome the constraints of observational data, our paper corroborates and unifies the results of the papers above.

Our paper has practical relevance for the mechanism design of online marketplaces, for trading or for philanthropic giving (e.g., Carroll 2019). We show that decentralized covert bargaining can be implemented in an orderly and anonymous way in a market populated by multiple agents. Many digital platforms currently admit some limited bargaining action.¹⁴ However, thanks to technological advances, decentralized bargaining features similar to the ones we study are becoming reality (Gaineddenova 2022).¹⁵ The bargaining protocols

¹³Bochet et al. (2023) and Jackson et al. (2023) study a situation where a single pair of individuals bargain over multiple goods that can be bundled together. Haruvy et al. (2020) add bargaining features to an ultimatum game. Gaineddenova (2022) studies a pricing mechanism where buyers make an offer revealing their private information and sellers could bargain over it.

¹⁴For instance, Ebay and many fundraising platforms (e.g., GoFundMe, GiveDirect, Kickstarter) allow for auction-format item listings. Search engines and digital social networks (including Taobao and Facebook) sell online advertising space through second price auctions (e.g., Edelman et al. 2007; Abraham et al. 2020). Some platforms for in-person services (e.g., Care.com) initiate bilateral bargaining among clients: the platform establishes the match, but users negotiate terms and prices in private.

¹⁵In the ride-hailing platform studied by Gaineddenova (2022) price setting is decentralized: riders offer a price for the ride and drivers may accept, decline or bargain over it. Thus, unlike Uber-like platforms, the intermediary has no control over price.

we propose are versatile, can be scaled up for large markets, and they allow for a flexible compensation fee for service providers (as discussed in Section 6). Should a market for multilateral bargaining platforms develop, our results provide useful insights for policy makers willing to regulate the industry.

This paper is organized as follows: Section 2 describes the experiment. Sections 3 to 5 describe our results on: efficiency and surplus; bids and roles; and the dynamics of simultaneous bargaining, respectively. Section 6 concludes. Online Appendix A details the experimental protocol and the instructions to participants. Online appendix B discusses the theoretical motivation. Online Appendix C presents ancillary results.

2 Experimental design

We now describe our experimental design. We start by presenting the general features of the game. We then describe the different experimental treatments. Online Appendix A presents the interface and its visual layout and the instructions given to participants.

2.1 General features

The game

In our experiment, each group of 6 players plays 8 trading games together. The players appear as nodes arranged in a hexagon, and they can form 10 transactions represented as links across nodes.¹⁶ The composition of groups remains constant across the 8 games, while the valuations of players change: the value v_{ij}^k that player k attributes to transaction ij varies across games. As each game unfolds, players form transactions by placing or revising bids, subject to the rules explained in Section 2.2.

Internal players

Because of the visual layout of the game (interface with links across nodes), subjects draw a natural distinction between transactions that involve them directly, and transactions among others. When a player is visually represented at one of the endpoints of a link, we call this player an internal player. Other players are called external.

Valuations v_{ij}^k for internal players $k = i$ or j are drawn uniformly at random in the integer interval $[-10, +10]$. Thus, values can be positive – the player gets a benefit if the transaction occurs, which makes her a buyer – or negative – she pays a production/reservation cost if the

¹⁶Out of the 15 possible pairings of 6 players, only 10 links appear on the screen. This is done to keep the length of the session within reason.

transaction occurs, which makes her a seller.¹⁷ The value can also be zero – in which case the player is neither a buyer nor a seller. In the current experiment we do not partition players into types: each player can simultaneously be a buyer and a seller over different transactions (e.g., Agranov and Elliot 2021). The protocol can however be amended to partition players into types (i.e. either buyers or sellers) without loss of generality. Bids of internal players are fully unconstrained: they can bid any positive integer (which represents an offer), any negative integer (which represents a request), or place no bid at all.¹⁸

External players

In 4 of 8 games, all external players $k \neq i, j$ have $v_{ij}^k = 0$. We call these two-sided games. They represent the standard scenario where transactions are between two parties, a buyer and a seller. In the remaining 4 games, the values for external players $k \neq i, j$ are drawn uniformly at random from the $[0, +10]$ interval. We call these many-sided games. This represent situations in which transactions involve multiple buyers, as in real-life fundraising and charitable contributions settings. Bids of external players are constrained to be non-negative integers. This restriction is imposed to ensure that external players cannot prevent internal players from trading, e.g., out of spite or envy, but it is unessential to the game.¹⁹

Transactions and gains

Let $b_{ij,t}^k$ denote the bid of player k on transaction ij at time t , and call $b_{ij,t} = \sum_k b_{ij,t}^k$ the *total bid* on transaction ij at time t . In our game, a transaction is *provisionally* formed at time t when the total bid is strictly positive, that is, when $b_{ij,t} \geq 1$.²⁰ Importantly, while negotiations are open, bids can be revised, up or down, at no cost.

In our game, subjects only get a monetary reward for the transactions that are concluded at the end of the game. These are the transactions for which the *final* total bid $b_{ij,T} \geq 1$, where T indicates the game's ending. If a transaction is concluded, the gain of a player is

¹⁷This imposes symmetric uniform values for buyers and sellers, as in Chatterjee and Samuelson (1983) and Myerson and Satterthwaite (1983).

¹⁸In our experiment we do not prevent an internal player with a positive or zero value to bid a negative integer (request), or an internal player with negative or zero value to bid a positive integer (offer).

¹⁹Other than that, these bids are unconstrained, i.e. all external players (even those with zero value) can place positive bids, for an amount of their choice, in all games (two-sided and many-sided).

²⁰No player needs to give an implicit consent (in the form of a bid) for a transaction. As long as other players bid enough to make the total bid strictly positive, it is concluded. But by making a sufficiently large negative bid, an internal player can *de facto* preclude a transaction.

computed as her value minus her bid, that is:

$$\begin{aligned} g_{ij,T}^k &= v_{ij}^k - b_{ij,T}^k \text{ if } b_{ij,T} \geq 1 \\ &= 0 \text{ if } b_{ij,T} < 1 \end{aligned}$$

Thus, in our game sellers make a positive gain if the transaction occurs and their request exceeds their cost. Buyers make a positive gain if the transaction occurs and their offer is smaller than their benefit. If a player offers more than her value or requests less than her cost, she makes a loss (i.e., negative gain). If a player bids exactly her value, her gain is 0. This mimics the Nash demand game (Nash 1953), where two players simultaneously announce their requests and receive the amount requested if the sum does not exceed the amount of resources available (otherwise they receive a low disagreement payoff).

The silent auctioneer

If a transaction is concluded at the game's end, the silent auctioneer (played by the computer) collects the total bid. Since the trading rule requires that $b_{ij,T} \geq 1$, for each deal the auctioneer collects at least a minimal fee of 1 unit, which we regard as a remuneration for the service provided. Any amount in excess of the minimal fee represents 'money left on the table' by the players, generated by the difference in prices. The idea is not entirely novel: in the double auction by McAfee (1992), for instance, the equilibrium price for the two sides of the market differs and the auctioneer serves as a budget balancer by soaking up this difference as pure profit.²¹ Our game extends this idea to a bargaining floor populated by multiple agents. Handing over the final total bid to the auctioneer ensures that negotiating parties preserve the privacy of their bids: mechanisms capping the auctioneer's fee would, *ex post*, reveal the bids made by others and the surplus generated by the transaction.²²

By allowing subject to renegotiate after a trade has been provisionally formed, we give them the chance to decrease the auctioneer's share. Whether they are able to do so is an empirical question that motivates our experimental design.

²¹McAfee (1992) observes that this feature mirrors what happens with financial intermediaries, such as professional stock exchange traders, who can profit by buying shares at the seller price and re-selling them at the buyer price.

²²In a standard auction setting with private values, the auctioneer's fees are fixed and the final selling price is observed by the trading parties (Larsen 2021).

Information

Throughout our experiment, players never observe the values or the bids of others. However, we tell them explicitly that values are independent across players and transactions.²³ During the game's unfolding, when a transaction is formed, players can only presume that it is profitable (i.e., that the sum of values is strictly positive). But there is no scope for them to learn which players value the transaction, and by how much.

Visual layout

Much attention went into designing an interface that presents all the relevant information in a compact yet intuitive manner. At each point of the game, the transactions that are provisionally formed – i.e., for which the total bid is 1 or above – are displayed as a thick solid line. The transactions that are currently not formed are displayed as a thin dotted line. A player can open a dialog box above each link to view the current state of play, with the value of this transaction to her and the gain she would obtain if the game were to stop at that moment. The reports are updated in real time based on current bids. This design, which is the result of careful research and development, presents the information in an instantaneously available and engaging format. For details, we refer to Appendix A.

2.2 Treatments

Each group plays 8 games in sequence. Each of these games is played under one of four treatments denoted T_A, T_B, T_C and T_D . These treatments introduce different trading protocols described in what follows.

No bargaining

All games start with a phase of multilateral sealed bids. Players have unlimited time to enter take-it-or-leave-it bids in the dialog boxes associated with each possible transactions. Once all players stop bidding, the game ends. If the total bid for a transaction is positive, that transaction is implemented. Players then observe their own gains for each transaction separately and in total. This can be seen as a trading floor in which buyers and sellers announce their secret reservation prices, and the auctioneer implements all transactions that generate a positive surplus by charging the prices announced.²⁴

²³Although we do not provide explicit information on the probability distribution of values, in the trial games we present participants with multiple examples of values drawn from the same distribution as in the main games.

²⁴This mechanism directly compares to most centralized matching algorithms (e.g. speed dating, UCAS) where a match occurs only if both sides make an offer. In our setting, a match is made if the two offers are

In Treatment T_A , this first phase marks the end of the game. Thus, under T_A the game only has one round, which we name round 0. In other treatments, the game continues with one or more rounds of covert bargaining. During round 0, players do not yet know which treatment they are in – i.e., the game may stop at the end of round 0 or it may continue with further rounds of bargaining. This ensures that players behavior in round 0 is comparable across treatments.

In a setting of private information, these sealed bids are a formidable challenge for subjects. In Appendix B we review the theoretical framework and derive the relevant efficiency bounds, showing that many profitable transactions do not occur (the estimated probability of trade is $\sim 29\%$). We use it as a benchmark to judge the efficiency gains that can be achieved with covert bargaining. It is also a natural way of launching the covert bargaining process, to which we now turn.

Sequential bargaining

The covert bargaining mechanism we propose could be applied in different ways. Since we cannot rely on previous experimental evidence, we focus on three protocols derived from our conjectures about the behavioral response of players.

In Treatments T_B and T_C , the game continues with a bargaining phase where all the action is focused on one transaction at a time (sequential bargaining). In these two treatments the bargaining is divided into multiple rounds, themselves divided into turns. In a given round, a turn is devoted to bidding on one specific transaction.²⁵ The sequence of play is as follows: when the turn begins, the transaction currently up for bidding is highlighted. Then the bargaining floor opens and all players can place, revise, or drop a bid on this specific transaction.²⁶ We place no limit on the number of bids: as long as bids change, the turn remains open. If there has been no change in *all* bids for a given amount of time, the turn ends.²⁷ Only at the end of the turn does the computer display whether the transaction is provisionally formed or not.²⁸ In the next turn, the game then moves to the next transaction. When all transactions are visited, the round is complete and another round begins.

compatible, and the pricing mechanism ensures that bids anonymity is preserved.

²⁵There are as many turns in a round as there are transactions in the game. The order of turns, i.e., the order in which transactions are auctioned, varies randomly across rounds.

²⁶The bids placed in round 0 serve as start-up bids for the first round of bargaining; in subsequent rounds, the start-up bids are the bids placed in the previous round.

²⁷The wait time is 10 seconds in the first round and 5 seconds thereafter. The remaining time is not displayed in the screen because it was found to be a distraction during the pilot. But the color scheme of the screen changes to gray to signal the approaching end of a turn.

²⁸If the sum becomes positive, the link is activated and it turns from a dotted to a solid line. If the transaction was formed in a previous round but the sum of the bids subsequently falls to 0 or below, the link is de-activated, i.e., it turns from a solid to a dotted line.

The two versions of the sequential bargaining protocol differ for their the stopping rule. Treatment T_B has a stopping rule based on bids, and the game ends in one of two ways: either because, at the end of the round, there has been no change in *bids* relative to the preceding round – what we call a *natural end*; or because the maximum number of rounds has been reached.²⁹ At the end of the game, the last bids are retained, final total bids calculated, and players are informed of their gains. Treatment T_C is nearly identical to T_B except that the stopping rule is based on *transactions* rather than bids. In this treatment the game reaches its natural end if there has been no change in provisionally formed *transactions* from one round to the next – even if there were changes in bids.

The rationale for these two protocols goes as follows. Both sequential treatments corral subjects into a systematic sequence of bargaining sub-games on one transaction at a time. This is cognitively easy, but it could potentially be perceived as repetitive and boring by players. Protocol T_C is intended to limit bidding wars and curtail the time that players take to experiment with frivolous bids and attrition strategies. But it is also possible that the faster stopping rule in T_C reduces the scope for discovery.³⁰

Simultaneous bargaining

In contrast to treatments T_B and T_C , treatment T_D implements an unconstrained bargaining floor in which players can update bids simultaneously on all transactions, with no specific sequencing imposed. Like in Treatments T_B and T_C , this treatment begins with bids from the sealed-bid phase. As the bargaining phase unfolds, each player can update any bid in any order. The silent auctioneer updates the information on transactions and gains in real time. This allows players to see which transactions would be formed if the game were to end at that point. The game stops when there has been no change in *all* bids for a set amount of time.³¹ Thus, under T_D , games have two rounds: the sealed-bid phase of round 0; and one long round of simultaneous bargaining.

²⁹In bargaining games with a predetermined end time, it is common to observe a bunching of bids just before the deadline. To mitigate this problem, we randomize the number of rounds at the end of T_B and T_C . If a game has not ended naturally by the end of the 6th round, the game is forcibly ended with a 50% chance in round 6 and a 50% chance in round 7. All remaining games stop by the end of the 8th round. Thus, the upper bound on the number of rounds varies from 6 to 8. Players are informed about the ending rules and their probabilities.

³⁰To illustrate, imagine that the dominant heuristic is for players to gradually increase their offer until the transaction is formed. Under T_B the game continues until they stop increasing their bid but in T_C the game ends if new bids fail to change transactions. Hence T_C may penalize bidding-up strategies that are too gradual.

³¹At the beginning of the phase, the wait time is 20 seconds since the last bid placed by any player. This gives subjects enough time to absorb the information coming from round 0. The wait time is subsequently reduced to 10 seconds. The timer is not openly displayed, but colors fade away to mark the passage of time from the moment the last bid was entered.

In the protocol T_D subjects follow multiple transactions simultaneously, which in principle increases cognitive load with respect to the sequential protocols. However, this could speed up the game, provided that subjects are able to maintain concentration and motivation. Furthermore, information discovery about profitable transactions is facilitated when bargaining is simultaneous: in T_D formed transactions are displayed on the screen as soon as the sum of bids turns positive; in contrast, in T_B and T_C subjects have to wait until the end of a round. The comparative performance of these three protocols is an open empirical question that we explore in Section 3.

2.3 Implementation

The laboratory experiment was conducted in 2019 at the Paris School of Economics in France. We ran six experimental sessions with a total of 22 unique groups of 6 players playing 8 games. Group composition remains unchanged throughout a session, but the identity of players is reshuffled from one game to another.³²

Operationally, we have drawn 4 matrices for two-sided games and 4 for many-sided games.³³ Examples of these matrices are reported at the end of Appendix A.1. For each group, the 4 two-sided matrices get randomly paired with treatments T_A to T_D , and the same goes for the 4 many-sided matrices. Thus, each group plays 8 matrices in total and 2 matrices (one two-sided and one many-sided) for each treatment, in a random order which is group-specific.³⁴ As a result, we observe $22 \cdot 8 = 176$ unique games equally distributed among the four treatments (44 games per treatment).

Sessions unfold as follows. After reading the instructions, players run three trial games to accustom themselves to the interface.³⁵ After the trial games, players answer a quiz to test their understanding of the instructions. The quiz is corrected immediately afterwards on the

³²Each player sees on the screen a circle with himself at the bottom (“ME” – followed by his current letter identifier) and the other 5 players around the circle, each identified with a letter. While ME stays always at the bottom, the other players’ letters are visualized in clockwise order (i.e., C will be always between B and D). We reshuffle letter identifiers at the end of each game. To illustrate, a player may see himself as “ME (D)” in one game and “ME (A)” in another, while all other identifiers have been similarly been reshuffled. This is done to minimize spillovers from one game to the next.

³³These values matrices were randomly drawn according to the rules detailed in Section 2.1, with three additional constraints: 1) each subject is an internal player for at least one transaction; 2) $v_{ij} = \sum_k v_{ij}^k \neq 0$ (i.e. the total value of a link across all players is either strictly positive or strictly negative) to mark a clear efficiency criteria; 3) $v_{ij}^k \neq 1$ (i.e. players are not indifferent between not making an offer or making a minimal offer of 1). As a result of randomization, the number of profitable transactions varies between 5 and 8 out of 10 across the retained matrices.

³⁴For example, group 1 may play treatment T_A with matrix 1 in the third game while group 2 plays the same matrix 1 with treatment T_C in the seventh game.

³⁵Players are informed that during the trials they will play treatments B, C and D, in that order. The value matrices we used for trial games were generated for this scope, and display 7 links instead of 10. The rules to draw the trial matrices as well as all other rules are the same as in the main games.

board. Once this is completed, players run a social value orientation task (Murphy et al. 2011) and then proceed to the main part of the experiment. At the end of the games, players complete a questionnaire with socio-demographic information and comprehension feedback, and they receive their payment. To determine subjects' earnings, we randomly draw 2 games for each group and players receive the monetary equivalent of their gain at the end of these games.³⁶ The average earnings are 24.5 euros for about 2 hours in the laboratory.

3 Efficiency and surplus

3.1 Ex-post efficiency

Since there are 176 unique games with 10 transactions each, we observe the outcome for $N = 1760$ transactions, divided equally across the four treatments (440 transactions per treatment). Out of these, 550 transactions are defined as non-profitable because their total value over all players is strictly negative, i.e. $v_{ij} = \sum_k v_{ij}^k < 0$. The remaining 1210 transactions are called profitable since their total value is strictly positive, i.e. $v_{ij} = \sum_k v_{ij}^k > 0$.³⁷ The exact number of profitable transactions slightly differs by treatment (from 276 in T_D to 322 in T_A) because value matrices are randomly assigned to games.

Only a negligible number of non-profitable transactions are formed by the game's end (8 out of 550). This provides reassurance that the majority of players understand the rules of the game. Thus our main analysis of efficiency and surplus, summarized by Figure 1, focuses on profitable transactions only.³⁸

The upper part of Figure 1 illustrates our results about efficiency. Our first efficiency measure is the percentage of profitable transactions that are formed at game's end.³⁹ Panel A shows this percentage for each of the four treatments. We see that only a small fraction of profitable transactions (27%) are formed in the sealed-bid treatment T_A – a proportion that is remarkably close to the theoretical probability of trade of 29% derived in Appendix B. The percentage of profitable transactions formed increases to 60% or more when covert bargaining is possible, i.e., in treatments T_B , T_C and T_D . Theory suggests that private information has a high cost in terms of efficiency (e.g., Myerson and Satterthwaite 1983; Sobel and Takahashi

³⁶If some players had incurred negative earnings, we would have subtracted these losses from their show-up fee. We had no such cases.

³⁷As explained in Footnote 33, we have ruled out transactions with $v_{ij} = 0$.

³⁸We present results in visual form for ease of interpretation. Similar findings (available upon request) are obtained if we use regression analysis to control for order effects (to capture possible learning across games) and session fixed effects.

³⁹5 profitable transactions out of 1210 generate negative gains for players ($\sum_k g_{ij,T}^k < 0$). We do not regard these 5 observations as efficient, to facilitate a clear-cut decomposition into players and auctioneer surplus (see below).

1983; Williams 1987; Vincent 1989; Ausubel et al. 2002). This prediction has so far been confirmed by the empirical evidence available on bargaining.⁴⁰ This is also what we find in our experimental setting: a large proportion of profitable deals never occur. However, covert bargaining mitigates the private information penalty to a large extent.

Panel B of Figure 1 depicts efficiency in terms of the percentage of the total feasible value that is achieved. More precisely, the height of each bar in panel B represents the sum of v_{ij} for completed transactions, divided by the sum of v_{ij} for all profitable transactions in the treatment group under consideration. It provides a finer picture of the magnitude of the surplus that is generated by our trading games. Results from panel B are in line with the ones in Panel A. In particular, it suggest that the sealed bids in T_A achieve on average 37% of the total feasible value, while this statistics is in the 70-80% range for the covert bargaining treatments.

3.2 The division of surplus

Efficiency can in turn be split into two components, that we plot in panels C and D. Panel C depicts the *surplus share for players*: for a given transaction ij , this is computed as the share of its total value earned by the players by the end of game, i.e., $g_{ij,T}/v_{ij}$ where $g_{ij,T} = \sum_k g_{ij,T}^k$ if the transaction occurs and zero otherwise. Conversely, the *surplus share for the auctioneer* in panel D is the share of total value appropriated by the auctioneer, i.e., $(v_{ij} - g_{ij,T})/v_{ij}$ if the transaction occurs and zero otherwise.⁴¹ Since the surplus shares for players and the auctioneer add up to one if the transaction occurs, the quantities in panels C and D sum up by construction to the efficiency measure in Panel A.

When we look at how the surplus is split between players and the auctioneer, we see that most of the efficiency gains stemming from bargaining go to players. In contrast, the auctioneer's surplus share remains stable and low across all treatments. Recall that, in our experiment, the silent auctioneer collects at least a minimal fee of 1 unit for each completed transaction. The yellow triangles in panel D represent the lower bound of auctioneer surplus if all realized transactions yielded the minimal fee of 1.⁴² The comparison between the

⁴⁰Using observational data, Larsen (2021) estimates an efficiency loss of 17-23% for two-sided uncertainty, while Freyberger and Larsen (2025) find that the fraction of failed negotiations on Ebay is 37%. Ambrus et al. (2018) find an efficiency loss of 14% in ransom negotiations. In an experimental setting, Goeree and Lindsay (2020) show that bargaining treatments with complete information and communication produce higher efficiency.

⁴¹Our metric to quantify total efficiency and surplus shares parallels the bargaining power measure reported by Larsen and Zhang (2021).

⁴²For a given treatment, the yellow triangle represents the average value of $1/v_{ij}$ over the transactions formed by the end of the game. This differs across treatments because value matrices are randomly allocated, so that the average value of profitable transactions differs.

surplus share achieved by the auctioneer and its lower bound shows that subjects were surprisingly efficient at not leaving money on the table.⁴³ This result is not trivial: under covert bargaining, the division of gains from trade is an open theoretical question, and the auctioneer could have acted as an additional source of friction. Our results show that this is not the case, and the secrecy of bids did not result in an advantage for the auctioneer.

Finally, in the lower part of Figure 1 (panels E and F) we decompose panel C by plotting separately the surplus share of buyers and sellers. Since buyers and sellers are in a fully-symmetrical position, the division of the surplus between the two groups is *a priori* undetermined.⁴⁴ Our results show that the quasi-totality of the surplus stemming from bargaining is appropriated by buyers, while the average surplus share of sellers is nearly zero. This reconfirms a previously documented earning gap in favor of buyers (Smith and Williams 1982; Sherstyuk et al. 2021). Digging into this result is the object of Section 4.

3.3 Comparison of bargaining protocols

Overall, our results suggest that all three versions of covert bargaining deliver comparable results in terms of efficiency and division of surplus. This is a notable empirical findings *per se*, as we could not rely on previous empirical evidence about it.

When we look at the results from Figure 1, we see that the two sequential treatments deliver nearly identical outcomes throughout. Our data also indicates that, unsurprisingly, treatment T_C cuts game time: in treatment T_B the average number of rounds is 6.4, while in treatment T_C the average number of rounds is 3.9. We conclude from this that the faster stopping rule in T_C speeds up bargaining without reducing the scope for discovery. This implies that T_C may be preferable to T_B in many contexts.

Our results also show that the simultaneous bargaining protocol T_D delivers an outcome that is at least as good as the sequential protocols. Players seem able to juggle a moderate number of simultaneous negotiations at once, and the benefit of T_D in terms of information discovery more than compensates for the increase in cognitive load.⁴⁵ Overall, all three protocols appear to foster beneficial exchanges and allows players to extract most of the surplus from trade. Their application to real-life trade platforms depends on the size and scope of the market, an issue we discuss in Section 6.

⁴³For corroborating evidence, 51% of deals in the lab occur with $b_{ij,T} = 1$, which is the minimum. Another 22% have $b_{ij,T} = 2$.

⁴⁴As we know from Muthoo (1990) if buyer and seller are allowed to continue making alternating offers and counter-offers without commitment, any division of the gains from trade can be sustained as a sub-game perfect equilibrium of a modified bargaining model (Rubinstein 1982). This theoretical indeterminacy motivates our empirical exploration, as discussed in Appendix B.

⁴⁵As a result of the value randomization, players hold non-zero values for *grossomodo* 4 transactions on average (3 in one-sided matrices).

3.4 Evidence on sub-samples of transactions

Figures 2 and 3 reproduce the analysis of Figure 1 separately for low- and high-value profitable transactions (below and above the median value of v_{ij}), respectively. Our main findings hold in both cases: total efficiency increases considerably when covert bargaining is possible, and most of the extra surplus stemming from it goes to buyers – while the auctioneer and the sellers do poorly in comparison. This pattern is accentuated in high-value transactions because the efficiency gains stemming from bargaining increase dramatically. In particular, we find that, in the three bargaining treatments, the percentage of profitable transactions formed (panel A) is around 40% for low-value transactions compared to 80% or more for high-value transactions. This is consistent with experimental evidence showing that deal rates increase with the size of the surplus (e.g., Camerer et al. 2019 in a context of one-sided private information). Figure 3 confirms that the quasi-totality of the additional surplus generated for high-value transactions is still appropriated by buyers, while the surplus share going to sellers and auctioneer remains consistently low.

In Figures A.6 and A.7 of Appendix C, we repeat the analysis separately for transactions where no more than two players hold non-zero values (two-sided transactions) *vs.* transactions where three or more players hold non-zero values (many-sided transactions). Similar findings are again obtained.

3.5 Renegotiation and efficiency

In the covert bargaining treatments, bidding comes without commitment and players have the opportunity to renegotiate deals. To explore this dimension, in Appendix Figure A.8 we compare the moment of activation (that is, the moment when players observe that a transaction is provisionally formed for the first time) and the end of the game.⁴⁶ The upper part of the figure plots efficiency as the percentage of profitable transactions formed at activation and game-end respectively (as in panel A of Figure 1). The central and lower parts of the figure decompose efficiency into the surplus share of players and auctioneer respectively (as in panels C and D of Figure 1).

Since deals can be broken during the renegotiation process, efficiency can only decrease from activation to game’s end in the three bargaining treatments. Results from Figure A.8 indicate that there is indeed a decrease in efficiency after discovery, especially in treatments T_B and T_D . However, this decrease is modest in magnitude and it is mostly associated with a large drop – in relative terms – of the auctioneer surplus. This reconfirms the finding

⁴⁶In T_D players see in real time whether the transaction is activated or not – while for sequential bargaining treatments, the moment of activation coincides with the end of the turn.

that bargaining without commitment allows players to extract surplus at the auctioneer's expenses. Conversely, the decrease in the surplus share of players is small in relative terms but nonetheless present in all bargaining treatments. This raises the possibility that excessive bargaining may introduce frictions, a point that we revisit in Section 5.

4 Bids and Roles

The efficiency of market protocols is ultimately rooted into the bidding strategies adopted by players. We now document the differences in bidding pattern depending on the role ascribed to the player for a given transaction – i.e., buyer or seller.

4.1 Descriptive statistics

Table ?? presents descriptive statistics on the bidding behavior of players. Our unit of observation is player k on transaction ij in game g . Since there are 10 transactions per game, we have $N = 132 * 8 * 10 = 10,560$ observations. We report statistics for the entire sample and then we split observations according to the role ascribed to the player for the transaction: *sellers* (with $v_{ij}^k < 0$, $N = 1,694$); *buyers* (with $v_{ij}^k > 0$, $N = 3,014$); and players with $v_{ij}^k = 0$ ($N = 5,852$).⁴⁷ Column (1) reports the mean of value v_{ij}^k , Column (2) reports the number of players who are active (that is, place at least one bid), and Column (3) reports the number of bids placed.⁴⁸ Columns (4) and (5) provide additional information about the bids placed – whenever player k has placed multiple bids over transaction ij in game g , we retain the last one. Column (4) counts the losing bids, that is, bids where $b_{ij,t}^k > v_{ij}^k$, meaning that they would result in a loss for the player if the game were to stop at that moment and the transaction was formed. Column (5) reports the number of bids equal to $v_{ij}^k - 1$: these bids are associated with a minimal positive gain of 1 unit if the transaction occurs.

Column (1) reassures the reader that values are randomized in a comparable manner across roles and treatments – albeit with values that are positive for buyers and negative for

⁴⁷Buyers are more numerous than sellers because four games allow for many-sided transactions. The relative frequency of buyers and sellers is not balanced across treatments because the matrices are randomly allocated.

⁴⁸To identify (unique) bids we proceed as follows. We exclude first bids equal to zero: since 0 is the default initial value proposed in the dialog box, first bids that are equal to zero do not represent an actual bid. We also exclude all repeated consecutive *identical* bids from the same player, since they are redundant. Then, for all treatments we only retain the last bid that a player places in the sealed-bid phase (round $r = 0$) and, for the sequential treatments T_B and T_C , we only retain the last bid that a player places within a given round $r \geq 1$. This is because in these cases only the last bids affect transactions, and outcomes are only revealed at the end of a round. In the simultaneous bargaining treatment T_D , we retain all the bids placed in the bargaining round.

sellers.⁴⁹ Column (2) shows that, as expected, players with zero value are rarely active: in only 499 (8.5%) of 5852 player-transaction pairs with zero value a bid was placed.⁵⁰ Buyers appear to be less active than sellers: 2576 or 85.5% of buyer-transaction pairs, versus 1636 or 96.6% for sellers, place at least one bid.

Overall, 13,331 bids were placed across all players and treatments combined (Column 3). Interestingly, we see that bids are less frequently placed by sellers than buyers: we observe 3704 bids by sellers, which gives 2.26 bids per active seller on average, compared to 8325 bids by buyers, which gives 3.23 bids per active buyer on average. Together with the evidence from Column (2), this indicates that sellers are more likely to place at least one bid, but less likely to bargain when compared to buyers.⁵¹

When we examine whether subjects occasionally place losing bids at the end of the game (Column 4), we see that only 301 bids (2.8% of all bids) are losing bids, placed either by sellers (145 bids) or by players with zero value (156 bids). No losing bid was placed by buyers. Furthermore, not all these bids turn into a loss because most of the corresponding transactions are not formed. This is reassuring since it suggests that players have little difficulty in understanding the game.

Column (5) reports the number of last bids seeking a minimal gain of 1 unit, that is, equal to $v_{ij}^k - 1$.⁵² These bids more frequently come from sellers (761, or 47% of active sellers) than buyers (423, or 16% of active buyers). This means that, by the end of the game, nearly half of the sellers settle on the most conservative reservation price and ask for a minimal gain of 1 unit, while only one sixth of sellers do so.

Summarizing, the results reported in Table ?? suggest that there are meaningful differences in the bidding behavior of players depending on their role. Sellers appear to bargain less (i.e., they place fewer bids) and less hard (i.e., they frequently settle for a minimal bid of $v_{ij}^k - 1$) than buyers. Given our experimental design, these differences cannot be driven by the distribution of values, which are randomly selected – nor can they be imputed to individual heterogeneity since all players simultaneously occupy both roles.

⁴⁹The small gap between mean values of sellers and buyers is due to the fact that we set $v_{ij}^k \neq 1$ to avoid indifference between placing a minimal offer or not (see Footnote 33).

⁵⁰This behavior is typically transitory, and not necessarily ‘irrational’. If we look at the last bid of these 499 active players with $v_{ij}^k = 0$, we remark that 156 of them are positive, 153 are negative, and 190 are zero (i.e., the player had previously placed a non-zero bid but she reverts to zero). Negative bids are possible because some internal players have $v_{ij}^k = 0$ but can place a request to consent to the transaction. Thus, negative and zero bids are compatible with monetary payoff maximization. Only a negligible minority of these players (156 over 5852, or 2.7%) place a bid which would result in a loss in case of deal.

⁵¹This is also confirmed by the regression analysis reported in Appendix Table A.2. Results show that, once we control for all potential confounding factors, sellers place consistently fewer bids than buyers.

⁵²This statistic is not informative for players with zero values, because by design the majority of their bids are losing bids (since they can only place positive bids unless $k = i$ or j).

4.2 Profit margins

The evidence from Section 4.1 introduces naturally a question about differential profit margins depending on player's role. To dig into the issue, we estimate the following model:

$$b_{ij,g}^k = \beta_0 + \beta_1 v_{ij}^{k-} + \beta_2 v_{ij}^{k+} + \beta_3 S + \beta_4 T_g + \lambda_s + \varepsilon_{ij,g}^k \quad (1)$$

where $b_{ij,g}^k$ represents a given bid (first or last) placed by player k on transaction ij in game g . The first regressor of interest v_{ij}^{k-} is the value of transaction ij if k is a seller, i.e., $v_{ij}^{k-} = v_{ij}^k$ if $v_{ij}^k < 0$; it is 0 otherwise. The second regressor of interest v_{ij}^{k+} is the value of ij if k is a buyer, i.e., $v_{ij}^{k+} = v_{ij}^k$ if $v_{ij}^k > 0$ and 0 otherwise. We control for game order S (from 1 to 8), treatment dummies T_g , and session fixed effects λ_s . Standard errors are wild-bootstrapped at the group level (Cameron et al. 2008), which is the highest level at which participants interact in the experiment.

Table ?? presents the regression results from Equation ???. Column (1) only includes initial bids placed on transaction ij during game g . Columns (2) and (3) focus on final bids, for all treatments and for bargaining treatments only, respectively.⁵³ For the two regressors of interest, the average of the requested profit margin implied by the bid is reported below the coefficient (in square brackets).

Results show that asking bids placed by sellers respond slightly more than 1 for 1 to negative values. For instance, in column (2) (last bids, all treatments combined), sellers demand on average 1.327 to sustain a negative value (cost) of -1 . If all these transactions were formed (which is not the case), these bids would yield a modest average profit margin of 32.7% for sellers. In contrast, buyers' bids respond much less than 1 for 1 to positive values: in Column (2) we see that buyers offer on average 0.302 for a positive value of $+1$. If all these transactions were formed, these bids would yield an average profit margin for buyers of $1 - 30.2\% = 69.8\%$, which is much higher than the corresponding figure for sellers.

These results show that players request considerably lower profit margins when they are in the position of sellers. This non-aggressive bidding behavior is consistent with the differential results on surplus (Table 1). In the economic literature, the behaviors of buyers and sellers have been compared through the lenses of market power, experience, information, and rationality (see Simonsohn and Ariely 2008 and Garratt et al. 2012 for recent evidence based on Ebay data). Our results are aligned with the overwhelming empirical evidence that sellers require a minimum selling price substantially higher than the maximum amount offered by buyers – the so-called willingness to accept (WTA) / willingness to pay (WTP)

⁵³The difference between columns (1) and (2) is driven by players placing multiple bids: if a player places zero or one bid on ij during game g , her dependent variable stays the same.

gap. In particular, from Kahneman et al. (1991) onward, the ratio WTA/WTP has often been estimated to be in the range 2-3 for ordinary private goods, which compares well with our estimates. We add to these findings by showing that this occurs not because the ask price is unreasonably high, but because the offer price is unreasonably low. The randomization of subjects across roles and values in our experiment confirm that this effect is entirely behavioral: it is not due to self-selection into roles or to a particular understanding of the game’s rules. Furthermore, it is only partially mitigated by bargaining dynamics: the gap between the implicit margins requested by sellers and buyers for first vs. last bids decreases, but it remains wide. This empirical regularity, which is often seen through the lenses of the endowment effect, arises in our setting where values are not attached to objects or actions. Based on this, we conjecture that the behavior we observe is likely to be rooted in common real-life practices that make subjects more socially accustomed to haggle when buying than when selling.⁵⁴ Thus, these findings are reconcilable with the experience hypothesis (List 2003), in the sense that the endowment effect is tied to market inexperience. Other explanations are conceivable as well: for instance, given the way the trading process is framed, players may see negative bids as compensations or protection against losses for links enforced by others, while they may see positive bids as an occasion for profit.

4.3 Internal vs. external buyers

Finally, we investigate whether buyers behave differently when they hold an internal vs. external position in the transaction. To recall, in our experiment, transactions are visually represented as links, which may create a distinction between the two internal players i, j represented at the extremities of the link, and external players $k \neq i, j$ who are not. All internal players can place positive or negative bids, regardless of whether they are buyers or sellers. In the many-sided games, some external players are also buyers, i.e., they have $v_{ij}^k > 0$. These external buyers cannot place negative bids – to ensure that they do not prevent others from trading. But otherwise they have the same payoff rule and incentives as internal buyers. It follows that differences in bidding behavior between internal and external buyers must arise either due to the constraint on negative bids, or due to the framing of the experiment. In particular, participants who are visually represented as external contributors may bid differently if they feel less engaged into or less generous for transactions among third parties.

⁵⁴This is aligned with the results of Ikica et al. (2023), who find that buyers bargain more aggressively than sellers in the early stages of double auctions (even in treatments with no information feedback). It is also compatible with the argument by Smith and Williams (1982) that most subjects have more experience with the role of a buyer than that of a seller.

In Appendix Tables A.3 and A.4, we compare the number and magnitude of bids placed by internal and external buyers ($N = 1,694$ and $1,562$ respectively). Table A.3 indicates that the two types of buyers displays comparable patterns of behavior regarding the number of bids placed. But taking all bids into accounts (Columns 1 and 2), the results presented in Table A.4 suggests that the implicit margin requested by external buyers is more modest than the one requested by internal buyers – 60.8% compared to 77.8% across all last bids – although if it remains twice as high as the average margin requested by sellers. However, our data suggests that this difference arises primarily because a small minority of internal buyers place negative bids (241 observations).⁵⁵ Once we remove these observations in Column 4, the implicit margins requested by external and internal buyers get closer: 54.4% for externals compared to 62.3% for internals. Thus, the framing of the experiment in terms of internal vs. external buyers seem to have little effect on the observed behavior.

Finally, our experiment requires players to engage in multiple transactions within a given game, which may increase complexity. In Table A.5, we examine whether bidding behavior is affected by the number of transactions subjects are simultaneously involved in. In Panel A we regress three outcomes representing the behavior of player k for link ij in game g : the number of bids $\# b_{ij,g}^k$, the request on last bid $v_{ij}^k - b_{ij,T}^k$ and a dummy taking value one if the request on last bid is equal to one unit, i.e. $v_{ij}^k - b_{ij,T}^k = 1$ ('minimal request'). The regressors of interest in columns 1 to 9 are the number of transactions for player k in game g with $v_{ij}^k < 0$ and $v_{ij}^k > 0$ respectively (excluding transaction ij itself). We split the sample into sellers (columns 1-3), internal buyers (columns 4-6) and external buyers (columns 7-9) respectively. The results show that the number of potential trades does not affect the behavior of sellers, whereas it does affect the behavior of both internal and external buyers, who bid less aggressively (i.e., request a smaller potential gain and are more likely to place a minimal request). The similar order of magnitude of the estimates for negative- versus positive-value trades suggests that subjects are influenced by the total number of ongoing trades rather than their type. This may reflect either an attention constraint or lower relative stakes, in the sense that subjects already expect gains from other trades and therefore reduce their effort on a given one.

In the second part of Table A.5 (Panel B), we restrict attention to the sub-sample of buyers and augment the specification of Panel A with a variable capturing the behavior of player k as a seller, namely the share of bids with a minimal request of 1 (i.e. such that $v_{ij}^k < 0$ and $v_{ij}^k - b_{ij,T}^k = 1$) placed across all games. Results suggest that players who place

⁵⁵In a setting of covert bargaining, this 'greedy' strategy of buyers (i.e. exploiting the poor informational setting and the positional advantage to request a compensation) may result in a large gains if the transaction is sponsored by others. We had allowed this feature to see how many players adopt this strategy, and remark that their number is small.

minimal-request bids as sellers are also more likely to do so as buyers. In principle, we could have found that some sellers adopt a minimal-request strategy to redirect their bargaining effort toward bidding more aggressively as buyers, but this does not appear to be the case. Instead, our findings suggest that certain subjects consistently rely on simple negotiation strategies throughout the game to navigate the complex experimental environment.

5 The dynamics of simultaneous bargaining

We end our analysis with an examination of the dynamics of the covert bargaining process. In what follows we focus on T_D because the simultaneous bargaining protocol provides ideal data for this purpose. In Online Appendix C we show that all the findings discussed below extend to sequential bargaining protocols as well.

5.1 Bidding up

Experimental evidence suggests that human subjects are able to resolve difficult strategic games by relying on simple heuristics. Little is known, however, about multilateral bargaining situations with little information and structure. Since we allow for renegotiation at no cost, the pool of bidding strategies that the players could adopt is potentially large. Two simple heuristics nonetheless come to mind as likely guides to bargaining without commitment: bidding down and bidding up.

In a bidding-down strategy, players initiate bargaining with a high bid (e.g., $v_{ij}^k - 1$) and gradually lower it. If adopted by all players, such strategy is an efficient way to reveal all profitable transactions upfront, and players can subsequently decrease their bids to capture a larger share of the surplus from the auctioneer. Conversely, players can follow a bidding-up strategy whereby they initially make a low bid and gradually raise it. A gradual bidding-up strategy has been documented by Bakus et al. (2020) on observational data, but theoretical support remains scarce.⁵⁶ If pursued by most players for a sufficiently long time, this bidding-up process should eventually reveal all profitable transactions while minimizing the surplus that goes to the auctioneer. If, however, some players lose patience or diverge from this strategy at some point, some profitable transactions may go unimplemented.

First, we focus on the individual overall strategy across games, to pinpoint subjects who use simple monotonic (or quasi-monotonic) strategies. Let us define an *increasing strategy* as a strategy in which all bids, or all bids but one, increase with respect to the previous one.

⁵⁶Following Rubinstein (1982), most theoretical analyses of strategic bargaining games predict immediate agreement. Notable exceptions are Compte and Jehiel (2004) and Keniston et al. (2021).

Similarly, in a *decreasing strategy* all bids, or all bids but one, decrease with respect to the previous one. We remark that players' behavior differs heavily before and after activation, which we define as the moment when the transaction is provisionally formed for the first time (and thus shown to be profitable to the participants). The share of players using a decreasing strategy remains of the same order of magnitude — 6% and 7% before and after activation. However, the share of players who follow an increasing strategy is large before activation (26%) and drops to 8% after activation. This provides initial evidence of a generalized bidding-up strategy in the initial phase of the game, while players explore to identify mutually beneficial exchanges—a point that we examine in detail in what follows.

During the simultaneous bargaining round of treatment T_D we observe 4,995 bids placed over 425 transactions.⁵⁷ We define a *bid run* as the ordered sequence of bids placed by all players on a given transaction ij , and we divide the sample of bids into four approximately equal quantiles based on the length of the bid run they belong to.⁵⁸ For each quantile, we estimate a model of the form:

$$b_{ij,t_{ij}} = \alpha + \sum \beta \lambda_{t_{ij}} + \lambda_{ij*m} + \varepsilon_{ij,t_{ij}} \quad (2)$$

where the *tick* variable t_{ij} represents the order in which a bid is placed on transaction ij by different players within the run, and $b_{ij,t_{ij}}$ is the total of all outstanding bids placed by all players on that transaction at tick time t_{ij} .⁵⁹ This outcome variable is provisional since it refers to a given point t_{ij} along the bargaining sequence and need not correspond to the end of the game, nor to a transaction that is provisionally formed. The regression includes fixed effects $\lambda_{t_{ij}}$ for each value of the tick variable, and transaction-per-matrix fixed effects λ_{ij*m} ($10 * 8 = 80$ effects) to control for potential confounding effects correlated with matrix structure. Estimating this separately per quantile corrects for the possible correlation between the length of a bid run and the pace of increase in bids.⁶⁰

⁵⁷Overall we observe 6002 bids over 440 transactions for T_D , but we discard all bids placed in the sealed-bid phase since they (cannot be ranked temporally and thus) do not help shedding light on the bidding dynamics. For 15 transactions out of 440, no bid was placed in Round 1.

⁵⁸The first quantile includes bids which belong to bid runs of length 1-9; these account for 23% of all bid observations and 58% of bid runs observations (i.e., out of the 425 unique bid runs that are still open in Round 1, 245 of them have less than 10 bids). The second quantile includes bids belonging to bid runs of length 10-19: this accounts for 30% of bid observations, and 26% of bid run observations (111 bid runs). The third quantile includes bids belonging to bid runs of length 20-34; these account for 24% of bid observations and 11% of bid runs (48 bid runs). Finally, the last quantile includes bid runs of length 35+; these account for 23% of bid observations and 5% of bid runs observations (21 bid runs).

⁵⁹The tick variable t_{ij} works like a time identifier in panel data. For example, if we observe 5 bids placed on transaction ij by 3 different players (3 - 5 - 4 - 3 - 4) in that order, t_{ij} would take values 1 to 5 to indicate the order in which these bids were placed.

⁶⁰To understand the issue, imagine that all subjects follow a bidding-up strategy until they reach a deal, and that players differ in the speed with which they increase their bids. In this case, bid runs that end quickly

Our estimate of interest is how the predicted values of the total outstanding bid $b_{ij,t_{ij}}$ evolves over bargaining time t_{ij} , which we plot in Figure ???. Fitted values are negative throughout because many (profitable and non-profitable) transactions never occur, and most deals are sealed by a small margin. With this caveat in mind, our estimates validate the hypothesis that total bids increase over time, conditional on the total duration of the bid run.

We also note an increase in noise in the upper tail of the fourth graph. It comes from bids in the upper 5% of bid run length (more than 35 ticks). This indicates that some players lengthen bid runs by diverging from a gradual bidding-up strategy, which causes the observed non-monotonicity in bidding dynamics. In the next subsection we shed light on this issue by examining how bidding dynamics evolve during a bid run.

5.2 Discovery and appropriation

Our experimental design of bargaining without commitment puts us in the unique position of being able to observe how players behave after a provisional agreement is obtained.⁶¹ We now provide evidence that, after a transaction is revealed to be profitable, some players reduce their bids in an effort to increase their gains.

To show this, we split the bids placed during the simultaneous bargaining round of treatment T_D into two groups: those placed before and after activation, which we define as the moment when the transaction is provisionally formed for the first time.⁶² Column (1) of Table ?? reports these summary statistics for the entire sample of bids (4,995 bids over 425 transactions). As done in Section 5.1, Columns (2) to (5) breaks this sample down according to the length of the bid run. Panel A reports statistics on the corresponding transactions: number; mean value; and percentage of transactions implemented at the end of the game.⁶³ Panel B reports statistics on the bids placed before activation (number of bids, average number of bids per transaction) and classifies them according to the history of play of player k on transaction ij . In particular, we split the bids into three shares: first-time bids; increasing bids (that is, bids that represent an increase relative to the previous bid);

are those with larger/faster increases in bids, while those that take longer must have smaller increases in bids. It follows that pooling observations over all bid runs of different lengths yields β coefficients that are higher for short runs but smaller for long runs. By estimating the model separately for different run lengths, we can estimate more precisely whether β s increase monotonically within runs.

⁶¹This sets us apart from most settings in which bargaining stops as soon as an agreement is obtained (e.g. Bakus et al. 2020).

⁶²Recall that, in treatment T_D , players can see the activation of links in real time.

⁶³The sample also includes non-profitable transactions, which explains the difference between the deal rate reported here and the efficiency numbers reported in Figure 1.

and decreasing bids (that is, bids that represent a decrease relative to the previous bid).⁶⁴ We also report the percentage of increasing bids that go up by 1 unit (a sub-group of all increasing bids) and the share of decreasing bids that go down by 1 unit (a sub-group of all decreasing bids). The same statistics are reported in Panel C for bids placed after activation.

From Panel A we see that long bid runs are associated with higher value transactions: the average total value in Q1 is 4.60 compared to 7.52 in Q4. They are also less successful in general: the average probability of terminating the run with a deal is remarkably stable (i.e., between 46% and 48%) in the first three quartiles but drops down to 29% in the fourth quartile.

Turning to Panel B, we note that the overwhelming majority of bids placed before activation are increasing bids – most of the time by a single unit. For instance out of all 1947 bids placed before activation (Column 1, all bid run lengths combined): 8% are first-time bids; 68% represent an increase relative to the previous bid; and only 24% represent a decrease relative to the previous bid. Interestingly, the majority of all bids (70%) represent a 1-unit change, up or down: 55% of all bids represent a 1-unit increase and 15% a 1-unit decrease, relative to the previous bid. These patterns remain qualitatively unchanged if we split bids according to the length of the run (Q1 to Q4). Overall, bids placed before activation seem to corresponds to a discovery phase during which subjects play a cautious bidding-up strategy to reveal profitable transactions.

Panel C examines the bids placed after first activation. In longer bid runs, there are disproportionately more bids placed after first activation. In Q1, the number of bids per link placed before and after activation is nearly the same – 2.57 and 2.15, respectively. In contrast, the relative frequency of bids placed after first activation explodes for higher quartiles: 8.44 versus 17.04 in Q3; and 10.76 versus 43 in Q4. This suggests that long bid runs arise when (some) players haggle hard after activation, and this in turn may explain why long bid runs are less successful on average. Furthermore, we see that players haggle down in this phase of the bid run: after activation, decreasing bids become more frequent (45% of all bids are now decreasing compared to 24% before activation - see Column 1). As in panel B, this pattern remains remarkably stable across quartiles.

To sum up the evidence, we find that, when they bargain covertly and without commitment, most players engage in a slow bidding-up strategy to reveal profitable transactions. While this behavioral strategy increases efficiency with respect to the sealed-bid scenario, it is gradual. This can become a disadvantage when bargaining is slow, costly or constrained (in terms of attention or time). We also uncover a form of friction arising from the fact that

⁶⁴In Table ?? we restrict to bids placed in Round 1, but we use the information from round 0 to determine whether the bid is above/below the previous one.

some players bargain too hard in an attempt to appropriate more surplus. This is particularly noticeable in long bid runs: after activation, a small minority of players continue to revise their bid in an effort to appropriate (the unknown) surplus away from other players and the auctioneer.

6 Concluding remarks

With a few notable exceptions, much of the existing literature on the efficiency of decentralized markets is theory-based (e.g., Kirschsteiger et al. 2005; Kagel et al. 2010; Condorelli et al. 2017; Agranov and Elliot 2021). Bargaining games with little structure and private information have received less attention than auctions, possibly because they are less amenable to clear-cut theoretical predictions. However, the study of decentralized bargaining floors can provide valuable insights to theorists and policy makers (Karagözoğlu 2019). To help fill this gap, we design a new class of bargaining protocols where bids remain confidential (covert bargaining) and a small window of renegotiation is allowed.

Our experimental results suggest that covert bargaining doubles the efficiency margins with respect to take-it-or-leave-it bids, mostly to the profit of players (and in particular, buyers) rather than the silent auctioneer. This is tied to the fact that participants experiment through a cautious and gradual strategy of incremental bids – a strategy that is expedited in an unconstrained bargaining environment. Still, not all profitable transactions are realized, which confirms the private information penalty documented in previous theoretical and empirical studies. We are able to associate this penalty with two main behavioral drivers. First, buyers haggle hard as they try to extract an unreasonable share of surplus. This behavior stands in contrast with that of sellers, who ask for a more modest reservation price. Second, a small minority of participants continue bargaining after a transaction is provisionally formed, and this jeopardizes the likelihood of reaching a deal.

This paper offers an exploration of novel trading environments with covert bargaining features that are technologically practicable, and may become reality in the close future. As such, our findings have practical relevance for the design of online marketplaces. The covert bargaining features we propose could be implemented in real-life trading floors. By imposing a minimal amount of structure to keep all information private, they circumvent the theoretical and practical complications that arise when participants are asked to report preferences or reservation values.⁶⁵ These bargaining protocols have several advantages:

⁶⁵Preference revelation on large platforms may be logically impractical. Note also that a platform may choose to keep donors' identities private, not only the donation amounts. In the context of crowdfunding and charitable contributions, donors may prefer privacy—either because they do not want to signal their disposable income or because they do not want to reveal their association with a particular cause (e.g., a

they are versatile, they can scale up to operate in middle-sized or large markets, and they allow for an adjustable remuneration of the trading platform. As for the versatility, our experimental design was deliberately general but the platform/auctioneer could switch off some of the features. For instance, our protocols could regulate a ‘standard’ market where participants are either buyers or sellers and all transactions are bilateral, as well as a market of crowdfunding/ philanthropic giving (i.e. where the one seller fixes a monetary objective that is attained through bids/donations, whose amount is not revealed). Also, explicit consent could be requested from all parties involved. As for the scale-up, sequential protocols could be implemented in markets with many participants but few items for sale. Similarly, the simultaneous protocol is suited for large markets with many participants and many items, as long as each participant is only interested in a small number of items – which they could be asked to specify beforehand in order to be allowed to bid. Finally, all protocols allow for an adjustable remuneration of the trading platform.⁶⁶ A large positive threshold mechanically increases the provider’s margin, which incentivizes service provision. A negative threshold can be used to subsidize certain transactions, such as those generating positive externalities for society at large. Should a market for multilateral bargaining platforms develop, our results could provide useful insights for policy makers willing to regulate the industry – notably on which mechanisms are more efficient.

While our sealed-bid auction benchmark is admittedly simple — preserving full privacy but eliminating the discovery process inherent to bargaining — one could in principle design alternative benchmarks that relax secrecy while retaining a bargaining structure (in the spirit of alternating-offers by Rubinstein 1982), though it is important to note that secrecy is a core feature of our mechanism, as revealing how much surplus is left for the intermediary would itself affect behavior, much as in real-world market-making settings where gains are intentionally kept opaque.

Our design lays the groundwork for future extensions in two directions. First, in the current experiment, transactions are assumed independent and participants do not compete over them. This allows us to abstract from congestion and competition considerations, and it provides us a straightforward efficiency metric. But it only describes bargaining markets where the demand is (locally) thin. Yet, the bargaining protocols we propose are adapted for markets of non-independent goods, since agents may retract from sub-optimal early deals (Niederle and Roth 2009). We know from earlier work that laboratory subjects perform well in matching games of substitutes without transfers (e.g., Comola and Fafchamps 2018) and

political party or any organization linked to a sensitive trait such as religion or sexual orientation).

⁶⁶In our current design, the silent auctioneer fixes a threshold of total bids of 1 to implement a transaction. This minimum fee serves to remunerate the platform, and it can be set to any number as desired.

in bargaining games over bundles of goods (Bochet et al., 2023; Jackson et al., 2023). It should therefore be possible to extend covert bargaining to non-independent transactions, making it closer to the complex features of real-life markets.

Second, even though transactions are represented visually to subjects as a graph, our experiment does not examine the topology of the resulting network or its externalities, since there is no resale or transfer across nodes. Yet, our work opens a new avenue of inquiry into network formation games (e.g., Currarini and Morelli 2000; Baccara et al. 2012; Agranov et al. 2021). Link formation with transfers has a natural interpretation in terms of a buyer-seller network (e.g., Mutuswami and Winter 2002; Choi et al. 2017). Extending our framework to allow players to derive utility from indirect links would provide valuable information in strategic settings such as diffusion games, information flows, and competition.

References

Abdulkadiroglu, A., N. Agarwal, P. Pathak (2017), ‘The Welfare Effects of Coordinated Assignment: Evidence from the New York City High School Match,’ *American Economic Review*, 107(12): 3635-89.

Abraham, I., S. Athey, M. Babaioff, M. Grubb (2020), ‘Peaches, Lemons, and Cookies: Designing Auction Markets with Dispersed Information’, *Games and Economic Behavior*, 124(C): 454-477.

Agranov, M., C. Cotton, C. Tergiman (2020), ‘Persistence of power: Repeated multilateral bargaining with endogenous agenda setting authority’, *Journal of Public Economics*, 184: 104-126.

Agranov, M., A. Dianat, L. Samuelson, L. Yariv (2021), ‘Paying to Match: Decentralized Markets with Information Frictions’, CESifo Working Paper Series 8820, CESifo.

Agranov, M., M. Elliott (2021), ‘Commitment and (In)Efficiency: A Bargaining Experiment’, *Journal of the European Economic Association*, 19(2): 790-838.

Agranov, M., C. Tergiman (2014), ‘Communication in Multilateral Bargaining’, *Journal of Public Economics*, 118: 75-85.

Ali, N. (2006), ‘Waiting to settle: Multilateral bargaining with subjective biases’, *Journal of Economic Theory*, 130(1): 109-137.

Ambrus, A., E. Chaney, I. Salitskiy (2018), ‘Pirates of the Mediterranean: An empirical investigation of bargaining with asymmetric information,’ *Quantitative Economics*, 9(1): 217-246.

Ambrus, A., S. Lu. (2015), ‘A Continuous-Time Model of Multilateral Bargaining’, *American Economic Journal: Microeconomics*, 7 (1): 208-49.

Andreyanov, P., E. Caoui (2022), ‘Secret reserve prices by uninformed sellers’, *Quantitative Economics*, 13, 1203–1256.

Ausubel, L., P. Cramton, R. Deneckere (2002), ‘Bargaining with incomplete information’, in

Handbook of Game Theory (Amsterdam: Elsevier Science BV), 3, 1897-1945.

Baccara, M., A. Imrohoroglu, A. Wilson, L. Yariv (2012), 'A Field Study on Matching with Network Externalities', *American Economic Review*, 102(5): 1773-1804.

Backus, M., T. Blake, B. Larsen, S. Tadelis (2020), 'Sequential Bargaining in the Field: Evidence from Millions of Online Bargaining Interactions', *The Quarterly Journal of Economics*, 135(3): 1319-1361.

Backus, M., T. Blake, S. Tadelis (2019), 'On the Empirical Content of Cheap-Talk Signaling: An Application to Bargaining', *Journal of Political Economy*, 127: 1599-1628.

Baron, D., J. Ferejohn (1989), 'Bargaining in legislatures', *American Political Science Review*, 83(4): 1181-1206.

Bennett, E. (1997), 'Multilateral Bargaining Problems', *Games and Economic Behavior*, 19(2): 151-179.

Bloch, F., M. Jackson (2007), 'The formation of networks with transfers among players', *Journal of Economic Theory*, 133(1): 83-110.

Bochet, O., M. Khanna, S. Siegenthaler (2023), 'Beyond Dividing the Pie: Multi-Issue Bargaining in the Laboratory', *Review of Economic Studies*, forthcoming.

Bodoh-Creed, A. (2020), 'Risk and Conflation in Matching Markets', Haas School of Business, UC Berkeley (mimeo).

Bodoh-Creed, A., J. Boehnke, B. Hickman (2021), 'How Efficient are Decentralized Auction Platforms?', *Review of Economic Studies*, vol. 88(1): 91-125.

Burrows, P., G. Loomes (1994), 'The Impact of Fairness on Bargaining Behaviour', In: Hey J.D. (eds) *Experimental Economics*.

Cason, T., L. Gangadharan (2006), 'Emissions variability in tradable permit markets with imperfect enforcement and banking', *Journal of Economic Behavior & Organization*, 61(2), 199-216.

Camerer, C., G. Nave, A. Smith (2019), 'Dynamic Unstructured Bargaining with Private Information: Theory, Experiment, and Outcome Prediction via Machine Learning', *Management Science* 65(4):1867-1890.

Cameron, C., J. Gelbach, D. Miller (2008), 'Bootstrap-Based Improvements for Inference with Clustered Errors', *The Review of Economics and Statistics*, 90(3): 414-427.

Carroll, G. (2019), 'Robustness in Mechanism Design and Contracting', *Annual Review of Economics*, 11, 139-166.

Cason, T., C. Plott (2014), 'Misconceptions and game form recognition: challenges to theories of revealed preference and framing', *Journal of Political Economy*, 122(6): 1235-1270.

Chatterjee, K., W. Samuelson (1983), 'Bargaining under incomplete information', *Operations Research*, 31, 835-851.

Che, Y., Y. Koh (2016), 'Decentralized College Admissions', *Journal of Political Economy*, 124(5): 1295-1338.

Choi, S., A. Galeotti, S. Goyal (2017), 'Trading in Networks: Theory and Experiments', *Journal of the European Economic Association*, 15(4): 784-817.

Comola, M., M. Fafchamps (2018), 'An Experimental Study of Decentralized Networked Markets', *Journal of Economic Behavior & Organization*, 145(C): 567-91.

Compte, O., P. Jehiel (2004), 'Gradualism in Bargaining and Contribution Games', *Review of Economic Studies*, 71: 975-1000.

Condorelli, D., A. Galeotti, L. Renou (2017), 'Bilateral Trading in Networks', *The Review of Economic Studies*, 84(1): 82-105.

Coursey, D. , J. Hovis, W. Schulze (1987), 'The disparity between willingness to accept and willingness to pay measures of value', *Quarterly Journal of Economics*, 102, 679–690.

Croson, R., T. Boles, J. Murnighan (2003), 'Cheap talk in bargaining experiments: Lying and threats in ultimatum games', *Journal of Economic Behavior & Organization*, 51(2):143-159.

Currarini, S., M. Morelli (2000), 'Network formation with sequential demands', *Review of Economic Design*, 5: 229-49.

Edelman, B., M. Ostrovsky, M. Schwarz (2007), 'Internet Advertising and the Generalized Second Price Auction', *American Economic Review* 97(1): 242-59.

Forsythe, R., J. Kennan, B. Sopher (1991), 'An Experimental Analysis of Strikes in Bargaining Games with One-Sided Private Information', *American Economic Review*, 81(1): 253-78.

Freyberger, B., J. Larsen (2025), 'How Well Does Bargaining Work in Consumer Markets? A Robust Bounds Approach', *Econometrica*, 93(1): 161-194

Gaineddenova, Renata (2022). 'Pricing and Efficiency in a Decentralized Ride-Hailing Platform', Harvard University (mimeograph).

Galeotti, F., M. Montero, A. Poulsen (2019), 'Efficiency Versus Equality in Bargaining', *Journal of the European Economic Association*, 17(6): 1941-1970.

Garratt, R., M. Walker, J. Wooders (2012), 'Behavior in second-price auctions by highly experienced eBay buyers and sellers', *Experimental Economics*, 15: 44-57.

Georgantzís, N., D. Navarro-Martínez (2010), 'Understanding the WTA–WTP gap: Attitudes, feelings, uncertainty and personality', *Journal of Economic Psychology*, 31(6): 895-907.

Goeree, J., L. Lindsay (2020), 'The Exposure Problem and Market Design', *Review of Economic Studies*, 87: 2230-2255.

Güth, W., S. Huck, P. Ockenfels (1996), 'Two-level ultimatum bargaining with incomplete information: An experimental study', *Economic Journal*, 106: 593-604.

Haeringer, G., M. Wooders (2011), 'Decentralized job matching,' *International Journal of Game Theory*, 40(1): 1-28.

Haruvy, E., E. Katok, V. Pavlov (2020), 'Bargaining Process and Channel Efficiency', *Management Science*, 66(7): 2845-2860.

Horton, J., W. Kerr, C. Stanton (2017), 'Digital Labor Markets and Global Talent Flows', NBER Chapters, in: *High-Skilled Migration to the United States and Its Economic Consequences*,

71-108.

Huang, J. J. Kessler, M. Niederle (2022), 'Fairness Concerns Are Less Relevant When Agents Are Less Informed', working paper.

Ikica, B., S. Jantschgi, H. Nax, D. Nuñez Duran, B. Pradelski (2023), 'Competitive Market Behavior: Convergence and Asymmetry in the Experimental Double Auction', *European Economic Review*, forthcoming.

Isoni, A., G. Loomes, R. Sugden (2011), 'The Willingness to Pay-Willingness to Accept Gap, the Endowment Effect, Subject Misconceptions, and Experimental Procedures for Eliciting Valuations: Comment', *American Economic Review*, 101(2): 991-1011.

Jackson, M., H. Sonnenschein, Y. Xing, C. Tombazos, O. Al-Ubaydli (2023), 'The Efficiency of Negotiations with Uncertainty and Multidimensional Deals', Johns Hopkins Carey Business School Research Paper No. 21-14.

Jackson, M., I. Kremer (2006), 'The Relevance of a Choice of Auction Format in a Competitive Environment', *Review of Economic Studies*, 73(4): 961-981.

Johnson E., C. Camerer, S. Sen, T. Rymon (2002), 'Detecting failures of backward induction: Monitoring information search in sequential bargaining', *Journal of Economic Theory*, 104(1): 16-47.

Kagel, J., Y. Lien, P. Milgrom (2010), 'Ascending Prices and Package Bidding: A Theoretical and Experimental Analysis', *American Economic Journal: Microeconomics*, 2(3): 160-85.

Kagel, J., A. Roth (2000), 'The dynamics of reorganization in matching markets: a laboratory experiment motivated by a natural experiment', *Quarterly Journal of Economics*, 115(1): 201-35.

Kagel, J., K. Wolfe (2001), 'Tests of fairness models based on equity considerations in a three-person ultimatum game', *Experimental Economics*, 4(3): 203-219.

Kahneman, D., A. Tversky (1979), 'Prospect Theory: An Analysis of Decision under Risk', *Econometrica*, 47(2): 263-291.

Kahneman, D., L. Knetsch, R. Thaler (1991), 'Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias', *Journal of Economic Perspectives*, 5 (1): 193-206.

Karagözoglu, E. (2019), 'On 'Going Unstructured' in Bargaining Experiments', In: Laslier, JF., Moulin, H., Sanver, M., Zwicker, W. (eds) *The Future of Economic Design*.

Keniston, D., B. Larsen, S. Li, J. Prescott, B. Silveira, C. Yu (2021), 'Fairness in Incomplete Information Bargaining: Theory and Widespread Evidence from the Field', NBER working paper No. 29111

Kirchsteiger, G., M. Niederle, J. Potters (2005), 'Endogenizing market institutions: An experimental approach', *European Economic Review*, 49: 1827-53.

Klemperer, P. (1996), 'Auctions vs. Negotiations', *American Economic Review*, 86: 180-94.

Krishna, V., R. Serrano (1996), 'Multilateral bargaining', *Review of Economic Studies*, 63: 61-80.

Kriss, P., R. Nagel, R. Weber (2013), 'Implicit vs. explicit deception in ultimatum games with

incomplete information', *Journal of Economic Behavior & Organization*, 93: 337-346.

Kroll, E., R. Morgenstern, T. Neumann, S. Schosser, B. Vogt (2014), 'Bargaining power does not matter when sharing losses – Experimental evidence of equal split in the Nash bargaining game', *Journal of Economic Behavior & Organization*, 108: 261-272.

Larsen, B. (2021), 'The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions', *Review of Economic Studies*, 88(2): 851-882.

Larsen, B. and A. Zhang (2021), 'Quantifying Bargaining Power Under Incomplete Information: A Supply-Side Analysis of the Used-Car Industry', SSRN working paper 3990290.

Lee, S. (2017), 'Incentive Compatibility of Large Centralized Matching Markets', *The Review of Economic Studies*, 84(1): 444-463.

List, J. (2003), 'Does market experience eliminate market anomalies?', *Quarterly Journal of Economics*, 118: 41-71.

Luhan, W., O. Poulsen, M. Roos (2019), 'Money or morality: fairness ideals in unstructured bargaining', *Social Choice and Welfare*, 53: 655-675.

McAfee, R. (1992), 'A dominant strategy double auction', *Journal of Economic Theory*, 56 (2): 434-450

Milgrom, P., I. Segal (2020), 'Clock Auctions and Radio Spectrum Reallocation', *Journal of Political Economy*, 128(1): 1-31.

Milgrom, P., R. Weber (1982), 'The Value of Information in a Sealed Bid Auction', *Journal of Mathematical Economics*, 10: 105-14.

Mitzkewitz, M., R. Nagel (1993), 'Experimental results on ultimatum games with incomplete information', *International Journal of Game Theory*, 22(2): 171-198.

Morley, I., G. Stephenson (2015), 'The Social Psychology of Bargaining', Psychology Press.

Murphy, R., K. Ackermann, M. Handgraaf (2011), 'Measuring Social Value Orientation', *Judgment and Decision Making*, 6(8): 771-781.

Muthoo, A. (1990). "Bargaining without Commitment", *Games and Economic Behavior*, 2: 291-7.

Mutuswami, S., E. Winter (2002), 'Subscription mechanisms for network formation', *Journal of Economic Theory*, 106(2): 242-64.

Myerson, R., M. Satterthwaite (1983), 'Efficient mechanisms for bilateral trading', *Journal of Economic Theory*, 29: 265-281.

Nash, J. (1953), 'Two-person cooperative games', *Econometrica*, 21(1):128–140.

Navarro, N., R. Veszteg (2020), 'On the empirical validity of axioms in unstructured bargaining', *Games and Economic Behavior*, 121: 117-145.

Niederle, M., A. Roth (2009), 'Market Culture: How Rules Governing Exploding Offers Affect Market Performance', *American Economic Journal: Microeconomics*, 1(2): 199-219.

Ochs, J., A. Roth (1989), 'An experimental study of sequential bargaining', *American Economic Review*, 79(3): 355-384.

Plott, C., K. Zeiler (2005), 'The willingness to pay-willingness to accept gap, the "endowment effect", subject misconceptions, and experimental procedures for eliciting valuations', *American Economic Review*, 95: 530-545.

Pruitt, D. (2013), 'Negotiation Behavior', Academic Press, New York.

Roth, A., E. Peranson (1999), 'The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design', *American Economic Review*, 89 (4): 748-780.

Rubin, J., B. Brown (1975), 'The Social Psychology of Bargaining and Negotiation', New York: Academic.

Rubinstein, A. (1982), 'Perfect equilibrium in a bargaining model', *Econometrica*, 50: 97-109.

Shupp, R., J. Cadigan, P. Schmitt, K. Swope (2013), 'Institutions and information in multilateral bargaining experiments', *The B.E. Journal of Economic Analysis & Policy*, 13(1): 485-524.

Simonsohn, U., D. Ariely (2008), 'When Rational Sellers Face Nonrational Buyers: Evidence from Herding on eBay', *Management Science*, 54(9): 1624-1637.

Smith, V. L., A. Williams (1982), 'The effects of rent asymmetries in experimental auction markets', *Journal of Economic Behavior & Organization*, 3(1):99-116.

Sherstyuk, K., K. Phankitnirundorn, M. Roberts (2021), 'Randomized double auctions: gains from trade, trader roles, and price discovery', *Experimental Economics*, 24(4),1325-1364.

Sobel, J., I. Takahashi (1983), 'A Multi-Stage Model of Bargaining', *Review of Economic Studies*, 50: 411-426.

Srivastava, J. (2001), 'The role of inferences in sequential bargaining with one-sided incomplete information: Some experimental evidence', *Organizational Behavior and Human Decision Processes*, 85(1): 166-187.

Thaler, R. (1988), 'Anomalies: The Winner's Curse', *Journal of Economic Perspectives*, 2: 191-202.

Tremewan, J., C. Vanberg (2016), 'The dynamics of coalition formation – A multilateral bargaining experiment with free timing of moves', *Journal of Economic Behavior & Organization*, 130: 33-46.

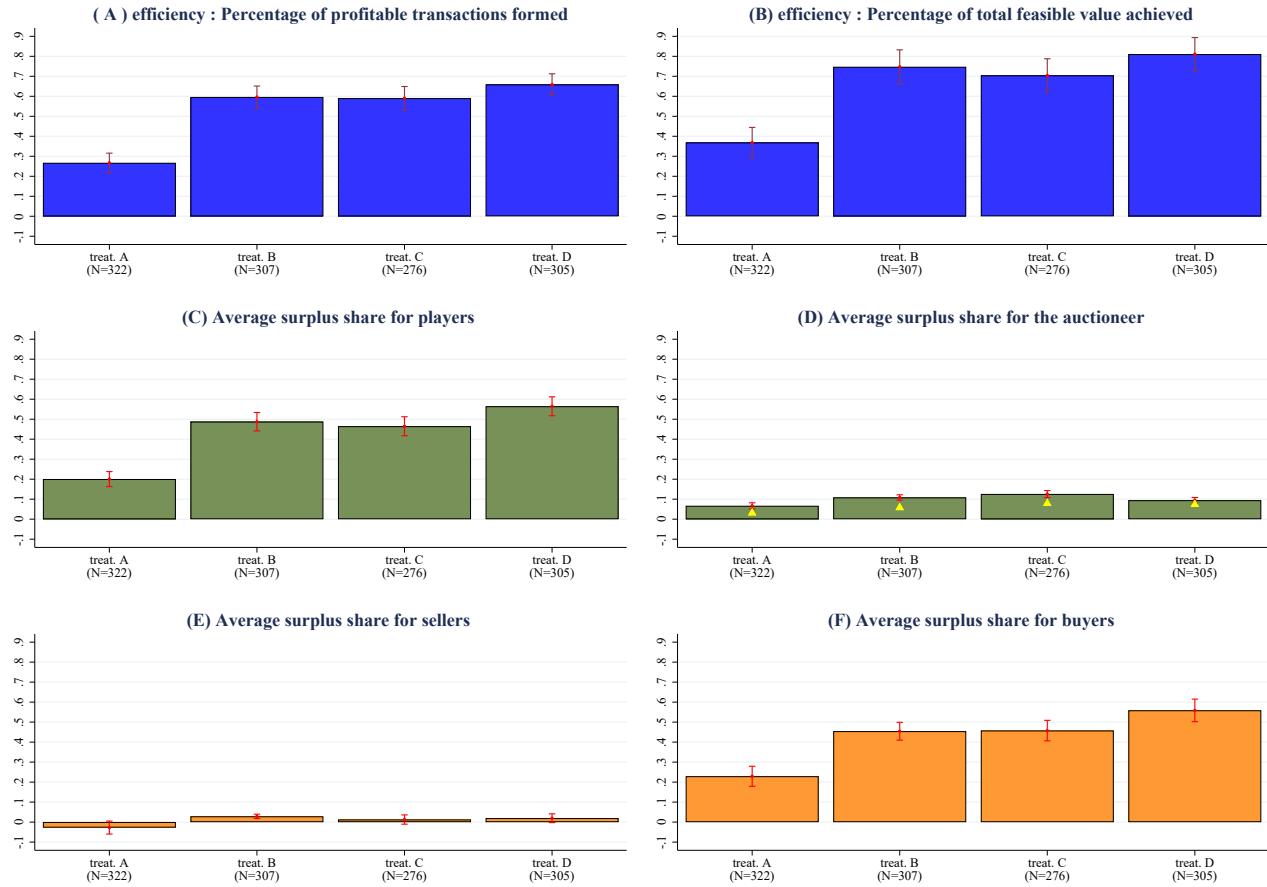
Valley, K., L. Thompson, R. Gibbons, M.H. Bazerman (2002), 'How Communication Improves Efficiency in Bargaining Games', *Games and Economic Behavior*, 38: 127-155.

Vincent, D. (1989), 'Bargaining with common values', *Journal of Economic Theory*, 48(1): 47-62.

Williams, S. (1987), 'Efficient Performance in Two Agent Bargaining', *Journal of Economic Theory*, 41: 154-172.

Tables and Figures

Figure 1: Efficiency and surplus



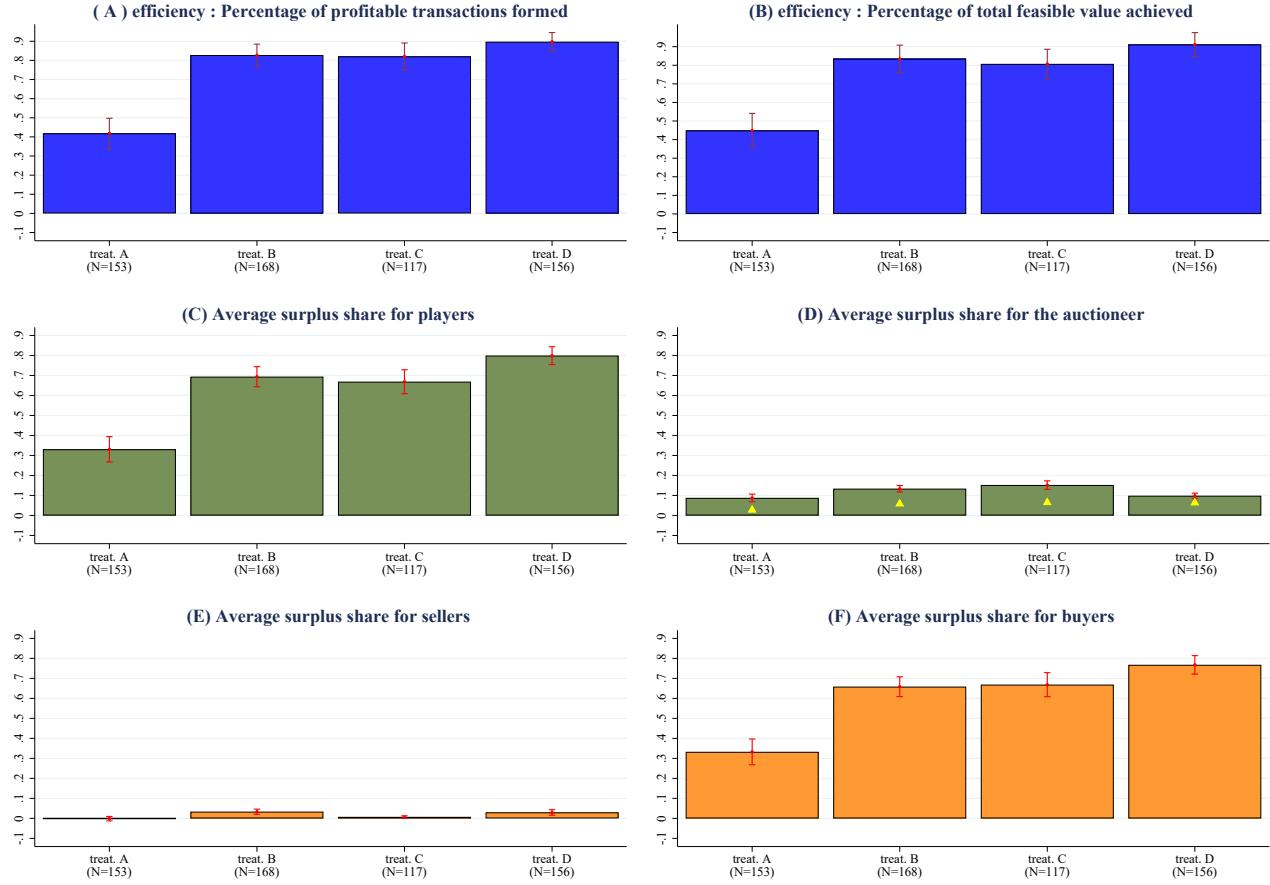
Note : Standard errors shown in red. Yellow triangles in panel D represent the auctioneer surplus if all transactions yielded the minimal fee of 1.

Figure 2: Efficiency and surplus, low-value transactions ($v_{ij} < 10$)



Note : Standard errors shown in red. Yellow triangles in panel D represent the auctioneer surplus if all transactions yielded the minimal fee of 1.

Figure 3: Efficiency and surplus, high-value transactions ($v_{ij} \geq 10$)



Note : Standard errors shown in red. Yellow triangles in panel D represent the auctioneer surplus if all transactions yielded the minimal fee of 1.

Table 1: Statistics by role

type	treat.	N	(1)	(2)	(3)	(4)	(5)
			mean v_{ij}^k	# active players	# bids placed	# losing bids	# bids $v_{ij}^k - 1$
<i>All</i>	<i>All</i>	10560	0.82	4711	13331	301	1251
$v_{ij}^k = 0$	<i>All</i>	5852	0.00	499	1302	156	67
	T_A	1262	0.00	63	63	29	9
	T_B	1501	0.00	167	497	45	29
	T_C	1681	0.00	130	269	50	15
	T_D	1408	0.00	139	473	32	14
$v_{ij}^k > 0$	<i>All</i>	3014	6.00	2576	8325	0	423
	T_A	903	6.00	654	654	0	33
	T_B	750	5.97	689	2232	0	147
	T_C	557	6.18	516	1387	0	74
	T_D	804	5.90	717	4052	0	169
$v_{ij}^k < 0$	<i>All</i>	1694	-5.55	1636	3704	145	761
	T_A	475	-5.81	439	439	56	154
	T_B	389	-5.26	385	993	25	193
	T_C	402	-5.37	393	795	35	211
	T_D	428	-5.68	419	1477	29	203

Table 2: Profit margins by role

	(1)	(2)	(3)
dep. var.: $b_{ij,g}^k$ (bid by player k on transaction ij in game g)			
bids	first	last	last
treatments	<i>All</i>	<i>All</i>	T_B, T_C, T_D
v_{ij}^{k-}	1.404***	1.327***	1.353***
margin	[40.4%] (0.124)	[32.7%] (0.125)	[35.3%] (0.143)
v_{ij}^{k+}	0.192***	0.302***	0.334***
margin	[80.8%] (0.014)	[69.8%] (0.017)	[66.6%] (0.020)
S	0.024 (0.021)	0.025* (0.014)	0.022 (0.020)
T_g	yes	yes	yes
λ_s	yes	yes	yes
Const.	-0.101 (0.344)	-0.451* (0.224)	-0.289 (0.202)
Obs	10,560	10,560	7,920
R-sq	0.287	0.341	0.307

Notes: OLS results reported. $v_{ij}^{k-} = v_{ij}^k$ if $v_{ij}^k < 0$ and 0 otherwise. $v_{ij}^{k+} = v_{ij}^k$ if $v_{ij}^k > 0$ and 0 otherwise. We control for game order S (from 1 to 8), treatment dummies T_g , and session-level fixed effects λ_s . Wild-bootstrapped p-values in parentheses, clustered at the unique group level. *** p<0.01, ** p<0.05, * p<0.1.

Figure 4: Dynamics of cumulative bids in T_D

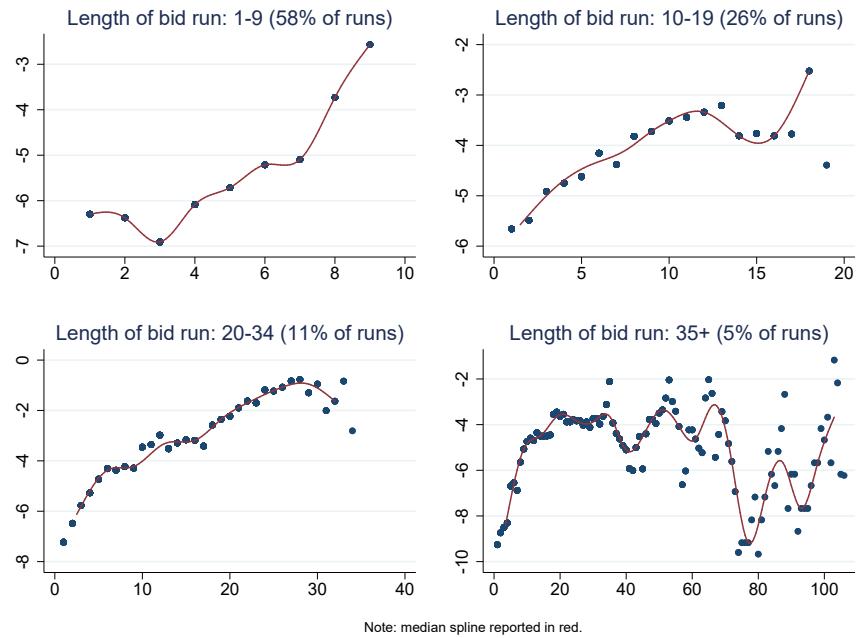


Table 3: Statistics on bids in T_D , before and after activation

	(1)	(2)	(3)	(4)	(5)
Quartile	all	Q1	Q2	Q3	Q4
length (n. of ticks)	all	1-9	10-19	20-34	35+
n. bids	4995	1155	1488	1223	1129
panel A: transactions					
n. transactions	425	245	111	48	21
mean v_{ij}	5.21	4.60	5.39	6.88	7.52
% deals	0.45	0.46	0.47	0.48	0.29
panel B: bids placed before activation					
n. bids	1947	629	687	405	226
bids per transaction	4.58	2.57	6.19	8.44	10.76
% unique/first bids	0.08	0.10	0.06	0.07	0.05
% of increasing bids	0.68	0.68	0.68	0.67	0.73
% increasing by 1 unit	0.55	0.55	0.56	0.50	0.58
% of decreasing bids	0.24	0.22	0.26	0.25	0.22
% decreasing by 1 unit	0.15	0.15	0.17	0.14	0.14
panel C: bids placed after activation					
n. bids	3048	526	801	818	903
bids per transaction	7.17	2.15	7.22	17.04	43.00
% unique/first bids	0.04	0.09	0.05	0.03	0.01
% of increasing bids	0.51	0.44	0.47	0.53	0.57
% increasing by 1 unit	0.44	0.38	0.43	0.47	0.47
% of decreasing bids	0.45	0.46	0.48	0.44	0.42
% decreasing by 1 unit	0.34	0.40	0.41	0.33	0.26

Online Appendix A: Experimental Instructions

A.1 Visual Layout

We now describe the main features of the experiment's visual layout. The four images below (A.1 to A.4) are screenshots from the game's interface that are included in the instructions for participants. Since the experiment is implemented in French, a legend with corresponding French and English terms is provided on the right side of each figure.⁶⁷ A complete translation of the instructions for players, with reference to these pictures, is given in Section A.2 below.

General interface

Players appear as nodes arranged in a circle (or hexagon) with the player herself always positioned at the bottom of the circle ('ME'). Each player is represented by a letter identifier, and both the letter identifiers and the position of the players are reshuffled from one game to the next. Transactions are visually represented as links between nodes. The links that are currently active – i.e., for which the total bid is currently positive – are displayed as a thick solid line. Those that are inactive are displayed as a thin dotted line. The links that cannot be formed are not shown, i.e., they have no line.

Gain tag

Above each link is a tag reporting the hypothetical *gain* the player *would* derive from it, based on her current bid, should the link be formed at the end of the game. This is simply computed as her *value* for the transaction minus the last *bid* she made on it.⁶⁸ These gains are color-coded: positive gains appear in green; negative gains appear in red. The activation status of the link is disclosed via the line pattern: the line representing the link and the lines contouring the tag turn from dotted to solid when the link is activated, and the words on the tag appear in bold. For instance, in Figure A.4, link A-C is activated. This design, which is the result of careful research, offers the advantage of presenting the state of play in an intuitive way: all the relevant information is instantaneously available to players in an easily understandable format.

⁶⁷In the instructions, we have chosen to represent only screenshots from the trial games, with 7 links instead of 10.

⁶⁸Note that this is not the same naming convention used in the analysis, where we define the gain as 0 whenever $b_{ij,t} < 1$ (see Section 2.1). However, for the sake of the interface and instructions we found it more appealing to display at any moment the difference between value and current bid (which corresponds to the hypothetical gain if the transaction took place), since the activation status is conveyed to players through line patterns as explained below.

Bargaining process

The bidding process is as follows. While a transaction is being auctioned (either because it's its turn in T_B and T_c , or because the player selects it in T_D), the (solid or dotted) line representing the link gets highlighted, the nodes at the link's end turn red, and a dialog box opens in the middle of the screen. This dialog box contains two pieces of information: the *value* for the player (which is set exogenously and cannot be changed); and the player's current *bid*. Players can bid on links involving themselves with either an offer – a positive number – or a request – a negative number. For example, in Figure A.2 and A.3, player E has opened the dialog box for link E-D which has value +3 to her, and she is contemplating an offer of 2 units (Figure A.2) or a request of 2 units (Figure A.3).⁶⁹ Players can also bid on transactions that do not involve them, but in this case their bid has to be positive, i.e., no 'request' option is available on the dialog box (Figure A.4, for link A-B). No constraint is imposed on the magnitude of bids made by subjects.

At the beginning of a game, no link is activated and the default value of all bids is set to zero. In games that include multiple rounds, the player's bid in the immediately preceding round becomes the default starting bid in the following round.

⁶⁹In Figures A.2 and A.3, player E has not validated her bid yet with the OK located within the dialog box. Thus, the tag on the link still reports the gain of 3 which is the gain associated with the default bid of 0. The link E-D is currently not activated, so the link and tag lines appear dotted.

Figure A.1

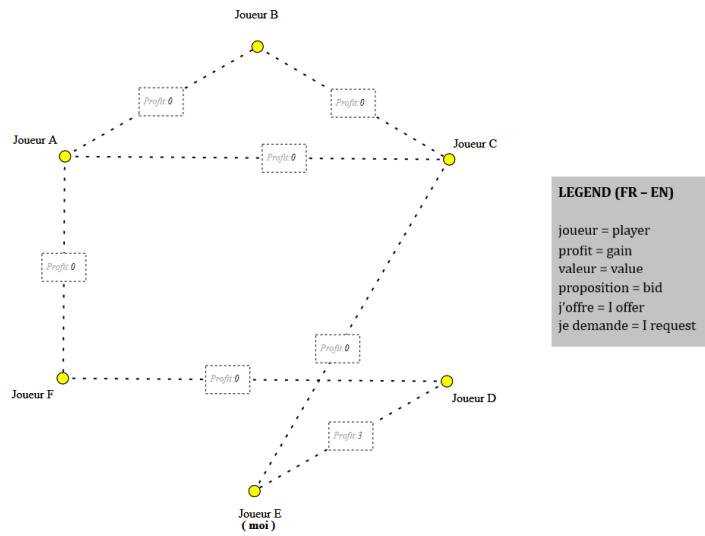


Figure A.2

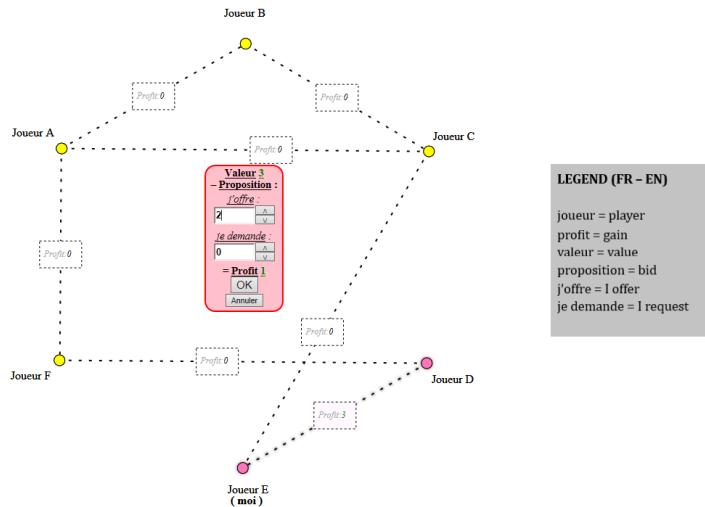


Figure A.3

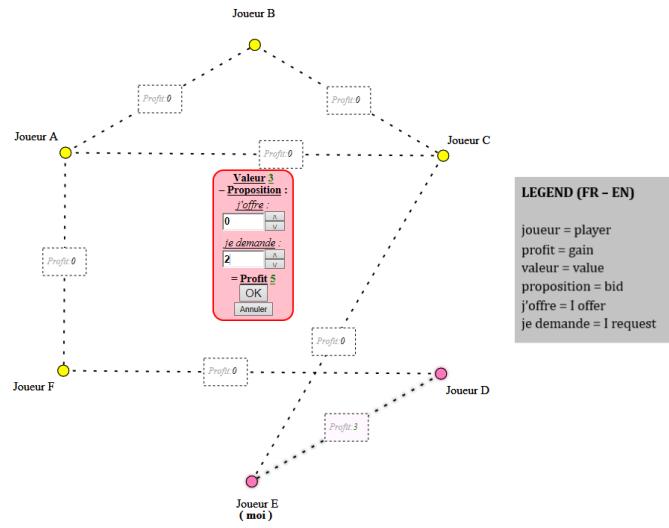
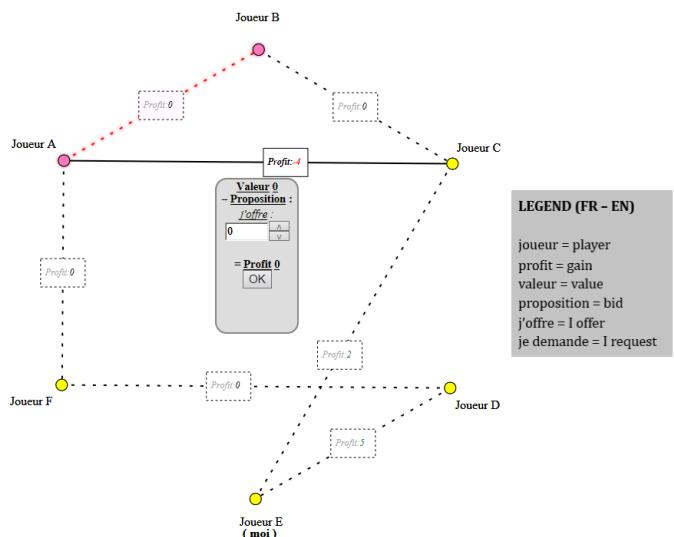


Figure A.4



Example of matrices

Example of two-sided matrix

player i	player j	gain 1	gain 2	gain 3	gain 4	gain 5	gain 6
1	2	5	6	0	0	0	0
1	3	-	-	-	-	-	-
1	4	-	-	-	-	-	-
1	5	-9	0	0	0	-6	0
1	6	10	0	0	0	0	-2
2	3	-	-	-	-	-	-
2	4	0	9	0	-4	0	0
2	5	0	-7	0	0	4	0
2	6	0	8	0	0	0	10
3	4	-	-	-	-	-	-
3	5	0	0	3	0	6	0
3	6	-	-	-	-	-	-
4	5	0	0	0	-9	2	0
4	6	0	0	0	9	0	0
5	6	0	0	0	0	-5	0

Example of multi-sided matrix

player i	player j	gain 1	gain 2	gain 3	gain 4	gain 5	gain 6
1	2	-9	-9	5	0	0	6
1	3	-8	0	8	6	5	9
1	4	-	-	-	-	-	-
1	5	7	3	0	8	-2	0
1	6	-8	2	0	8	0	-7
2	3	-	-	-	-	-	-
2	4	0	-8	0	-6	0	6
2	5	0	-6	0	8	-6	10
2	6	-	-	-	-	-	-
3	4	-	-	-	-	-	-
3	5	0	6	-1	0	-10	9
3	6	0	10	0	0	2	-7
4	5	5	0	0	5	-3	0
4	6	0	0	9	-5	4	0
5	6	-	-	-	-	-	-

A.2 Instructions to subjects⁷⁰

Thank you for participating in this experiment about decision making. Please turn off your phone and put it away. You are not allowed to communicate with other participants during the session unless you are invited to do so by an experimenter, or you will be disqualified from the payment. All your decisions stay anonymous.

Today you will play a game whose rules are explained in what follows. You will receive 10 euros for showing up on time. In addition, you can accumulate earnings during the session according to your decisions and those of the other participants. At the end of the session, your earnings will be converted into euros and paid out in cash in private. Your earnings will remain confidential.

Description of the game

General framework

This is a link-formation game between 6 players, who are located around a circle and labeled with the letters A, B, C, D, E, and F.

You are always the player at the bottom of the circle: your icon is indicated by 'ME' as well as by your letter identifier.

- *Example:* in Figure 1-3 you are player E.

All the links that can be formed are displayed with a dotted line on the graph.

- *Example:* in Figure 1-3 the links between E-F, and E-A cannot be formed.

When a link is formed, it appears with a thick, solid black line, and it is visible to all players.

- *Example:* in Figure 4, the link A-C is formed.

Gain

On each link you will see a tag with a dotted or solid frame, which contains a piece of information, the 'gain'. If you click on the tag, the two players at the link's ends are displayed in red and a dialog box appears (Figure 2).

⁷⁰Translated from French. Figures 1 to 4 in the instructions correspond to the Figures A.1 to A.4.

Value

The first piece of information in this dialog box is the ‘value’, which indicates how much you are remunerated if this link is formed by the end of the game. The value remains constant during the game, and it can be positive or negative (a negative value represents a loss).

- Please note that the values are independent for each player. So, if the E-D link has a value of 3 for you (as in Figure 2), it does not mean that it has value 3 for the others!
- You only know your own values, not those of the other players.

Bids

The second piece of information in the dialog box is the ‘bid’, which indicates the amount you are offering or requesting for the link to be formed. This is your decision variable during the game: you are free to place bids in order to form links.

- **If a link concerns you directly, you can place positive bids (offers)** – i.e., you offer to pay for the link to be formed – **or negative bids (requests)** – i.e., you ask compensation for the link to be formed. *Example:* in Figure 2, to form the E-D link, you can place either an offer or a request (i.e. a positive or a negative bid).
- **If a link does not directly concern you, you can only place positive bids (offers)** – i.e. you can only offer to pay for the link to be formed. *Example:* in Figure 4, to form the link A-B, you can only make an offer.
- You can place bids as you wish. This means that you are not obliged to place bids on links that directly concern you, and that you can place bids on links between third players.
- You can choose the amount of the bids as you wish, without limits (except those set by the above rules).
- You do not observe the bids of other players.

Link formation and gain

At any point in the game, if the sum of the bids by all players for a given link is strictly positive, then the link is formed. That is, if the sum of all positive bids (offers) exceeds the negative bids (requests), the link is formed. Still, players may be able to revise their bids later on, and thus as the game unfolds formed links can be deleted if the sum of bids becomes negative.

When a link is formed, the tag appears with a solid frame line (Figure 4: for link A-C both the link line and the gain tag appear with a solid rather than a dotted line).

Your final gain only depends on the links that are formed by the end of the game: it is zero if the link is not formed, and it is based on your last bid if the link is formed. There are four possible cases:

1. If a **link is not formed**: Final Gain = 0
2. If your current **bid is positive (offer)** and the link is formed: Final Gain = Value – Offer
3. If your current **bid is negative (request)** and the link is formed: Final Gain = Value + Request
4. If your current **bid is zero (no bid)** and the link is formed: Final Gain = Value

Example: In Figure 2 your value for the E-D link is +3. If you place a positive bid (offer) of 2, your gain for the E-D link would be $3 - 2 = 1$ if this link is formed by the end of the game. If you place a negative bid (request) of 2 (Figure 3), your gain would be $3 + 2 = 5$ if the link is formed by the end of the game.

Positive gains are displayed in green on the tag, negative gains are displayed in red.

- *Example:* in Figure 4 the gain for A-C is negative (red), while the gain for E-D is positive (green).

Your total gain of the game is the sum of the gains for all links that are formed at the end of the game.

- *Example:* if the B-C link is formed and then deleted before the end of the game, it will be considered ‘non-existent’, i.e., as not being formed, for the purpose of calculating the game’s total gains.

Phases of the game

The game is organized in two phases.

First phase

You start the first phase with no links formed. You now have time to place your bids in the order you wish. To place a bid, click on the link tag, the dialog box opens, and enter your bid.

Careful:

- you must click on the tag that contains the gain, not on the line!
- when you enter a bid, be sure to validate your choice by pressing the OK button within the dialog box. Validated choices appear in light blue.

Once you have finished placing your bids, click the OK button at the bottom of the screen to move on to the second phase of the game.

Second phase

Once you have completed the first phase of the game, you move on to the second phase. During this phase, you will have the opportunity (or not) to revise your bids according to one of 4 scenarios described below.

The relevant scenario is only announced at the beginning of the second phase. This means that in the first phase, you do not know if and how you will have the opportunity to revise your bids.

Scenario A

In Scenario A, the bids from the first phase are final. This means that:

- the links for which the sum of all bids in the first phase is positive are formed,
- the total gains of the game is calculated for each player on the basis of the links formed.

Scenario B

In Scenario B, the game starts with the bids from the first phase, but you have the opportunity to revise them, one link at a time.

Scenario B is organized into several rounds. In each round, the computer visits all the links (in a random order that changes from round to round), and one link at a time appears in red:

- when a link appears in red, the corresponding dialog box pops up and you can revise your bid freely. You can change your bid even if you did not place any bid in the first phase!
- you have at least 10 seconds to enter your new bid (longer if the other players also revise their bids). When time runs out, the dialog box turns gray.

Example: in Figure 4, link A-C is formed, link A-B is currently being revised, and time is elapsing (because the background color of the dialog box is gray).

Stopping rule:

- The game ends before the 6th round if during an entire round the bids of all players do not change;
- At the end of the 6th and 7th rounds the game ends with 50% probability, and at the end of the 8th round the game is forced to an end (even if some bids have changed within the round).

Scenario C

In Scenario C, you have the opportunity to revise the bids from the first phase one link at a time, as in Scenario B.

All rules are the same as in Scenario B, except for the stopping rule:

- The game ends before the 6th round, if during a whole round the links do not change;
- At the end of the 6th and 7th rounds the game ends with 50% probability, and at the end of the 8th round the game is forced to an end.

The difference is that in scenario B the game stops when the bids stay the same, and in scenario C it stops when the links stay the same.

Example: imagine that in round 2 you make an offer of 4 for link A-C and the link is formed, and in round 3 you lower your offer to 3. All other players do not change their bids in round 3, and the sum of bids for this link remains positive. In Scenario B, the game does not end yet, and you proceed to round 4. In Scenario C, the game ends with round 3.

Scenario D

In Scenario D, the game starts from the bids of the first phase, and you have the opportunity to revise your bids for all links simultaneously.

This means that there is no division between rounds and it is your responsibility to click on the tag of the links you want to revise. You can do this freely, in any order, and as many times as you like.

The links that are currently formed (because the sum of the bids is positive) are displayed in real time, i.e. you can see thick lines appearing and disappearing in real time.

If you do not make any changes for (at least) 20 seconds, the game stops. This time is longer if the other players also make changes. When the time runs out, the screen turns gray.

The session's unfolding

In today's session you will take a preliminary quiz, and then you will start with 3 trial games (to get used to the software). You will play scenarios B, C and D (in that order).

After the trial games, you will play 8 games.

In each game, the players you play with remain the same, but the letter identifiers will change. That is:

- You may be called D during the first game, and A during the second;
- You never know how the other players changed positions (i.e., the player named C during the first game may be named D in the second game, etc.).

End of the session

At the end of the session, you will be asked to answer a short final questionnaire. Your answers are anonymous and confidential.

After the questionnaire, you will receive information about your total gains in each of the 8 games. The computer will draw 2 games out of 8 and your earnings will be calculated based on the total gain from these two games. Your total gain in the trial games will not be taken into account.

The payment rule is as follows: 10 euros fixed payment + 0.5 euros for each point.

To get paid and leave the room, you have to wait (silently) until we call you.

Please review these instructions carefully. If you have any questions, please raise your hand. We will come to you immediately to answer your questions in private.

Online Appendix B: Theoretical motivation

In our experiment, players can engage simultaneously in multiple independent trading games. In what follows we focus on the benchmark scenario where two players negotiate one transaction under private information. We first discuss the case where only sealed bids (i.e. take-it-or-leave-it offers) are allowed, and derive the efficiency bounds if bids are placed by the buyer only, the seller only, or both players simultaneously. We then claim that a setting of bargaining without commitment eludes a theoretical characterization in terms of ex-post efficiency.

Consider a simple example with one seller, one buyer, and one unit of the unique good. The privately-known reservation values of the seller and buyer are denoted t_s and t_b , respectively: the seller needs to receive at least t_s to cover the cost of producing and selling the good, and the buyer needs to pay at most t_b to gain. Both t_s and t_b are independently drawn from the same uniform distribution which, without loss of generality, we posit to span the $[0, 1]$ interval. While t_s and t_b are private knowledge to the seller and buyer, respectively, both hold uninformative – i.e., uniform – priors over the $[0, 1]$ interval about the private valuation of the other.⁷¹ Seller and buyer both get a payoff of 0 in case of no sale. Trade is profitable whenever $t_s < t_b$, which happens 50% of the time.

Let us first consider the market design in which the seller makes a take-it-or-leave-it offer to the buyer at price p_s .⁷² The offer that maximizes the seller's profit solves the following well-known optimization problem (e.g., Myersson 1981):

$$\text{Max}_{p_s \geq t_s} (1 - p_s)(p_s - t_s)$$

where $1 - p_s = \text{Pr}(t_b \geq p_s)$ is the probability of selling the good. The optimal price offer is $p_s^* = \frac{1+t_s}{2}$. Since $p_s^* > t_s$ for all $t_s < 1$, ex post efficiency is not achieved: $\text{Pr}(t_s \leq t_b) > \text{Pr}(p_s^* \leq t_b)$ and half of the mutually beneficial trades do not occur.

The situation is similar if it is the buyer who makes a take-it-or-leave-it offer. In that case, the optimization problem of the buyer is:

$$\text{Max}_{p_b \leq t_b} p_b(t_b - p_b)$$

where $p_b = \text{Pr}(t_s \leq p_b)$ is the probability of buying. The optimal buying price offer is $p_b^* = \frac{t_b}{2}$, which similarly rules out half of the feasible trades. We also note that, in the seller

⁷¹As pointed out in the main text, subjects in our experimental setting are not explicitly told how valuations are drawn. But during the training games they are confronted with multiple examples drawn from the same interval as in the main games, thus it is not far fetched to assume that their form correct priors.

⁷²This is equivalent to setting a reservation price in an auction to a single buyer.

case, all price requests are $1/2$ or above while in the buyer case, all price offers are $1/2$ or below. Who makes the offer thus has massive implication for the distribution of surplus.

We now turn to a setting where both seller and buyer make simultaneously a take-it-or-leave-it offer and a trade occurs only if $p_s \leq p_b$. To facilitate comparison with our experimental results, we solve the Nash equilibrium for when the difference $p_b - p_s$ goes to the silent auctioneer – i.e., if a sale takes place, the seller gets the requested price p_s and the buyer pays the offered price p_b . In this case, the equilibrium offers are the joint solutions to two optimization problems of the form:

$$\begin{aligned} \text{Max}_{p_s \geq t_s} \Pr(\tilde{p}_b \geq p_s)(p_s - t_s) \\ \text{Max}_{p_b \leq t_b} \Pr(p_b \geq \tilde{p}_s)(t_b - p_b) \end{aligned}$$

where \tilde{p}_b is unknown to – but taken as given by – the seller and \tilde{p}_s is unknown to – but taken as given by – the buyer. Given that private values are uniformly distributed, $\Pr(\tilde{p}_b \geq p_s) = p_b^{\max} - p_s$ where p_b^{\max} is the maximum price that the buyer offers and $\Pr(p_b \geq \tilde{p}_s) = p_b - p_s^{\min}$ where p_s^{\min} is the lowest price the seller offers. This gives

$$p_s^* = \frac{t_s + p_b^{\max}}{2} \text{ and } p_b^* = \frac{t_b + p_s^{\min}}{2}$$

The lowest price the seller offers is when the cost of production is 0, hence $p_s^{\min} = p_b^{\max}/2$. The highest price the buyer offers is when $t_b = 1$, which implies that $p_b^{\max} = \frac{1+p_s^{\min}}{2}$. Combining the two to obtain the Nash equilibrium, we find that $p_s^{\min} = \frac{1}{3}$ and $p_b^{\max} = \frac{2}{3}$. This in turn allows us to calculate the highest value of t_s at which a sale happens. From the seller's participation constraint, it binds when $p_s^* = t_s$, that is when $t_s = p_b^{\max} = 2/3$. Above that price, the seller sells nothing since $2/3$ is the maximum price the buyer offers. Hence the seller makes price offers for $t_s \in [0, 2/3]$. Similar calculations for the buyer yield $t_b \in [1/3, 1]$. Numerical integration shows the probability of trade to be $\sim 29\%$ in this case, a result reminiscent of Chatterjee and Samuelson (1983).

In our experimental game, players have a brief window of opportunity to renegotiate.⁷³ This design has two important advantages: first, it depicts realistically several real-life trade settings (i.e. digital and banking services, real estate trades, online purchases) where agents can renegotiate and/or retract. Second, it widens the pool of behavioral strategies that the players could adopt in the bidding process. As we know from Muthoo (1990), if buyer and seller are allowed to continue making alternating offers and counter-offers without commit-

⁷³This is also true for the experiment by Camerer et al. (2019), where bargaining parties have 1.5 seconds to change their mind.

ment, any division of the gains from trade can be sustained as a subgame perfect equilibria of a modified bargaining model (Rubinstein 1982). This theoretical indeterminacy, which could be seen as a drawback from a purely theoretical perspective, motivates our empirical exploration. In our setting, bargaining can increase ex-post efficiency and the division of gains from trade is an open question. This key observation forms the starting point of our investigation.

Online Appendix C: Ancillary results

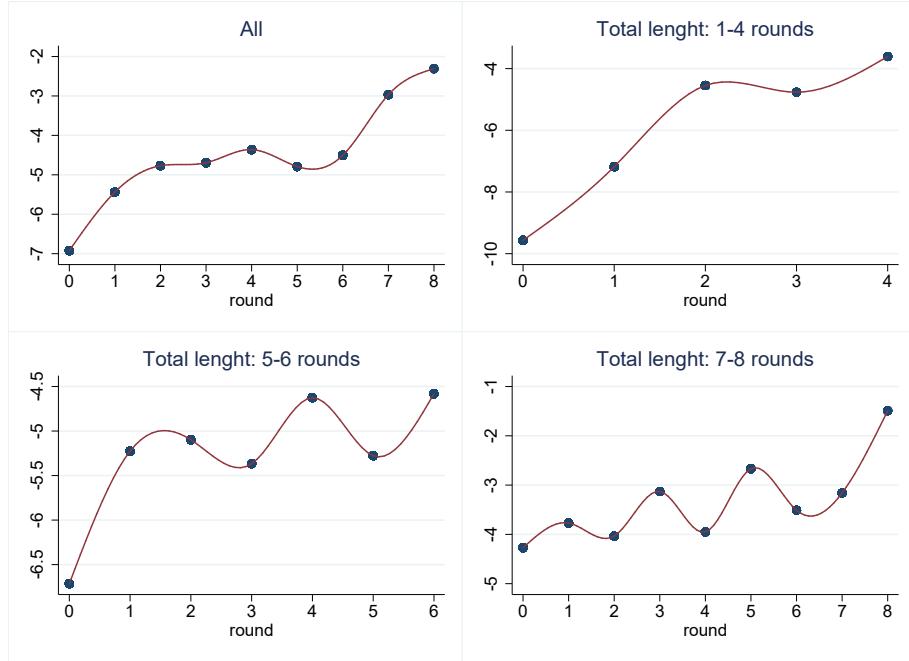
C.1 The Dynamics of Sequential Bargaining

We now investigate the bidding dynamics in sequential treatments T_B and T_C , where bargaining is organized into rounds. We observe 3,671 round outcomes over 880 unique transactions. In analogy with Section 5.1 we estimate a model of the form

$$b_{ij,r} = \alpha + \sum \beta \lambda_r + \lambda_{ij*m} + \varepsilon_{ij,t_{ij}} \quad (3)$$

where $b_{ij,r}$ is the cumulative bid on transaction ij at the end of round r . The regression includes a set of dummies λ_r representing the round number (from 1 to 8), and transaction-per-matrix fixed effects λ_{ij*m} ($10 * 8 = 80$ effects) to control for potential confounding effects correlated with matrix structure. Our estimate of interest is how the predicted values of the cumulative bid $b_{ij,t_{ij}}$ evolves over rounds, which we plot in Figure A.5. In the upper-left panel of the figure we retain all transactions and rounds. In the other three panels we split the sample according to the length of the ‘bid run’, namely, the length of the bargaining process for transaction ij measured in total number of rounds: 1-4 rounds (30% of runs), 5-6 rounds (48% of runs), or 7-8 rounds (22% of runs). Results show that cumulative bids increase with the number of rounds, in line with the dynamics observed in Figure ???. This is consistent with experimental subjects playing a bidding-up strategy to reveal profitable transactions. We also remark an increase in noise in the upper tail of the fourth graph (7-8 rounds), which compares well with the non-monotonicity of long bids in Figure ??.

Figure A.5: Dynamics of cumulative bids in T_B and T_C



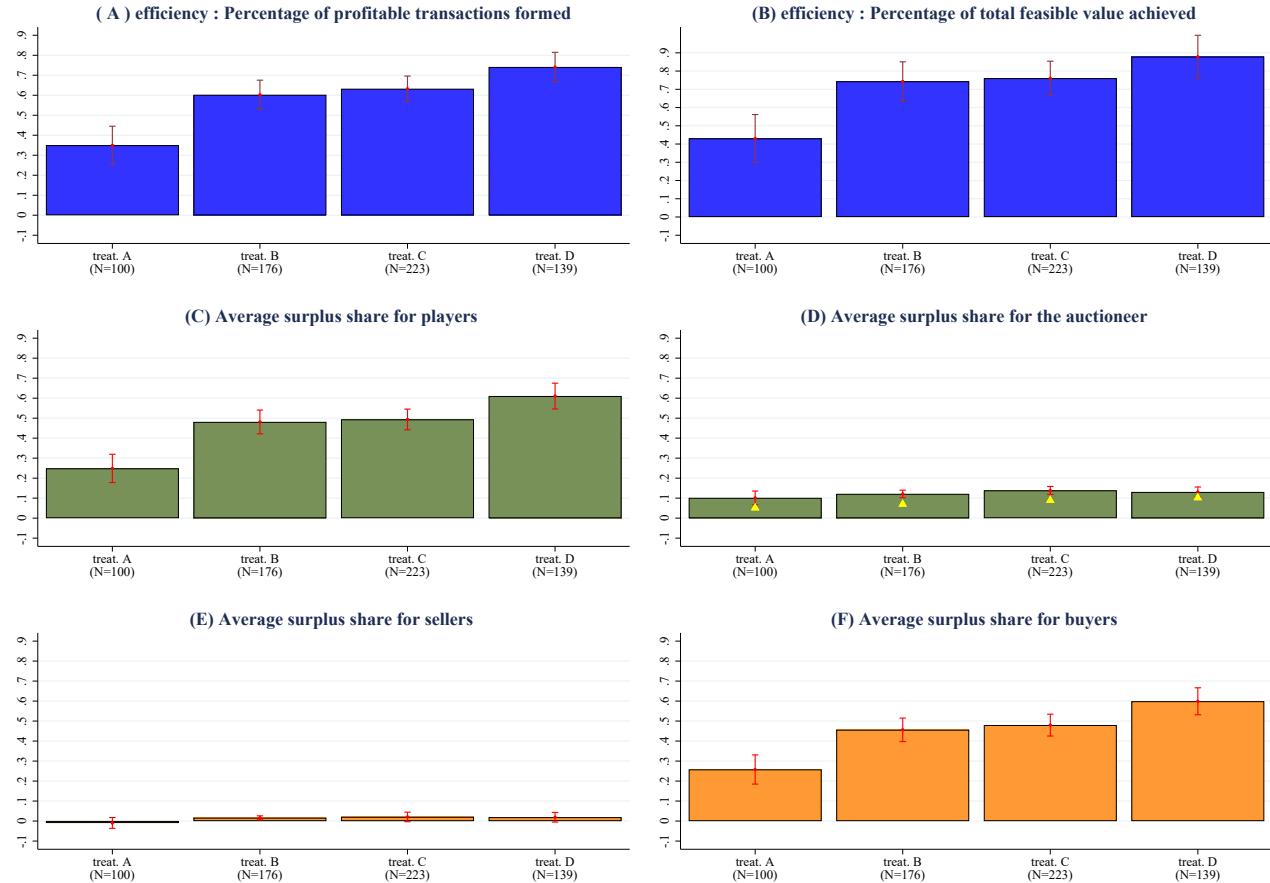
We conclude by briefly discussing how players react to the information of a transaction being formed in sequential treatments. Table A.1 below reproduce the content of Table ?? for bids placed in treatments T_B and T_C . It confirms all the main findings of Table ??, namely: most bids placed before activation are increasing, and often increasing by exactly one unit (67% and 49% respectively for all transactions confounded). After activation the situation reverses, and most bids are decreasing (54%), to a large extent by exactly one unit (42%). That is, approximately half of subjects respond to a formed transaction by decreasing their bid in the hope to extract more surplus. These statistics remain remarkably stable across quantiles (from Q1 to Q3, defined as above), and are in line with evidence on discovery and appropriation discussed for the simultaneous bargaining treatment.

Table A.1: Statistics on bids in T_B and T_C before vs. after activation

	(1)	(2)	(3)	(4)
Quartile	all	Q1	Q2	Q3
length (n. of rounds)	all	1-4	5-6	7-8
n. bids	4318	659	2577	1082
panel A: transactions				
n. transactions	880	260	430	190
mean v_{ij}	4.22	2.62	4.41	5.97
% deals	0.4	0.28	0.43	0.51
panel B: bids placed before activation				
n. bids	2804	505	1668	631
bids per transaction	3.19	1.94	3.88	3.32
% unique/first bids	0.10	0.13	0.10	0.08
% of increasing bids	0.67	0.69	0.65	0.71
% increasing by 1 unit	0.49	0.50	0.47	0.53
% of decreasing bids	0.22	0.18	0.24	0.21
% decreasing by 1 unit	0.16	0.12	0.16	0.17
panel C: bids placed after activation				
n. bids	1514	154	909	451
bids per transaction	1.72	0.59	2.11	2.37
% unique/first bids	0.09	0.16	0.07	0.10
% of increasing bids	0.37	0.32	0.38	0.39
% increasing by 1 unit	0.28	0.19	0.28	0.32
% of decreasing bids	0.54	0.53	0.55	0.51
% decreasing by 1 unit	0.42	0.39	0.43	0.41

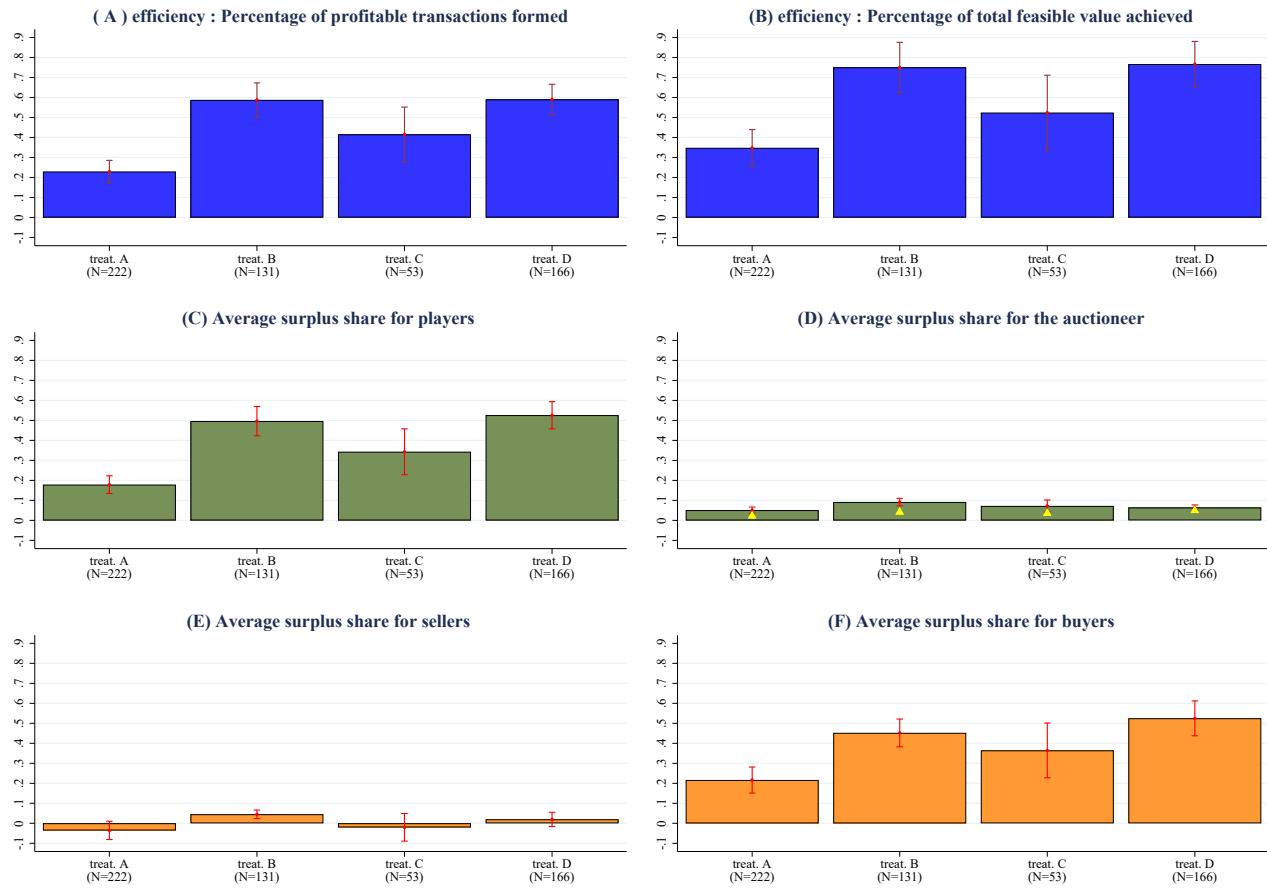
C.2 Supplementary Material

Figure A.6: Efficiency and surplus, two-sided transactions



Note : Standard errors shown in red. Yellow triangles in panel D represent the auctioneer surplus if all transactions yielded the minimal fee of 1.

Figure A.7: Efficiency and surplus, many-sided transactions



Note : Standard errors shown in red. Yellow triangles in panel D represent the auctioneer surplus if all transactions yielded the minimal fee of 1.

Figure A.8: The cost of renegotiation

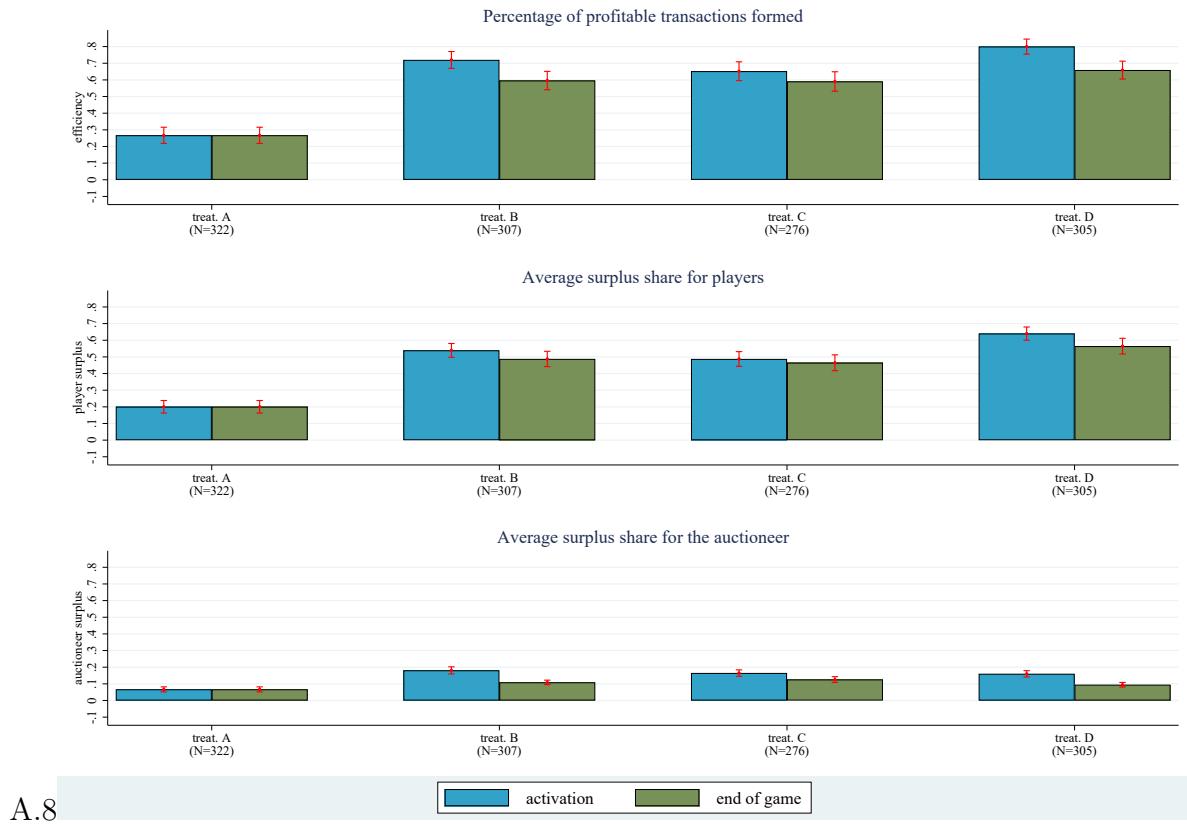


Table A.2: Number of bids

	(1)	(2)
Dep var: # $b_{ij,g}^k$ (number of bids by player k on transaction ij in game g)		
treatments	<i>All</i>	T_B, T_C, T_D
$seller_{ij}^k$	1.648*** (0.101)	1.884*** (0.132)
$buyer_{ij}^k$	2.188*** (0.143)	2.748*** (0.184)
$ v_{ij}^k $	0.065*** (0.014)	0.087*** (0.019)
S	-0.038 (0.044)	-0.042 (0.052)
T_g	yes	yes
λ_s	yes	yes
Const.	-0.770*** (0.158)	0.109 (0.132)
Obs	10,560	7,920
R-sq	0.266	0.278

Notes: OLS results reported. The dummy $seller_{ij}^k$ equals one if $v_{ij}^k < 0$, and the dummy $buyer_{ij}^k$ equals one if $v_{ij}^k > 0$ ($v_{ij}^k = 0$ omitted). We control for $|v_{ij}^k|$, game order S (from 1 to 8), treatment dummies T_g , and session-level fixed effects λ_s . Wild-bootstrapped p-values in parentheses, clustered at the unique group level. *** p<0.01, ** p<0.05, * p<0.1.

Table A.3: Number of bids, split of buyers

	(1)	(2)
Dep var: # $b_{ij,g}^k$ (number of bids by player k on transaction ij in game g)		
treatments	<i>All</i>	T_B, T_C, T_D
$seller_{ij}^k$	1.647*** (0.100)	1.886*** (0.132)
$buyer_{ij}^{k,int}$	2.231*** (0.148)	2.592*** (0.177)
$buyer_{ij}^{k,ext}$	2.140*** (0.233)	2.982*** (0.338)
$ v_{ij}^k $	0.065*** (0.014)	0.087*** (0.020)
S	-0.039 (0.043)	-0.039 (0.049)
T_g	yes	yes
λ_s	yes	yes
Const.	-0.762*** (0.144)	0.881*** (0.278)
Obs	10,560	7,920
R-sq	0.266	0.279

Notes: OLS results reported. The dummy $seller_{ij}^k$ equals one if $v_{ij}^k < 0$, $buyer_{ij}^{k,int}$ equals one if $v_{ij}^k > 0$ and $k = i$ or j , and $buyer_{ij}^{k,ext}$ equals one if $v_{ij}^k > 0$ and $k \neq i, j$ ($v_{ij}^k = 0$ omitted). We control for $|v_{ij}^k|$, game order S (from 1 to 8), treatment dummies T_g , and session-level fixed effects λ_s . Wild-bootstrapped p-values in parentheses, clustered at the unique group level. *** p<0.01, ** p<0.05, * p<0.1.

Table A.4: Magnitude of bids, split of buyers

	(1)	(2)	(3)	(4)
Dep var: $b_{ij,g}^k$ (bid by player k on transaction ij in game g)				
bids	first	last	last	last, censored
treatments	<i>All</i>	<i>All</i>	T_B, T_C, T_D	T_B, T_C, T_D
v_{ij}^{k-}	1.403***	1.326***	1.351***	1.357***
margin	[40.3%] (0.124)	[32.6%] (0.125)	[35.1%] (0.144)	[35.7%] (0.143)
$v_{ij}^{k+,int}$	0.119*** margin	0.222*** [88.1%] (0.022)	0.246*** [75.4%] (0.023)	0.377*** [62.3%] (0.016)
$v_{ij}^{k+,ext}$	0.274*** margin	0.392*** [72.6%] (0.013)	0.465*** [53.5%] (0.023)	0.456*** [54.4%] (0.024)
S	0.031 (0.022)	0.032** (0.014)	0.034* (0.018)	0.021 (0.018)
T_g	yes	yes	yes	yes
λ_s	yes	yes	yes	yes
Const.	-0.186 (0.348)	-0.544** (0.230)	-0.357* (0.205)	-0.171 (0.177)
Obs	10,560	10,560	7,920	7,679
R-sq	0.288	0.343	0.310	0.329

Notes: OLS results reported. $v_{ij}^{k-} = v_{ij}^k$ if $v_{ij}^k < 0$ and 0 otherwise.

$v_{ij}^{k+,int} = v_{ij}^k$ if $v_{ij}^k > 0$ and $k = i$ or j , and 0 otherwise. $v_{ij}^{k+,ext} = v_{ij}^k$ if $v_{ij}^k > 0$ and $k \neq i, j$, and 0 otherwise. We control for game order S (from 1 to 8), treatment dummies T_g , and session-level fixed effects λ_s . Wild-bootstrapped p-values in parentheses, clustered at the unique group level. *** p<0.01, ** p<0.05, * p<0.1.

Table A.5: Multiple transactions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A									
Dep var:	# $b_{ij,g}^k$	$v_{ij}^k - b_{ij,T}^k$	$1\{v_{ij}^k - b_{ij,T}^k = 1\}$	# $b_{ij,g}^k$	$v_{ij}^k - b_{ij,T}^k$	$1\{v_{ij}^k - b_{ij,T}^k = 1\}$	# $b_{ij,g}^k$	$v_{ij}^k - b_{ij,T}^k$	external buyers
# trans. $v_{ij}^k < 0$	-0.056 (0.079)	-0.161 (0.397)	0.008 (0.020)	-0.057 (0.100)	-0.460*** (0.108)	0.058*** (0.008)	0.201 (0.165)	-0.279*** (0.075)	0.039*** (0.010)
# trans. $v_{ij}^k > 0$	-0.018 (0.031)	-0.682 (0.407)	0.008 (0.009)	0.040 (0.096)	-0.424*** (0.085)	0.022*** (0.007)	0.014 (0.066)	-0.230*** (0.061)	0.030*** (0.009)
Observations	1,636	1,636	1,636	1,421	1,421	1,421	1,155	1,155	1,155
R-squared	0.170	0.037	0.035	0.173	0.059	0.042	0.220	0.099	0.097
Panel B									
				(10)	(11)	(12)	(13)	(14)	(15)
Dep var:	# $b_{ij,g}^k$	$v_{ij}^k - b_{ij,T}^k$	$1\{v_{ij}^k - b_{ij,T}^k = 1\}$	# $b_{ij,g}^k$	$v_{ij}^k - b_{ij,T}^k$	$1\{v_{ij}^k - b_{ij,T}^k = 1\}$	# $b_{ij,g}^k$	$v_{ij}^k - b_{ij,T}^k$	external buyers
# trans. $v_{ij}^k < 0$	-0.051 (0.108)	-0.390*** (0.108)	0.055*** (0.007)	0.190 (0.166)	-0.259*** (0.082)	0.034*** (0.011)			
# trans. $v_{ij}^k > 0$	0.042 (0.095)	-0.397*** (0.083)	0.021*** (0.007)	0.004 (0.065)	-0.212*** (0.057)	0.025** (0.009)			
% minimal selling bids	-0.160 (0.507)	-1.940*** (0.580)	0.081** (0.038)	0.435 (0.365)	-0.822*** (0.276)	0.213*** (0.053)			
Observations	1,421	1,421	1,421	1,155	1,155	1,155	1,155	1,155	1,155
R-squared	0.174	0.085	0.047	0.221	0.111	0.111	0.121	0.121	0.121

Notes: OLS results reported. # $b_{ij,g}^k$ is the number of bids placed by player k on transaction ij in game g . $v_{ij}^k - b_{ij,T}^k$ refer to the last bid on transaction ij in game g , and $1\{v_{ij}^k - b_{ij,T}^k = 1\}$ is a dummy taking value one if $v_{ij}^k - b_{ij,T}^k = 1$. The first two regressors of interest represent the number transactions with $v_{ij}^k < 0$ and $v_{ij}^k > 0$ respectively for player k in game g (other than transaction ij itself). The regressor % minimal selling bids is the share of bids such that $v_{ij}^k < 0$ and $v_{ij}^k - b_{ij,T}^k = 1$ placed by player k across all games. We control for game order S (from 1 to 8), treatment dummies T_g , and session-level fixed effects λ_s . Wild-bootstrapped p-values in parentheses, clustered at the unique group level. *** p<0.01, ** p<0.05, * p<0.1.