

Experimental Evidence on Semi-structured Bargaining with Private Information*

Margherita Comola[§] and Marcel Fafchamps[¶]

September 9, 2021

Abstract

We conduct a laboratory experiment to study a decentralized market where goods are differentiated and evaluations are private. We implement different semi-structured bargaining protocols based on deferred acceptance, and we compare their performance to the benchmark scenario of a sealed-bid auction. We show that bargaining dramatically improves efficiency, mainly to the benefit of players rather than the silent auctioneer. A protocol of unconstrained simultaneous bargaining performs best, doubling the proportion of deals (66%) relative to the benchmark (27%). This is because participants seek to reveal information through a gradual bidding-up strategy that favors bargaining environments. 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 thank Maxim Frolov and Shaden Shabayek for outstanding research assistance. We are grateful to Mariagiovanna Baccara, Francis Bloch, Yann Bramoullé, Olivier Compte, Nicolas Jacquemet, Luca Merlino, Arno Riedl, Jean-Christophe Vergnaud and numerous seminar and conference participants for helpful comments and discussions. Margherita Comola acknowledges the financial support of the ANR Dynamite and the EUR ANR-17-EURE-0001 grants.

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

[¶]Freeman Spogli Institute for International Studies, Stanford University.
fafchamp@stanford.edu

1 Introduction

A large research effort in economics focuses on understanding matching markets, that is, markets in which transactions occur on the basis of the availability and preferences of the participants, rather than on the basis of price only. In order to function, matching markets require either an institution that centralizes the matching process, or the technology to operate the market in a decentralized manner. Two distinct strands of the literature have developed around this dividing line. The first strand seeks to develop a centralized mechanism where matches are assigned on the basis of preference rankings provided by individuals (e.g., Roth 1984; Abdulkadiroglu et al. 2009; Kagel et al. 2010).¹ The second strand of the literature, to which our paper belongs, studies the functioning of decentralized markets, where participants are free to interact within the bounds of a protocol regulating transactions.

We design a laboratory experiment to study a decentralized market where goods are differentiated and information is private. Our experiment is sufficiently versatile to describe many offline and online market interactions, including settings where transactions involve more than two parties. We compare different protocols that are feasible, intuitive, and do not require agents to report their valuations. We show that protocols allowing for bargaining in a semi-structured manner dramatically improve ex-post efficiency relative to the benchmark scenario of a sealed-bid auction, and we trace the observed outcomes back to the behavioral strategies of participants.

Our experiment consists of an interactive trading floor which is populated by multiple players and goods, and is regulated by a silent, automated auctioneer. The values of the transactions, which can be positive or negative, are independent across players. The bids that subjects place can be positive (i.e., offers) or negative integers (i.e., requests). The auctioneer decides the protocol regulating how the market unfolds: this protocol varies across games and is the object of our investigation. At the end of a game, a transaction is completed only if the sum of the final bids is strictly positive, in which case subjects gain the difference between their value and

¹One practical difficulty in this literature is the need for participants to reveal their preferences over the set of alternatives, which renders the mechanism vulnerable to manipulation of reported rankings (e.g., Roth and Peranson 1999; Lee 2017; Bodoh-Creed 2020).

what they bid. The auctioneer receives a minimal fee of one unit as remuneration, and collects the leftover surplus if any.²

The experimental design relies on two key features. First, participants have symmetrically incomplete information: values and bids are kept private and no communication is allowed. This is a realistic assumption for many real-life markets where participants are uninformed about the characteristics and reservation price of others. It also allows us to abstract from considerations related to fairness: players ignore the amount of surplus generated in each transaction and how much of it is appropriated by other players but the experimenter knows.³ Second, we rule out competition over different goods: players make a gain for each successful transaction where they bid less (more) than their positive (negative) value.⁴ This allows us to abstract from congestion issues, and it gives a clear-cut prediction that only ‘profitable’ transactions for which the sum of all values is strictly positive should take place.

We investigate two main research questions: (1) what is the aggregate efficiency of different market protocols; and (2) what are the behavioral strategies used by participants. With respect to the first question, we compare four different protocols: in the sealed-bid protocol, that we use as a benchmark, players place sealed bids simultaneously and no bargaining is allowed (e.g. Block and Jackson 2007; Haeringer and Wooders 2011). The other three protocols allow for bargaining with a flavor of deferred acceptance: participants can update their bids over time and gains are determined solely by the bids held by the end of the game. We name these protocols

²For example, the transaction between i and j may have value -5 to i and +10 to j , meaning that i incurs a cost of 5 while j derives a benefit of 10 if the transaction takes place. Say that the final bids of i and j are -6 and 8 respectively. In this case, the transaction takes place since $8 - 6 = 2 > 0$. Player j gains $10 - 8 = 2$ and player i gains $-5 + 6 = 1$. The leftover surplus, computed as the difference between the bids, i.e., $8 - 6 = 2$, goes to the auctioneer.

³As the two values are independent and private, i cannot infer j 's value from her own. In our earlier example, 1 additional unit of surplus is appropriated by the auctioneer above his minimal fee of 1. The total sum of bids would have been strictly positive even if i lowered his offer (7 rather than 8) or if i raised his request (-7 rather than -6), but players do not know this. This feature is rooted in the assumption of symmetrically incomplete information in multilateral bargaining.

⁴This can be rationalized as a decentralized market where utility is separable and each agent's budget constraint is not binding over the set of goods he is interested in. For games where agents compete over a limited number of links, see Demange et al. (1986) and Comola and Fafchamps (2018).

‘semi-structured’ because they impose minimal constraints to the bargaining process, i.e., just enough to be implementable in an orderly way in a market that does not allow for information exchange. In two of these protocols, transactions are negotiated sequentially; the third unconstrained protocol allows participants to negotiate freely and simultaneously over all transactions.

Results from our experiment suggest that many profitable transactions are never formed. This matches theoretical predictions about the incomplete information penalty which induces delays and inefficiencies in trade (e.g., Myerson and Satterthwaite 1983; Williams 1987; Vincent 1989; Ausubel et al. 2002; Niederle and Yariv 2007). However, we find that all three bargaining protocols improve aggregate efficiency with respect to sealed bids. Unconstrained bargaining performs the best, doubling the proportion of completed transactions (‘deals’) from 27% to 66%. The extra surplus stemming from bargaining is mostly appropriated by players, while the margin for the auctioneer remains stably low.

We then exploit the richness of our data to trace our results back to the behavioral strategies adopted by participants. We show that most players adopt a cautious discovery strategy which consists in gradually increasing their bids until a transaction is formed and freezing afterwards. This bidding-up strategy, which overcomes the lack of information about others’ values and bids, favors bargaining treatments (especially the unconstrained one). We identify two main types of frictions impeding efficiency that relate to the player’s role and to deviations from the bidding-up strategy, respectively. Parties to a transaction are randomly selected to be a buyer – i.e., player with a positive value – or a seller – i.e., player with a negative value. We find that players acting as buyers bargain harder and longer than sellers: the implicit profit margin asked by buyers to approve a transaction is 78% against 33% for sellers. This fits the well-documented fact that selling prices are higher than buying prices – a disparity that 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; Georgantzis and Navarro-Martínez 2010; Isoni et al. 2011; Cason and Plott 2014). Our results provide new evidence on this commonly

observed empirical pattern: it is associated with unreasonably low offers, not with unreasonably high requests. In some games, we also introduce donors, that is, players who get a positive value from a transaction among third parties. This feature mimics initiatives where fundraising and charitable contributions help foster positive externalities.⁵ We show that the implicit profit margin requested by donors stands in an intermediate position between buyers and sellers. Finally, we also document that deals are delayed or prevented by a small minority of subjects who deviate from the prevalent bidding-up strategy in an attempt to appropriate a larger share of the unknown surplus.

The contribution of our paper is twofold: it advances the knowledge of bargaining games with limited structure; and it informs the design of efficient templates for online marketplaces (e.g., Carrol 2019). Unstructured bargaining is a topic of major relevance that has received surprisingly little attention by economists.⁶ Many off-line markets for differentiated goods and services allow for unregulated negotiations between parties: e.g., collectibles, art work, antiquities, second-hand cars, real estate, construction contracts, repair services, and business partnerships. The last decade has also witnessed a tremendous expansion of digital marketplaces that implement a variety of different protocols: e.g., platforms for car sharing, tendering for services, or philanthropic giving.⁷ While unstructured bargaining is ubiquitous in real-life situations, its theoretical indeterminacy has slowed down economic research. The complex strategy space of these games is difficult to model using non-cooperative game theory, and the absence of a general workhorse model has proved detrimental to empirical work.⁸ As a result, applied economists have focused their attention

⁵For an experiment incorporating this feature, see Deck and Johnson (2004).

⁶There is a sizable literature in social psychology on bargaining abilities in real-life situations – see Rubin and Brown (1975), Pruitt (2013), and Morley and Stephenson (2015).

⁷For instance, Ebay and many fundraising platforms (e.g., GoFundMe, GiveDirect, Kickstarter) currently allow for auction-format item listings. Search engines and digital social networks (including Taobao and Facebook, recently) 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 unstructured bargaining: the platform establishes the match, but users negotiate terms and prices in private.

⁸This is particularly true for bargaining games involving three or more players, such as games with donors in our experiment. Most models of multilateral bargaining have a unanimity or majority

either on auctions where bargaining is one-sided,⁹ or on highly-structured bilateral bargaining protocols where theory gives clear and testable predictions.¹⁰ Evidence on bargaining behavior with other protocols is limited, as the literature has focused instead on specific dimensions such as: fairness (e.g., Forsythe et al. 1991; Kroll et al. 2014; Galeotti et al. 2019; Luhan and Ross 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); and congestion (e.g., Kagel and Roth 2000; Che and Koh 2016; Abdulkadiroglu et al. 2017) – all dimensions that we purposefully set aside in our experimental design.

Our paper is closest to two recent studies that use observational data to analyze bargaining patterns. To the best of our knowledge, Larsen (2021) provides the only empirical analysis on the ex-post efficiency of bargaining outcomes under private information. By estimating implicit value bounds in sequential auctions of used cars, he concludes that 17–24% of profitable negotiations fail – a result that compares well to our findings on the magnitude of efficiency losses. In a related study, Bakus et al. (2020) use data on Ebay’s Best Offer platform to study the dynamics of bilateral bargaining. They document a strategy of gradual bid increases that aligns with our results but faces limitations imposed by the observational nature of their data: values cannot be observed and participants self-select into buyer or seller roles. Our experiment unifies the findings of Larsen (2021) and Bakus et al. (2020) by establishing novel behavioral evidence on the process of bidding with private information. Our results also suggest that bargaining games can be run in an decentralized and

closing rule (e.g., Baron and Ferejohn 1989; Ali 2006; Frechette 2009; Tremewan and Vanberg 2016; Agranov et al. 2021). In our setting, partial agreement is possible, making it harder to characterize theoretically (e.g., Bennet 1997; Ambrus and Lu 2015).

⁹There is a large and prominent theoretical literature on auctions (e.g., Milgrom and Weber 1982; Thaler 1988; Klemperer 1996; Bulow and Klemperer 2002; Jackson and Kremer 2006; Milgrom and Segal 2020). This literature has recently been complemented by observational studies on online auctions (e.g., Horton et al. 2017; Bodoh-Creed et al. 2021).

¹⁰Notable 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), Camerer et al. (1993), Mitzkewitz and Nagel (1993), Burrows and Loomes (1994), Güth et al. (1996), Güth and Van Damme (1998), Kagel and Wolfe (2001), Srivastava (2001), Johnson et al. (2002), Croson et al. (2003), and Kriss et al. (2013).

possibly anonymous way online, with just a minimal amount of structure to keep information private. This guarantees a profit margin for the platform, and it ensures that players do not need to report their values, thereby eschewing the thorny issue of manipulation of reported preferences.

This paper is organized as follows: Section 2 illustrates the experimental design. Sections 3 to 5 describe our results on: aggregate efficiency; bidding strategies by player’s role; and the dynamics of the unconstrained bargaining process, respectively. Section 6 concludes. Online Appendix A details the experimental protocol and the instructions to participants. Online Appendix B presents ancillary results.

2 Experimental design

In this section we describe the experimental design. We start by presenting the general features of the game. We then proceed to describe the different experimental treatments. Online Appendix A illustrates the details of the interface and its visual layout, and it reports the instructions for participants.

2.1 General features

Values and Bids

Subjects play games in groups of 6. The composition of each group remains constant across games. Each player k has a vector of values $v^k = [v_{ij}^k]$, where v_{ij}^k is her value for transaction ij . When the transaction involves the player herself, i.e., when $k = i$ or j , the value v_{ij}^k can be positive, null or negative. A positive value represents a benefit for the player if the transaction occurs, while a negative value represents a cost – e.g., production cost or reservation value. When the transaction does not involve the player herself, i.e., when $k \neq i, j$, the value must be null or positive ($v_{ij}^k \geq 0$). All values are independently drawn (in respect of the constraints above) and are uncorrelated across players *and* transactions, and players know that.¹¹

¹¹This design abstracts from complications related to substitute goods. Also, we do not require partitioning players into types (e.g. buyers and sellers, men and women): in the current version

As the game unfolds, players have the opportunity to place bids. Formally let $b_{ij,t}^k$ denote the bid of player k on transaction ij at a given point in time t during the game. When the transaction involves the player herself, i.e., when $k = i$ or j , that player can bid either a positive integer (which represents an offer) or a negative integer (which represents a request). When a transaction involves other players the bids can only be null or positive, i.e. $b_{ij,t}^k \geq 0$ when $k \neq i, j$.¹² We place no limit on the amount a player can bid.

Gains

Let us call $b_{ij,t} = \sum_k b_{ij,t}^k$ the *total bid* on transaction ij at a given point in time t . In our game, a transaction is formed as long as the total bid is strictly positive, that is, if $b_{ij,t} > 1$. This requires that the offers strictly exceed the requests while covering a minimum fee of 1 unit for the silent auctioneer (which is played by the computer).¹³

The gain of player k for transaction ij at time t is computed as $g_{ij,t}^k = v_{ij}^k - b_{ij,t}^k$ if $b_{ij,t} \geq 1$; it is 0 otherwise. In our game, players with a negative value make a gain if the transaction occurs and their request exceeds the cost. On the other hand, players with a positive value make a gain if the transaction occurs and their offer is smaller than their benefit. If players offer more than their value or request less than their cost, they make a loss (i.e., negative gain). If players bid exactly their value, they do not receive any gain, as in a Vickrey auction.

Importantly, while negotiations are open, gains are provisional because bids and transactions can still change. In our game, subjects only get a monetary reward for the transactions which take place by the end of the game, that is, if the *final* total bid $b_{ij,T} > 1$, where T indicates the game's ending.

of the protocol, each player can simultaneously have positive and negative values for different transactions (e.g., Agranov and Elliot 2021). However, this design can be restricted to accommodate 'gendered' players without loss of generality.

¹²This ensures that a player cannot prevent others from trading, e.g., out of spite or envy.

¹³A player need not give an implicit consent (in the form of a bid) for a transaction that involves her – others may bid enough to form it on their own. But by making a sufficiently large negative bid, a player can *de facto* preclude a transaction involving her.

Visual layout

Much attention went into designing an interface that presents all the relevant information to subjects in a compact yet intuitive manner. Players appear as nodes arranged in a hexagon, and transactions are represented as links across nodes. At each point of the game the transactions that are currently formed – i.e., for which the total bid is 1 or above – are displayed as a thick solid line. The transactions that are not currently formed are displayed as a thin dotted line. A player can open a dialog box above each link reporting the current state of play, with the *value* of this transaction to her and the *gain* updated in real time based on her current bid. Players never observe the values or bids of others. This design, which is the result of careful research and development, presents all the relevant information in an instantaneously available and engaging format. For further details we remand to Online Appendix A.

2.2 Treatments

Each group plays 8 games in sequence. Each of these games is played under one of four experimental treatments identified as T_A, T_B, T_C and T_D . These treatments introduce different market protocols that we describe in what follows.

Sealed-bid auction

All games start with a simultaneous sealed-bid auction. Players have unlimited time to enter or revise bids in the dialog boxes associated with the transactions. Once all players stop bidding, this phase ends. The silent auctioneer then sums up the bids on each transaction to determine whether the sum is strictly positive – if so, that transaction is formed. Players then observe their own gains (for each transaction and in total).

In Treatment T_A , this first phase marks the end of the game. Thus, under T_A the game only has one round, that we name round 0. In other treatments, the game continues with one or more rounds of 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. This ensures that players take round 0 equally seriously in all treatments. The sealed bid auction represents a formidable challenge for subjects who have no information on the values of others. This forces players to abstract from strategic or altruistic motives, as they can only bid on the basis of their own value. We use the sealed-bid auction as a benchmark to judge the efficiency gains that can be achieved by allowing discovery via bargaining, to which we now turn.

Sequential bargaining

In Treatments T_B and T_C , the game continues with several rounds of sequential bargaining with deferred acceptance (e.g., Gale and Shapley 1962). In these two treatments, all bargaining action is focused on one transaction at a time. The difference between T_B and T_C lies in the stopping rule that we discuss below.

In both treatments T_B and T_C the bargaining phase is divided into multiple rounds and each round is divided into turns. In a given round, each turn is devoted to bidding for one specific transaction.¹⁴ The sequence of play within a turn is as follows: when the turn begins, the transaction currently up for auction is highlighted. The bids placed in round 0 serve as start-up bids for the first round or bargaining; in subsequent rounds, the start-up bids are the last bids placed in the previous round. Then the bargaining floor opens and all players can place, revise, or drop a bid on this specific transaction as they wish. 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.¹⁵ At the end of the turn, the computer calculates whether the sum of bids is strictly positive or not, and the result is displayed in real time.¹⁶ The game then moves to the next transaction. When all transactions are

¹⁴There 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.

¹⁵The 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.

¹⁶If the sum becomes positive, the transaction is formed on the screen, i.e., 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

visited, the round is complete and another round begins.

Both sequential treatments impose some structure to the bargaining game by restraining players into independent sub-games but for a given transaction agents are able to bargain freely. The stopping rule differentiates Treatments T_B and T_C . Treatment T_B has a stopping rule based on bids. In this treatment, the game ends in one of two ways: either because at the end of one 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.¹⁷ At the end of the game, the last bids are retained, the final gains are determined, and players are informed of their gains, as in treatment T_A .

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 *transactions* formed from one round to the next – even if there were changes in bids. Treatment T_C aims at speeding up bargaining by curtailing the time players can take to experiment with frivolous bids and attrition strategies.

In both sequential treatments, the number of rounds may vary from 2 to 8.

Unconstrained bargaining

Treatment T_D is characterized by an unconstrained bargaining floor in which players can update bids simultaneously on all transactions, with no specific sequencing imposed and minimal structure. As in Treatments T_B and T_C , this treatment also begins with bids from the sealed-bid phase. As the bargaining phase unfolds, a player can update any of her bids in any order. The silent auctioneer updates the information on transactions and gains in real time. This allows players to see which

solid to a dotted line.

¹⁷In bargaining games with a predetermined end time, it is common to observe a bunching of offers and counter-offers 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. Players are informed about the forcible ending rules and their probabilities.

transactions would be completed if the game were to end at that point. We place no limit on the number of bids: the game stops when there has been no change in bids for a set amount of time since the last bid placed by any player.¹⁸ Thus, under T_D , games only have two rounds: the sealed bid auction round 0, and one round of simultaneous bargaining.

2.3 The value matrix

In our experiment each group plays 8 games by cycling across 8 different value matrices $V = [v^1, \dots, v^6]$ with the following features. In all matrices, 10 transactions appear on the screen and can be bid on.¹⁹ Values are set through independent random draw subject to two constraints: 1) we impose $v_{ij} \neq 0$ (the total value across all players is either strictly positive or strictly negative), which marks a clear efficiency criteria; 2) we set $v_{ij}^k \neq 1$ so that no player is indifferent between not making any offer or making a minimal offer of 1.

In all eight matrices, values for players who are directly involved in a transaction ($k = i$ or j) are randomly drawn in the interval $[-10, +10]$.²⁰ The values for players who are not part of the transaction ($k \neq i, j$) are set to zero in four matrices (that we call ‘without externalities’). The remaining four matrices are ‘with externalities’, in the sense that we randomly draw the values for $k \neq i, j$ players in the $[0, +10]$ interval.

As a result of randomization, the number of transactions with total value $v_{ij} > 0$ varies between 5 and 8 across all matrices – the others 2 to 5 transactions have $v_{ij} < 0$.

¹⁸At the beginning of the phase, the wait time is 20 seconds from the last bid 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 time the last bid was entered.

¹⁹Out of the 15 possible pairings of 6 players, only 10 links appear on the screen with a dialog box. This is done to keep the length of the session within reason.

²⁰Chatterjee and Samuelson (1983) and Myerson and Satterthwaite (1983) also impose symmetric uniform values for buyers and sellers.

2.4 Implementation

The laboratory experiment was conducted in 2019 at the Paris School of Economics (Paris, France). We ran six experimental sessions with a total of 22 unique groups of 6 players. Group composition remains unchanged throughout a session. Each group plays 8 games, corresponding to the 8 payoffs matrices – four with externalities and four without externalities. The identity of players is reshuffled from one game to another.²¹

Each of the four matrices without externalities gets randomly paired with treatments T_A to T_D . We do the same for the four matrices with externalities. Thus, each groups plays two matrices (one with and one without externalities) for each treatment, in a random order which is group-specific.²² We therefore observe $22 \cdot 8 = 176$ unique games equally distributed among the four treatments (44 unique games per treatment).

Sessions unfold as follows. After reading the instructions, players run three trial games to get used to the interface.²³ After the trial games, players answer a quiz to test their understanding of the instructions. The quiz is corrected immediately afterwards on the blackboard. Once this is completed, players run a social value orientation task (e.g., Murphy et al. 2011) and then proceed to the main part of the experiment, which comprises the eight games. At the end of the games, players complete a questionnaire with socio-demographic information and comprehension feedback, and they proceed to payment. To determine subjects' earnings, we randomly draw 2 games out of 8 for each group, and players receive the monetary

²¹Each 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 spillover from one game to the next.

²²For 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.

²³Players 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 have 7 links instead of 10. All other features are the same as in the main games.

equivalent of their gain at the end of these 2 games.²⁴ The average earnings are 24.5 euros for about 2 hours in the laboratory.

3 Efficiency

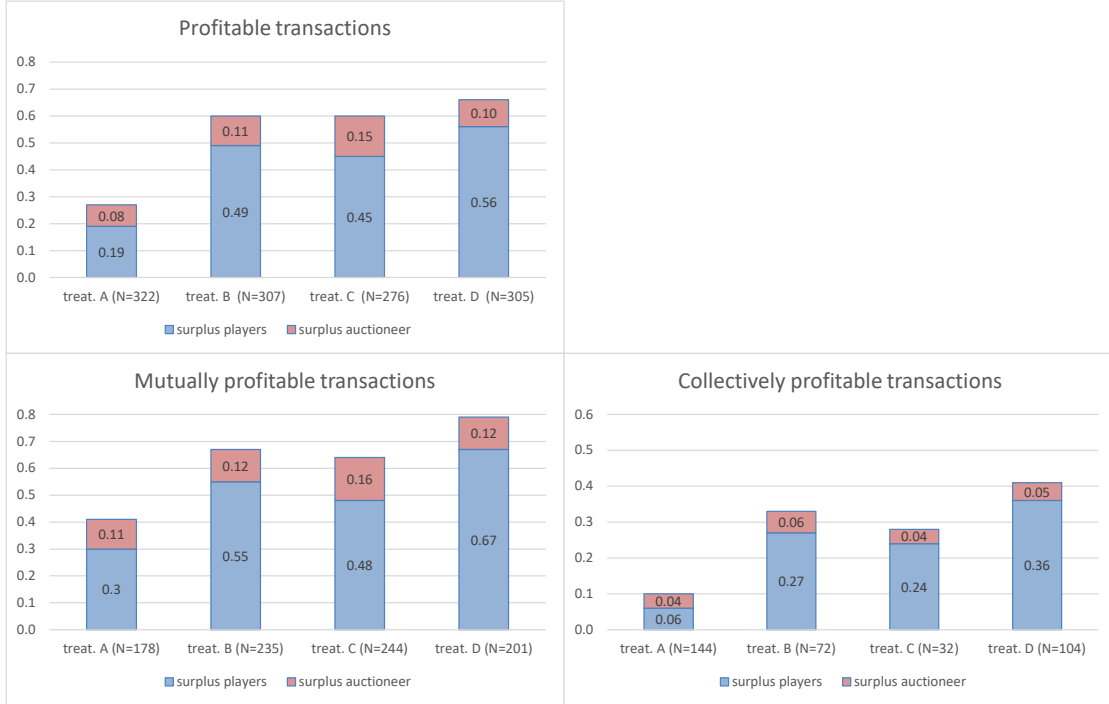
We start by examining our findings about market efficiency. We classify transactions by their value. Transaction ij is defined as *profitable* if its total value over all players is strictly positive, i.e. if $v_{ij} = \sum_k v_{ij}^k > 0$. We further distinguish between two types of profitable transactions. A transaction ij is called *mutually profitable* if the combined value of that transaction for i and j is strictly positive, i.e., if $v_{ij}^i + v_{ij}^j > 0$.²⁵ It is also possible for a transaction to be profitable ($v_{ij} > 0$) without being mutually profitable ($v_{ij}^i + v_{ij}^j \leq 0$). This requires that the combined benefit for players $k \neq i, j$ is large enough to compensate for the loss of players i and j , i.e., it requires that $\sum_{k \neq i, j} v_{ij}^k > -(v_{ij}^i + v_{ij}^j)$. We call these transactions *collectively profitable*. Since, by construction, every value matrix has 10 transactions and we have $N = 176$ unique games, we observe $N = 1760$ transactions, divided equally across the four treatments – totaling 440 transactions per treatment. Within a given treatment, however, transactions are unequally distributed by type because value matrices are randomly assigned to treatments.

Figure 1 illustrates our main result on aggregate efficiency. For each transaction type – i.e., profitable, mutually profitable, and collectively profitable – the height of the bar depicts the percentage of transactions formed at the end of the game, by treatment. Each bar is then split into two components: the *surplus for players* represents the share of total value appropriated by the players by the end of game, and it is defined as $g_{ij,T}/v_{ij}$ where $g_{ij,T} = \sum_k g_{ij,T}^k$. Conversely, the *surplus for the auctioneer* represents the share of total value appropriated by the auctioneer and it is defined as $(v_{ij} - g_{ij,T})/v_{ij}$. The surplus of players and the auctioneer are zero for

²⁴If some players had incurred negative earnings, we would have subtracted these losses from their show-up fee. We had no such cases.

²⁵It is still possible that other players hold positive values (i.e. that $\sum_{k \neq i, j} v_{ij}^k > 0$) but the transaction is classified as mutually profitable as long as $v_{ij}^i + v_{ij}^j > 0$.

Figure 1: Efficiency and surplus



all the transactions that are not formed by the game’s end and they mechanically sum up to one when a transaction occurs.

Figure 1 shows that only a small fraction of profitable transactions occur in the sealed-bid auction of T_A . But the overall efficiency rate increases dramatically when bargaining is possible, i.e., in treatments T_B, T_C and T_D . This is true across all profitable transactions and separately for mutually and collectively profitable transactions too. Theory suggests that private information has a high cost in terms of efficiency (e.g., Myerson and Satterthwaite 1983; Williams 1987; Vincent 1989; Ausubel et al. 2002). This prediction has so far been confirmed by the evidence available from observational data.²⁶ This also what we find in our experimental setting: a large proportion of profitable deals never occur.

²⁶Larsen (2021) estimates an efficiency loss of 17-23% for two-sided uncertainty, while Ambrus et al. (2018) find an efficiency loss of 14% in ransom negotiations.

The penalty is especially noticeable for collectively profitable transactions, which are formed only 10% of the time in treatment T_A and 41% of the time in treatment T_D . This can be attributed to two factors. First, collectively profitable transactions require the coordination of more players. Secondly, on average, their total value is lower than the value of mutually profitable transactions.²⁷ We find evidence in our data that transactions with lower values are less likely to be formed, across all treatments and transaction types.²⁸ This finding is consistent with experimental evidence showing that deal rates increase with the size of the surplus (e.g., Camerer et al. 2019).

When we examine how the surplus is split between players and the auctioneer, we see that the auctioneer’s surplus remains stable and low across all treatments and transaction types. Players receive most of the efficiency gains stemming from an increase of the percentage of deals in the bargaining treatments. Subjects are surprisingly efficient at not “leaving money on the table” for the silent auctioneer.²⁹ But, in their effort to achieve this, players may bid too little and miss some profitable transactions, a point that we revisit in detail below.

Our results show that the unconstrained bargaining protocol of treatment T_D delivers the best outcomes in all cases. In T_D , subjects are forced to follow multiple transactions simultaneously, which in principle increases complexity and raises cognitive load. In contrast, T_B and T_C corral subjects into a systematic sequence of bargaining sub-games, which is cognitively easier but slower, and it may make it difficult to maintain concentration and motivation. Furthermore, information discovery about the valuations of others is facilitated in the unconstrained bargaining protocol: T_D reveals information faster since formed transactions are displayed on the screen as soon as the sum of bids is positive; in contrast, in T_B and T_C subjects have to wait until the end of a round to find out whether their bids formed the transaction

²⁷The independent randomization scheme across players implies that, by construction, the expected value of $v_{ij} > 0$ is smaller when players i and j incur in a cost ($v_{ij}^i + v_{ij}^j \leq 0$).

²⁸The deal rate for profitable transactions switches from 32% to 74% for values below and above median, respectively.

²⁹A striking 51% of the transactions occurring in the lab have a sum of bids $b_{ij,T} = 1$, which is the minimum. Another 22% have $b_{ij,T} = 2$.

or not. Results from Figure 1 suggest that the information discovery benefit of T_D more than compensates for the reduction in cognitive load provided by the sequential bargaining processes.

We end this discussion of efficiency by comparing the two sequential bargaining protocols. To recall, in treatment T_B the stopping rule is based on changes in bids, while in T_C it is based on changes in transactions. Protocol T_C is intended to limit bidding wars, but it is also possible that a faster stopping rule reduces the scope for information discovery, thereby leading to fewer transactions.³⁰ Our results show that speeding up play by applying a looser convergence rule ends up reducing players' gains, which are slightly higher in T_B than T_C . Protocol T_C nonetheless cuts game time in half and games converge naturally, without reaching their round limit (See Online Appendix B.1). It is therefore conceivable that T_C may be preferable in practice to T_B , depending on the context. The magnitude of the difference between the two, however, is small compared to the difference with either T_A – which is worse – or T_D – which is better.

In Online Appendix B.1 we report detailed statistics on efficiency and surplus by treatment and transaction type, and we also discuss the result of a regression analysis that confirms the results above. Additional evidence reassures the reader that all treatments are comparable in the sealed-bid phase,³¹ and that the percentage of non-profitable transactions that are formed is consistently negligible.

4 Bidding Strategies and Roles

The efficiency of a market protocol ultimately depends on the bidding strategies adopted by players. To orient the analysis in this direction, in this Section we first

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 penalizes bidding-up strategies that are too gradual.

³¹Subjects never know in round 0 whether it is final or whether it will be followed by the other rounds of treatments T_B, T_C and T_D . Hence behavior in round 0 should be comparable across all treatments, which appears to be the case.

check whether bids are consistent with a basic rationality criterion. We then analyze the differences in bidding pattern depending on the role ascribed to the player for a given transaction – i.e., buyer, seller, or donor.

4.1 Losing bids

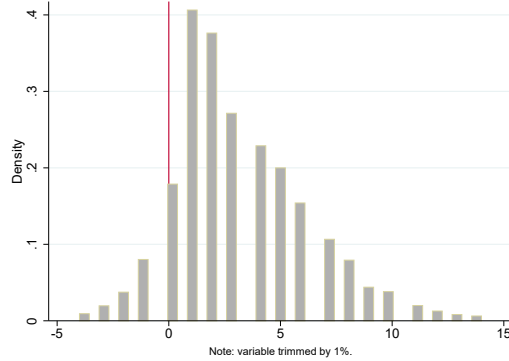
We observe 13,331 unique bids across all players and treatments combined (1,156 for treatment T_A , 3,722 for T_B , 2,451 for T_C and 6002 for T_D).³² We first examine whether subjects occasionally place losing bids, that is, bids where $b_{ij,t}^k > v_{ij}^k$, meaning that they would result in a loss for player k if the game were to stop at that moment t . Overall, only 1003 bids (7.5% of all bids) are losing bids. Note that these bids need not turn into a loss, because bids and deals are revised as bargaining unfolds. We plot in Figure 2 a histogram of $v_{ij}^k - b_{ij,t}^k$ across all bids and treatments combined. It shows that for the large majority of bids, $b_{ij,t}^k > v_{ij}^k$ by a wide margin. This is reassuring, as it suggests that players have little difficulty in understanding the way gains are calculated.

Interestingly, only 86 losing bids (out of 1003) are cases in which a buyer over-bids. In contrast, 321 are cases where the seller under-bids, that is, does not ask enough to cover her cost. The large majority of these losing bids (596) are situations in which a player places a positive bid on a transaction that has zero value for her. Since doing so may help others form a transaction, this could indicate generosity towards other players.³³

³²To 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 ($r = 0$) 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 unique bids placed in the bargaining round ($r = 1$).

³³Regression analysis, not shown here to save space but available upon request, confirms that negative gains are overwhelmingly associated to transactions with a negative and especially zero value. Losing bids are no less likely in later rounds or games, suggesting that they are not just early mistakes.

Figure 2: Histogram of $v_{ij}^k - b_{ij,t}^k$



4.2 Buyers, Sellers and Donors

We now examine whether players display different bidding patterns depending on their ascribed value for a given transaction. To do so, we classify players into four roles: players with no value ($v_{ij}^k = 0$), *sellers* ($v_{ij}^k < 0$ & $k = i$ or j), *buyers* ($v_{ij}^k > 0$ & $k = i$ or j), or *donors* ($v_{ij}^k > 0$ & $k \neq i, j$).

Our unit of observation is now player k for transaction ij in game g . Since there are 10 transactions per game, the number of observations is $N = 132 \cdot 8 \cdot 10 = 10,560$. Our data show that, in 4,711 out of 10,560 observations, the player has placed at least one bid.³⁴ As already noted in Section 4.1, a small minority of players bid on transactions that have no value to them, but otherwise most observed bids are placed by players for whom $v_{ij}^k \neq 0$.³⁵

In Table 1 we report descriptive statistics on the mean value and gain at the end of game, by role and treatment. Results show that sellers gain consistently less than buyers and donors: the average gain for sellers is 0.07, compared to 2.77 for buyers and 1.44 for donors. This is true across all treatments, including the sealed-bid auction T_A . This result cannot be imputed to individual heterogeneity since,

³⁴In 2137 cases, only one bid was placed (2137 unique bids in total) and in 2574 cases multiple bids were placed (11194 unique bids in total).

³⁵Overall, we have 5,852 observations out of 10,560 where $v_{ij}^k = 0$, and for only 399 of them the player placed at least one non-zero bid. Conversely, for the 4708 observations where $v_{ij}^k \neq 0$ we observe at least one bid in 90% of cases – and nearly all players who fail to bid are donors.

by design, roles are randomly assigned for a given transaction and players occupy simultaneously multiple roles within the same game. It also cannot be driven by the distribution of values v_{ij}^k across players since, by design, values are randomized in the same way across roles, as Table 1 displays – albeit with positive values for buyers and donors, and negative values for sellers.

Table 1: Gain by role

	treatment	N	(1) mean (v_{ij}^k)	(2) mean ($g_{ij,T}^k$)
$v_{ij}^k = 0$	<i>All</i>	5852	0	0.00
	T_A	1262	0	-0.00
	T_B	1501	0	0.01
	T_C	1681	0	-0.01
	T_D	1681	0	0.01
$v_{ij}^k < 0$ & $k = i$ or j	<i>All</i>	1694	-5.55	0.07
	T_A	475	-5.81	-0.09
	T_B	389	-5.27	0.20
	T_C	402	-5.37	0.05
	T_D	428	-5.68	0.14
$v_{ij}^k > 0$ & $k = i$ or j	<i>All</i>	1562	5.99	2.77
	T_A	319	5.94	1.65
	T_B	422	5.96	3.01
	T_C	431	6.18	2.80
	T_D	390	5.83	3.41
$v_{ij}^k > 0$ & $k \neq i, j$	<i>All</i>	1452	6.02	1.44
	T_A	584	6.03	0.82
	T_B	328	5.99	2.04
	T_C	126	6.18	1.19
	T_D	414	5.96	1.91

We now investigate the extent to which these discrepancies in average gains can be imputed to the differential frequency and magnitude of bids across roles. We start

by focusing on the frequency of bids and estimate a regression of the form

$$NBids_{ij,g}^k = \beta_0 + \beta_1 D_{ij}^{k-} + \beta_2 D_{ij}^{k+} + \beta_3 D_{ij}^{k,out} + \beta_4 S + \beta_5 T_g + \lambda_s + \varepsilon_{ij,g}^k \quad (1)$$

where $NBids_{ij,g}^k$ is the total number of bids placed by player k on transaction ij by the end of game g . Three dummy variables capture the player's role: seller ($D_{ij}^{k-} = 1$ if $v_{ij}^k < 0$ and $k = i$ or j); buyer ($D_{ij}^{k+} = 1$ if $v_{ij}^k > 0$ and $k = i$ or j); or donor ($D_{ij}^{k,out} = 1$ if $v_{ij}^k > 0$ and $k \neq i, j$). The omitted category is for $v_{ij}^k = 0$. We control for game order S (from 1 to 8), treatment dummies T_g , and session-level 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.

Coefficient estimates for model 4.2 are reported in Table 2, columns (1) to (5). Results show that buyers place consistently fewer bids than sellers or donors. This pattern is driven by the three bargaining protocols where subjects have an opportunity to revise bids. In Columns (6) to (10) we include the absolute magnitude of the value $|v_{ij}^k|$ as additional regressor, to ensure that the observed pattern is not driven by differences in values across roles. Similar results are obtained. Overall, Table 2 suggests that sellers are less prone to haggling when given the opportunity, but they are more likely to stick to their initial request.³⁶

Next we examine the magnitude of the bids by estimating the model

$$b_{ij,g}^k = \beta_0 + \beta_1 v_{ij}^{k-} + \beta_2 v_{ij}^{k+} + \beta_3 v_{ij}^{k,out} + \beta_4 S + \beta_5 T_g + \lambda_s + \varepsilon_{ij,g}^k \quad (2)$$

where $b_{ij,g}^k$ represents a given bid (first or last) placed by player k on transaction ij in game g . The three regressors of interest represent the player's role. The first regressor v_{ij}^{k-} is the value of transaction ij if k is a seller, i.e., $v_{ij}^{k-} = v_{ij}^k$ if $v_{ij}^k < 0$ and $k = i$ or j ; it is 0 otherwise. The second regressor 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 $k = i$ or j , and 0 otherwise. The third regressor $v_{ij}^{k,out}$ is the value of ij if player k is a donor, i.e., when $k \neq i, j$. We control for game

³⁶Results from a Tobit regression yield the same conclusions. They are available upon request.

Table 2: Number of bids

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	<i>All</i>	T_A	T_B	T_C	T_D	<i>All</i>	T_A	T_B	T_C	T_D
D_{ij}^{k-}	2.010*** (0.095)	0.872*** (0.019)	2.196*** (0.091)	1.825*** (0.084)	3.154*** (0.304)	1.647*** (0.100)	0.767*** (0.027)	1.893*** (0.144)	1.703*** (0.109)	2.138*** (0.325)
D_{ij}^{k+}	2.622*** (0.143)	0.765*** (0.030)	2.655*** (0.129)	2.368*** (0.144)	4.455*** (0.406)	2.231*** (0.148)	0.658*** (0.039)	2.311*** (0.155)	2.228*** (0.166)	3.411*** (0.426)
$D_{ij}^{k,out}$	2.533*** (0.275)	0.620*** (0.039)	2.525*** (0.186)	2.240*** (0.301)	5.025*** (0.849)	2.140*** (0.233)	0.511*** (0.036)	2.181*** (0.207)	2.101*** (0.293)	3.958*** (0.700)
$ v_{ij}^k $						0.065*** (0.014)	0.018*** (0.004)	0.058*** (0.018)	0.023 (0.021)	0.179*** (0.043)
S	-0.039 (0.043)	0.001 (0.005)	-0.077*** (0.026)	-0.005 (0.022)	-0.140 (0.149)	-0.039 (0.043)	0.001 (0.005)	-0.076*** (0.026)	-0.005 (0.022)	-0.139 (0.149)
T_B	1.144*** (0.079)					1.151*** (0.079)				
T_C	0.907*** (0.097)					0.909*** (0.098)				
T_D	1.971*** (0.243)					1.975*** (0.244)				
λ_s	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
Const.	-0.756*** (0.144)	-0.010 (0.038)	0.464*** (0.119)	0.086 (0.108)	0.927 (0.795)	-0.762*** (0.144)	-0.011 (0.038)	0.459*** (0.120)	0.084 (0.108)	0.923 (0.794)
Obs	10,560	2,640	2,640	2,640	2,640	10,560	2,640	2,640	2,640	2,640
R-sq	0.264	0.591	0.428	0.498	0.233	0.266	0.596	0.430	0.498	0.238

Notes: Wild-bootstrapped p-values in parentheses, clustered at the unique group level. *** p<0.01, ** p<0.05, * p<0.1.

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, as before.

Table 3 presents the regression results from Equation 2. Column (1) only includes initial bids placed on transaction ij during game g . Column (2) to (6) only includes final bids, for all treatments combined and for each one separately.³⁷ Column (7) only includes final bids by players who placed more than one bid for transaction ij in game g . Results show that asking bids placed by sellers respond slightly more than 1 for 1 to their (negative) values. For instance, in column (2) for all treatments combined, sellers demand on average 1.326 to sustain a negative value (cost) of -1 . If all these transactions were formed (which is obviously not the case), they would yield a modest profit margin of 32.6% for sellers. Buyers' bids respond much less than 1 for 1 to (positive) value: in Column (2) we remark that buyers offer on average 0.222 for a positive value of $+1$. If all these transactions were formed, they would yield an average margin for buyers of $1 - 22.2\% = 77.8\%$ – which is much higher than the margin requested by sellers. Donors are somewhere in between, offering on average a larger proportion of their value than buyers. Based on the coefficients reported in Column (2), donors would get a profit margin of $1 - 39.2\% = 60.8\%$ if all these bids resulted in deals. These findings are internally consistent with the evidence that sellers frequently request minimal gains: in our data, one third of all bids placed by sellers are in fact equal to $v_{ij}^k - 1$. This shows that, differently from buyers, a sizable share of sellers choose to fix the most conservative reservation price.³⁸ These results suggest that players bid less aggressively when placed in the position of seller: they tend to place fewer bids, and bids which imply a lower profit margin for themselves. In contrast, buyers and donors bid more often and their bids aim to achieve a higher profit margin. As a result, sellers make smaller gains by the end of the game.

³⁷The difference between columns (1) and (2) is driven by players placing multiple bids on transaction ij within game g : if a player places one or no bid, the dependent variable stays the same.

³⁸Bids by sellers equal to $v_{ij}^k - 1$ are not concentrated in early rounds but equally spread across all rounds and games. A similar behavior is not observed for buyers or donors.

Table 3: Magnitude of bids

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
bids	first	last	last	last	last	last	last
treatments	all	all	T_A	T_B	T_C	T_D	all
sample	all	all	all	all	all	all	non-zero
v_{ij}^{k-}	1.403*** (0.124)	1.326*** (0.125)	1.275*** (0.092)	1.487*** (0.199)	1.341*** (0.185)	1.249*** (0.082)	1.129*** (0.070)
v_{ij}^{k+}	0.119*** (0.022)	0.222*** (0.021)	0.128*** (0.020)	0.205*** (0.034)	0.283*** (0.027)	0.244*** (0.027)	0.370*** (0.033)
$v_{ij}^{k,out}$	0.274*** (0.013)	0.392*** (0.019)	0.271*** (0.014)	0.457*** (0.035)	0.444*** (0.067)	0.480*** (0.035)	0.645*** (0.034)
S	0.031 (0.022)	0.032** (0.014)	0.041 (0.044)	0.032 (0.043)	0.026 (0.038)	0.033 (0.037)	0.058* (0.029)
T_B	-0.005 (0.133)	0.258** (0.107)					
T_C	-0.124 (0.219)	0.222 (0.210)					0.359 (0.280)
T_D	-0.042 (0.138)	0.417*** (0.091)					0.304 (0.288)
λ_s	yes	yes	yes	yes	yes	yes	yes
Constant	-0.186 (0.348)	-0.544** (0.230)	-0.416 (0.304)	-0.171 (0.281)	-0.275 (0.250)	-0.431 (0.383)	-1.021*** (0.302)
Observations	10,560	10,560	2,640	2,640	2,640	2,640	2,574
R-squared	0.288	0.343	0.502	0.397	0.198	0.443	0.576

Notes: Wild-bootstrapped p-values in parentheses, clustered at the unique group level.

*** p<0.01, ** p<0.05, * p<0.1

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 that is substantially higher than the maximum amount offered by buyers – the so-called willingness to accept (WTA) - willingness to pay (WTP) gap. We refine this finding showing that this may happen not because

the ask price is unreasonably high, but because the *offer* price is unreasonably low. Furthermore, the randomization of subjects across roles and values demonstrate that this effect is entirely behavioral – it is not due to self-selection in buyer or seller role. This suggests that this empirical regularity is rooted in common real-life practices that make subjects more socially accustomed to haggle when buying than when selling.

5 The dynamics of unconstrained bargaining

This section explores the aggregate dynamics of the bidding process in T_D . We focus on T_D because the unconstrained bargaining protocol provides ideal data for this purpose. In Online Appendix B we show that our main findings for T_D extend to sequential bargaining protocols T_B and T_C as well.

5.1 Bidding up

At a first glance one can imagine two simple heuristics to guide bargaining in an environment where information is symmetrically private and acceptance is deferred: bidding down or bidding up. In a bidding-down strategy, player k initiates bargaining by bidding high – e.g., the reservation price $v_{ij}^k - 1$. In a sealed-bid auction, this choice is quite conservative and yields low gains to players. But when renegotiation is possible, this could be an efficient way to jump-start the bargaining process since it reveals upfront all profitable transactions without disclosing private values. Players can then subsequently lower their offer or increase their request to capture a larger share of the surplus from the silent auctioneer. The main drawback of this bidding-down strategy is that players have an incentive to defect unilaterally.

Conversely, players can follow a bidding-up strategy whereby they initially make a low bid and gradually raise it. If pursued by all players for a sufficiently long time, this process should eventually reveal all the profitable transactions while simultaneously minimizing the share of the surplus that goes to the auctioneer. The main drawbacks from the bidding-up strategy is that: (1) it requires multiple sequential bids, which

can become a disadvantage when bargaining is constrained or time-consuming; and (2) it fails to reveal all profitable transactions if some players refrain from raising their bid, e.g., in an attempt to extract more surplus.

We do not find evidence of a widespread bidding-down strategy in our data: bids equal to $v_{ij}^k - 1$ are relatively infrequent and mostly placed by sellers.³⁹ In contrast, our data support a generalized bidding-up strategy, the evidence for which we now discuss in detail. In treatment T_D we observe 4,995 unique bids placed over 440 transactions during the simultaneous bargaining round. We define a *bid run* as the ordered sequence of bids placed by all players on a given transaction ij during game g , and we estimate a model of the form:

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

where the *tick* variable t_{ij} represents the order in which bids on transaction ij are placed by different players within a given bid run, and $b_{ij,t_{ij}}$ is the total bid at tick time t_{ij} .⁴⁰ 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 currently 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 any possible confound correlated with matrix structure.

To correct for the possible correlation between the length of a bid run and the pace of increase in bids, we divide bid lengths into four approximately equal quantiles, and estimate model (3) separately for each quantile.⁴¹ Our estimate of interest is

³⁹These bids are about 17% of all bids across all treatments, and the overwhelming majority (e.g. 81% of them in round 0) are placed by sellers. This is consistent with the evidence discussed in Section 4.2.

⁴⁰The 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.

⁴¹To 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 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

how the predicted values of $b_{ij,t_{ij}}$ evolves over bargaining time t_{ij} which we plot in Figure 3. Fitted values are negative throughout because many transactions never occur, and most deals are sealed by a small margin – i.e., subjects appropriate most of the surplus from the silent auctioneer (see Section 3). 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. Putting together this result with the findings discussed in Section 4.2, we know that this gradual increase in bids mostly comes from the buyer side. A gradual bidding-up strategy has also been documented by Bakus et al. (2020) on the basis of observational data, but theoretical support for this strategy is still scarce.⁴²

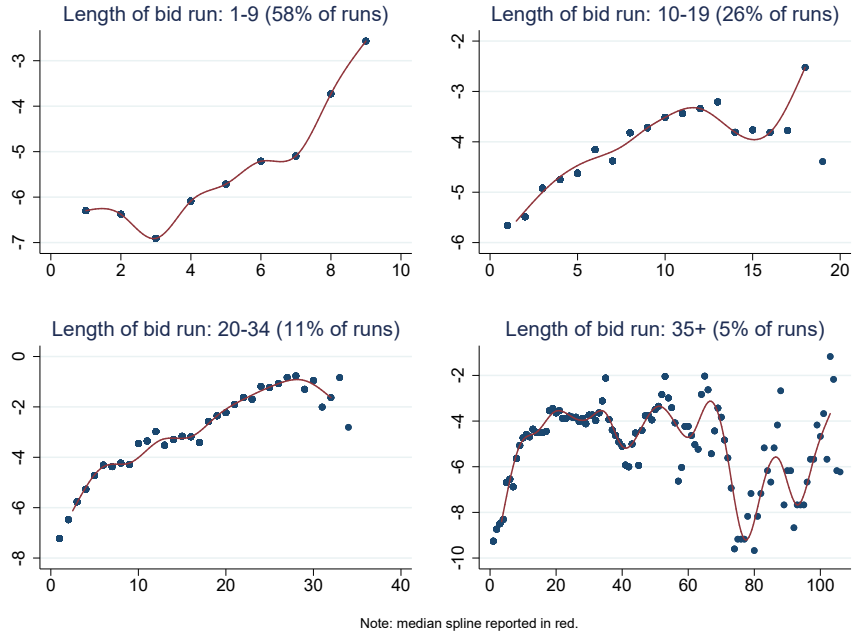
Finally, we remark an increase in noise in the upper tail of the fourth graph, which comes from bids in the upper 5% of bid run length (more than 35 ticks). This suggests that, towards the end of long runs, some players diverge from a gradual bidding-up strategy, causing the observed non-monotonicity of bidding dynamics. This behavior appears to be highly detrimental for efficiency: the average probability of ending the bid run with a deal is remarkably stable in the first three quantiles (i.e., between 46% and 48%). But it drops down to 29% in the last quantile, for bid runs with more than 35 ticks. In the next subsections we investigate this issue in more detail by examining how bidding dynamics change in different phases of the game.

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.

The first quantile includes unique 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).

⁴²Following Rubinstein (1982), most theoretical analyses of strategic bargaining games predict immediate agreement. One notable exception is the model by Compte and Jehiel (2004) who derive gradualism in a bargaining game under complete information with state-dependent outside options.

Figure 3: Bidding dynamics in treatment T_D



5.2 Discovery

We now split bid runs in two phases: before the transaction is formed for the first time, which we call *first activation*; and after that. We call *discovery run* the bidding phase leading up to the first activation. Table 4 breaks down unique discovery runs ($N = 270$) by their total duration in minutes and presents summary statistics.⁴³ Unsurprisingly, discovery runs involving more players last longer, and short runs are more likely to involve a single bid. More interestingly, we see that the majority of players placing multiple bids increase their bid between the beginning and the end of the discovery run – consistent to a bidding-up discovery phase. For longer discovery runs, a significant share of players either reduce their bid over time or go back to their original bid. But they still represent a minority of bidding players – i.e., from

⁴³We drop all the bid runs on transactions that were formed at the end of round 0 since, by definition, discovery has already taken place by the time subsequent rounds begin.

4.2% in the shortest runs, to 33% in runs over 2 minutes long. This is consistent with the view that, when most players follow a bidding-up discovery strategy, deviation from that strategy by some players tends to lengthen the time to the first activation.

Table 4: Discovery runs

Duration of the discovery run (in minutes)	<1	1-2	>2
# players bidding (mean)	1.59	2.13	2.51
# bids per active player (mean)	1.99	3.43	5.21
% bidding players who place one bid	55.6%	23.8%	15.9%
% bidding players with multiple bids: last higher than first	39.9%	54.3%	50.5%
% bidding players with multiple bids: last lower than first	2.4%	13.0%	20.7%
% bidding players with multiple bids: last = first	1.8%	9.1%	13.0%
<i>obs.</i>	104	83	83

To explore this idea further, we limit our attention to players who place multiple bids in discovery runs and we focus on *subsequent bids*, that is, the bids placed after the initial one. Table 5 provides summary statistics on 1397 unique subsequent bids placed by 383 players. As in Table 4, results are broken down by the total duration of the discovery run.

Table 5: Subsequent bids in discovery runs

Run duration (in minutes)	<1	1-2	>2
# of subsequent bids	166	407	824
% of increasing bids	88.6%	74.5%	66.4%
% of increasing bids: by 1 unit	73.5%	65.4%	50.9%
% of decreasing bids	11.4%	25.6%	33.6%
% of decreasing bids: by 1 unit	5.4%	18.2%	20.6%

Results indicate that the overwhelming majority of these subsequent bids are increasing bids – most of the time by a single unit. This is again consistent with a slow and cautious bidding-up strategy. Short bid runs are dominated by small

increasing bids. We also observe a non-negligible proportion of decreasing bids – still mostly by one unit. This proportion is higher in long discovery runs, consistent with the idea that discovery slows down when some players stop bidding-up before the first activation is reached.

The bidding-up discovery strategy followed by the majority of players is easily rationalized: it allows them to find a low bid for which the transaction occur. But it also means that, when this discovery process is curtailed in some way, as in sequential treatments that reach the limit of rounds or as when some players place decreasing bids, players may never reach this threshold and discovery fails. As a result, many profitable transactions do not occur, reducing aggregate efficiency.

5.3 Appropriation

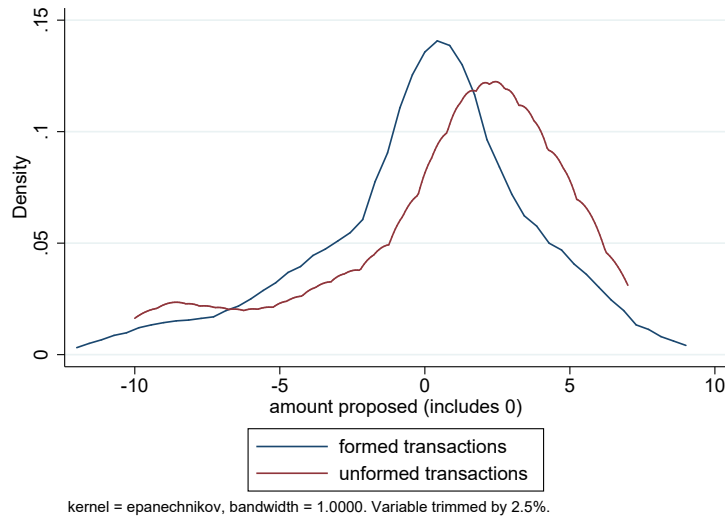
So far we have seen how subjects play a cautious bidding-up strategy to discover which transactions are profitable. Once first activation has taken place, however, some players may continue to haggle in an effort to increase their own gains. Since our experimental design implements deferred-acceptance bargaining, this puts us in a unique position to see whether bargaining continues after discovery.⁴⁴ In what follows we look into behavior after discovery ends and we provide evidence that, after first activation, some players reduce their bid. This appropriation strategy may be an attempt to get a larger share of the surplus, either by forcing others to increase their own bids, or by appropriating surplus from the silent auctioneer.

We start by noting that, in treatment T_D , not a single bid run stops immediately after discovery: once a link has been activated for the first time, we always observe more bids placed afterwards. Bids become less frequent, however: only 1030 unique bids (21% of total bids) are placed after first activation. The overwhelming majority of them (83%) are decreasing bids, i.e., they are lower than the bid placed by the same player at the time of first activation.⁴⁵ This suggests a desire to appropriate a

⁴⁴This sets our setting apart from that of Bakus et al. (2020), in which bargaining stops as soon as an agreement is obtained.

⁴⁵Out of these 1030 bids, only 106 bids (10%) represent an increase relative to the previous bid, while 856 (83%) represent a decrease. The remaining 68 (6%) are first-time bids by that player,

Figure 4: Bids on formed vs. unformed transactions



larger share of the surplus.⁴⁶ This interpretation is corroborated by Figure 4 below, which plots the Kernel density of bids placed on unformed transactions versus formed transactions (3965 and 1030 unique bids respectively). The Figure shows that bids on unformed transactions are higher and more frequently positive.⁴⁷

To summarize, the evidence we provide suggests that the behavior of subjects falls into two broad categories. A majority of subjects play a slow bidding-up strategy until a transaction is revealed to be profitable, and they freeze their bidding afterwards. A minority instead seeks to appropriate surplus by refraining from increasing their bidding on unformed transactions or by reducing their bids on formed

on that transaction. These first-time bids tend to be negative (-0.91 on average) even though the average value of these transactions is positive (1.34 on average). This suggests that these players bid to appropriate a share of a transaction that has just been revealed to be profitable.

⁴⁶This strategy benefits the player in most cases: only 180 (17%) of the 1030 bids trigger deletion. In the remaining 850 cases (83%) the transaction holds but the player now gets a higher gain.

⁴⁷Our data also show that decision time is longer for bids on formed transactions. For unformed transactions, the average time elapsed between the current and previous bid of a given player is 20 seconds and the median elapsed time is 6 seconds. In contrast, for formed transactions, the average time elapsed between bids is 30 seconds and the median elapsed time is 10 seconds. Decision time is often thought to proxy for cognitive load and is known to be higher for decisions that are more risky or include conflicting objectives.

ones, taking the activation as a signal that other players are willing to contribute to it.

In this Section, we have relied on results from treatment T_D because the unconstrained bargaining protocol provides richer data for analysis. We do, however, investigate the same issues in sequential treatments T_B and T_C as well, and we find consistent evidence of the same stylized facts – bidding-up during the discovery phase; and attempts at appropriation by a subset of players – see Online Appendix B for details.

6 Conclusion

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). Semi-structured bargaining games receive less attention than auctions, possibly because they are less amenable to clear-cut theoretical predictions. To help fill this gap in the literature, we design a laboratory experiment to quantify the efficiency gains of bargaining and the behavioral strategies of players in a decentralized market where valuations are heterogeneous and information is symmetrically private.

We implement three semi-structured bargaining protocols based on the principle of deferred acceptance, and we compare their ex-post efficiency to a benchmark scenario of a sealed-bid multilateral auction. We find that a protocol of unconstrained simultaneous bargaining doubles the efficiency margins with respect to sealed bids, mostly to the profit of players rather than the silent auctioneer. Thanks to the richness of our data, we are able to trace these results back to the behavioral strategies of participants. We find that players circumvent the lack of information by experimenting through a cautious and gradual strategy of incremental bids, which is facilitated and expedited in an unconstrained bargaining environment. Still, not all profitable transactions occur, which confirms the existence of the private information penalty documented in previous theoretical and empirical studies. We attribute this penalty to two main behavioral drivers. First, buyers haggle too hard as they try to extract

an unreasonable share of surplus. Their behavior stands in contrast to that of sellers, who tend to fix a more modest reservation price and stick to it. Second, a minority of participants deviate from the bidding-up strategy and/or continue bargaining after a provisional deal is reached, jeopardizing its chances of success.

Our findings have practical relevance for the design of online marketplaces for trading or philanthropic giving. Most of the existing digital platforms admit some limited bargaining action, but the current technology would allow them to integrate more complex bargaining features of the kind studied here. The bargaining protocols we propose are flexible and intuitive, and can scale up to operate in large markets.⁴⁸ Our design also offers the advantage of allowing for the remuneration of the trading platform, which incentivizes service provision.⁴⁹ Should a market for multilateral bargaining platforms develop, our results also provide useful insights for policy makers willing to regulate the industry – notably on the efficiency frontier.

Our current design lays the groundwork for future extensions in two directions. First, the experiment we present here does not model substitution between goods: all transactions have independent values; and participants do not compete over them. However, the advantage of allowing multiple trades at the same time is that players may consider certain goods to be substitutes.⁵⁰ But we know from earlier work (e.g., Comola and Fafchamps 2018) that laboratory subjects perform well in matching games of substitutes without transfers. It may therefore be possible for subjects to do well in a decentralized algorithm that combine both features: competition over goods and haggling over price. This would further integrate our results about bargaining behavior into the complex features of real-life markets.

⁴⁸For instance, sequential protocols could be implemented in markets with many participants but few items for sale. Similarly, the unconstrained protocol could be implemented in large markets with many participants and many items, as long as each participant is only interested in a limited number of items – which they could be asked to specify beforehand in order to be allowed to bid.

⁴⁹In our current design, the silent auctioneer fixes a threshold of total bids of 1 to implement a transaction. This minimum fee serves as remuneration for the platform, and it can be set to any number as desired. A large positive threshold increases the auctioneer’s surplus on deals. A negative threshold can be used to subsidize certain transactions, such as those generating positive externalities for society at large.

⁵⁰When haggling occurs sequentially and acceptance is not deferred, agents can find it hard to resist an early offer that is sub-optimal (e.g., Li and Rosen 1998; Niederle and Roth 2009).

Second, even though transactions are represented visually to subjects as a graph, our experiment does not speak about the topology of the resulting network nor of its externalities, since there are no resale or transfer across nodes. Still, our work feeds back into theory by opening 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 for utilities over indirect links would provide valuable information on 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.

Abdulkadiroglu, A., P. Pathak, A. Roth (2009), ‘Strategy-proofness versus efficiency in matching with indifference: Redesigning the NYC high school match’, *American Economic Review*, 99(5): 1954-78.

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.
- Ausubel, L., P. Cramton, R. Deneckere (2002), ‘Bargaining with incomplete information’, in *Handbook of Game Theory* (Amsterdam: Elsevier Science BV), Vol. 3, 1897-1945.
- Baccara, M., A. Imrohorglu, 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.
- 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.
- Bulow, J., P. Klemperer (2002), ‘Prices and the Winner’s Curse’, *RAND Journal of Economics*, 33 (1): 1-21.
- Burrows, P., G. Loomes (1994), ‘The Impact of Fairness on Bargaining Behaviour’, In: Hey J.D. (eds) *Experimental Economics*.
- Camerer, C., E. Johnson, T. Rymon, S. Sen (1993), ‘Cognition and framing in sequential bargaining for gains and losses’, *Frontiers of Game Theory* (MIT Press, Cambridge, MA), 27-47.
- 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.

Deck, C., C. Johnson (2004), 'Link bidding in laboratory networks', *Review of Economic Design*, 8(4): 359-372.

Demange, G., D. Gale, M. Sotomayor (1986), 'Multi-Item Auctions', *Journal of Political Economy*, 94(4): 863-72.

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.
- Fréchette, G. (2009), ‘Learning in a multilateral bargaining experiment’, *Journal of Econometrics*, 153(2): 183-195.
- Gale, D., L. Shapley (1962), ‘College Admissions and the Stability of Marriage’, *American Mathematical Monthly*, 69(1): 9-15.
- 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.
- Güth, W., S. Huck, P. Ockenfels (1996), ‘Two-level ultimatum bargaining with incomplete information: An experimental study’, *Economic Journal*, 106: 593-604.
- Güth, W., E. Van Damme (1998), ‘Information, strategic behavior, and fairness in ultimatum bargaining: An experimental study’, *Journal of Mathematical Psychology*, 42(2): 227-247.
- Haeringer, G., M. Wooders (2011), ‘Decentralized job matching,’ *International Journal of Game Theory*, 40(1): 1-28.
- 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.
- 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., 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.
- 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

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.

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.

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

Li, H., S. Rosen (1998), 'Unraveling in Matching Markets', *American Economic Review*, 88: 371-87.

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 un-

structured bargaining’, *Social Choice and Welfare*, 53: 655-675.

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.

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.

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.

Niederle, M., L. Yariv (2009), ‘Decentralized Matching with Aligned Preference’, NBER Working Papers 14840.

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. (1984), ‘The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory’, *Journal of Political Economy*, 92: 991-1016.

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.

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.

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.

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.⁵¹ 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 transactions 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, *should the link be formed at the end of the game*. This is simply computed as the *value* of the transaction to the player minus the last *bid* she made on it.⁵² These gains are color-coded: positive gains appear in green; negative gains appear

⁵¹In the instructions, we have chosen to represent only screenshots from the trial games, which only have 7 links instead of 10.

⁵²Note that this is not the same naming convention used in the analysis, where we define the gain as 0 whenever $b_{i,j,t} < 1$ (see Section 2.1). However, for the sake of the interface and instructions we found it more appealing to display the difference between value and bid, since the activation status is conveyed to players through line patterns as explained below.

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 much careful research and development, 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.

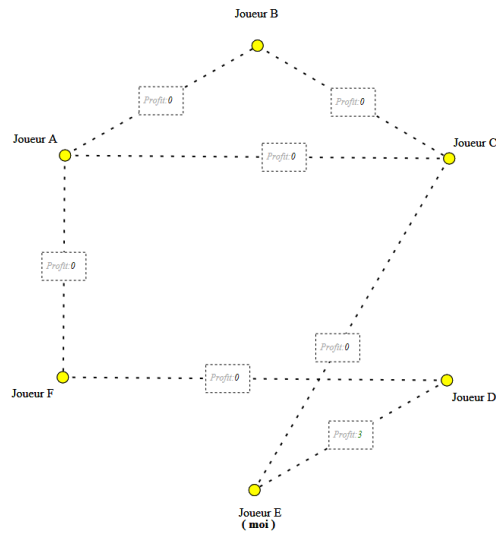
Bargaining process

The bidding process is as follows. While a link is being auctioned, the (solid or dotted) line representing it becomes thicker, the nodes at the link's end turn pink, and a dialog box opens in the middle of the screen. This dialog box contains two pieces of information: the value of the link, 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, and he is contemplating an offer of 2 units (Figure A.2) or a request of 2 units (Figure A.3).⁵³ 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 to them (Figure A.4). No constraint is imposed on the magnitude of bids made by subjects.

At the beginning of a game, the default value of all bids is set to zero. This is done so that each game starts with a blank slate, i.e., with no formed transaction. In games that include multiple rounds, the player's bid in the immediately preceding round becomes the default starting bid in the following round.

⁵³In Figures A.2 and A.3, player E has not validated his 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 implicit gain associated with the default bid of 0. The link E-D is currently not activated, so the link and tag lines appear dotted.

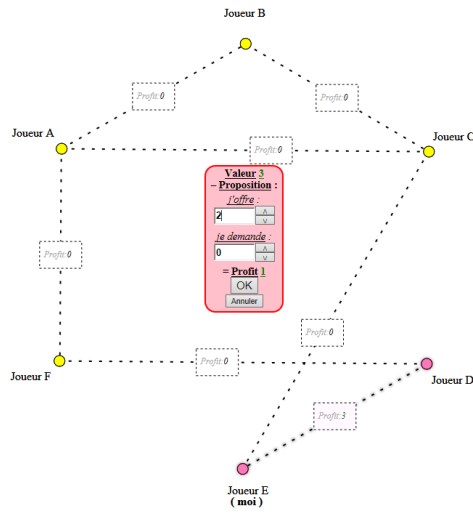
Figure A.1



LEGEND (FR - EN)

joueur = player
 profit = gain
 valeur = value
 proposition = bid
 j'offre = I offer
 je demande = I request

Figure A.2



LEGEND (FR - EN)

joueur = player
 profit = gain
 valeur = value
 proposition = bid
 j'offre = I offer
 je demande = I request

Figure A.3

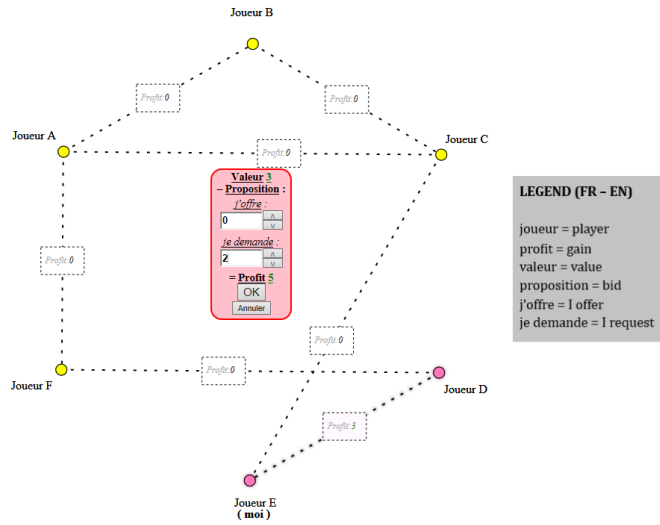
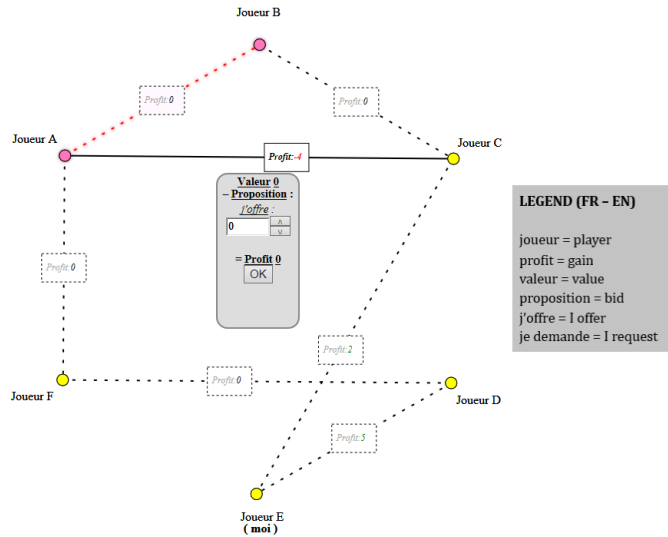


Figure A.4



A.2 Instructions to subjects⁵⁴

“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 are 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.

⁵⁴Translated from French. Figures 1 to 4 in the instructions correspond to the Figures A.1 to A.4 here.

The gain

On each link you will see a tag with a dotted or solid frame, and with 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).

Value

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

- Please note that the values are different for each player. Thus, 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 everyone!
- 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 make 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 a link between third players. *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 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 all players' bids for a given link is strictly positive, then the link is (provisionally) formed. That is, if the sum of all positive bids (offers) exceeds the negative bids (requests), the link is formed. However, 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 depends only 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 final 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 (even if some bids have changed within the round).

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 is up, 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: Ancillary results

B.1. Analysis of aggregate efficiency

Table 6 expands on the statistics reported in Figure 1. To do so, we report efficiency and surplus for the players and the auctioneer, respectively, for all the transaction types (all, profitable, non-profitable, mutually profitable, collectively profitable) and for both round 0 (sealed-bid phase) and the last round of the game.

As discussed in Section 3, we find that efficiency increases when bargaining is possible, and that the increase is mostly driven by the surplus of players. This is true across all transaction types and all bargaining treatments. Additionally, Table 6 shows that the player behavior is comparable in the sealed-bid phase across treatments, and that the percentage of non-profitable transactions formed remains negligible across all treatments and rounds.

To validate the results presented in Table 6, we run the linear regression

$$outcome_{ij,T} = \beta_0 + \beta_1 T_g + \beta_2 S + \lambda_s + \lambda_{ij*m} + \varepsilon_{ij,T} \quad (4)$$

where $outcome_{ij,T}$ represents the outcome of transaction ij at the end of the game, T_g is the vector of treatment dummies, and S is the order in which the game is played (from 1 to 8). We include session-level fixed effects λ_s , and transaction-per-matrix fixed effects λ_{ij*m} ($10 * 8 = 80$ effects). The latter control for variation arising from the random allocation of matrices to different treatments, and it absorbs transaction-specific attributes such as the total value of a link across players. Standard errors are wild-bootstrapped at the group level. In Tables 7, 8, 9 we consider three outcomes of interest, all computed for the end of game. The first outcome $g_{ij,T} = \sum_k g_{ij,T}^k$ is defined as the total gain realized by all players on transaction ij (which is zero if the transaction is not formed). The second outcome is a binary variable $deal_{ij,T}$ which equals one if the transaction occurs by the end of the game. The third variable is $players\ surplus_{ij,T}$ defined as the ratio $g_{ij,T}/v_{ij}$. The results indicate that all three bargaining protocols increase total surplus by increasing the gains, the number of profitable transactions, and the surplus for players. This is true across all trans-

Table 6: Efficiency of transactions

T	Round	transactions	all	non profitable	profitable	mutually profitable	collectively profitable
		# of transactions	440	118	322	178	144
A	0	efficiency	0.2	0	0.27	0.41	0.1
		surplus players	0.14	-	0.19	0.3	0.06
		surplus auctioneer	0.06	-	0.08	0.11	0.04
		# of transactions	440	133	307	235	72
B	0	efficiency	0.29	0.02	0.41	0.48	0.18
		surplus players	0.22	-	0.3	0.36	0.09
		surplus auctioneer	0.07	-	0.11	0.12	0.09
		# of transactions	440	164	276	244	32
C	0	efficiency	0.23	0.02	0.36	0.4	0.09
		surplus players	0.17	-	0.26	0.29	0.04
		surplus auctioneer	0.06	-	0.1	0.11	0.05
		# of transactions	440	135	305	201	104
D	0	efficiency	0.22	0.07	0.31	0.42	0.12
		surplus players	0.16	-	0.23	0.31	0.08
		surplus auctioneer	0.06	-	0.08	0.11	0.04
		# of transactions	440	135	305	201	104
D	last	efficiency	0.47	0.03	0.66	0.79	0.41
		surplus players	0.41	-	0.56	0.67	0.36
		surplus auctioneer	0.06	-	0.10	0.12	0.05

Table 7: Results for gains $g_{ij,T}$

	(1)	(2)	(3)	(4)	(5)
transaction	all	not profitable	profitable	mutually profitable	collectively profitable
T_B	2.232*** (0.000)	-0.095* (0.080)	3.085*** (0.000)	3.460*** (0.000)	2.803*** (0.004)
T_C	1.939*** (0.000)	-0.068 (0.127)	2.685*** (0.000)	2.915*** (0.000)	2.854** (0.017)
T_D	3.029*** (0.000)	-0.120 (0.221)	4.230*** (0.000)	4.912*** (0.000)	3.277*** (0.000)
S	0.119** (0.030)	-0.027** (0.044)	0.189*** (0.008)	0.180* (0.07)	0.270*** (0.006)
$\lambda_{ij* m}$	yes	yes	yes	yes	yes
λ_s	yes	yes	yes	yes	yes
Obs	1,760	550	1,210	858	352
R-sq	0.707	0.109	0.662	0.672	0.476

Notes: Wild-bootstrapped p-values in parentheses, clustered at the unique group level. *** p<0.01, ** p<0.05, * p<0.1

actions, as well as for mutually and collectively profitable transactions separately. The number of non-profitable transactions stays unaffected. Treatment T_D is best across the board, suggesting that easier discovery helps efficiency. There is also some evidence of learning across games as the session progresses.

Next we compare the two sequential bargaining treatments. As anticipated, we find that games in T_B last the longest: in our data the average number of rounds is 6.4 and only 2 of the 44 games end naturally. This is undoubtedly the result of the stopping rule based on changes in bids. In contrast, in treatment T_C where the stopping rule is based on transactions, the average number of rounds is 3.9 and 35 games out of 44 end naturally. The results shown in Tables 7, 8, 9 indicate that, although slightly more transactions are formed under T_C , efficiency is slightly higher in T_B and players make higher gains per transaction. Speeding up play by applying a stricter convergence rule comes at a cost to players. The magnitude of the difference, however, is small compared to the contrast with T_A or T_D . Given that treatment T_C

Table 8: Results for $deal_{ij,T}$

	(1)	(2)	(3)	(4)	(5)
transaction	all	not profitable	profitable	mutually profitable	collectively profitable
T_B	0.206*** (0.000)	0.026* (0.077)	0.270*** (0.000)	0.302*** (0.000)	0.251*** (0.002)
T_C	0.211*** (0.000)	0.013 (0.102)	0.292*** (0.000)	0.331*** (0.000)	0.214** (0.039)
T_D	0.266*** (0.000)	0.024 (0.174)	0.357*** (0.000)	0.414*** (0.000)	0.268*** (0.000)
S	0.014*** (0.008)	0.008** (0.048)	0.018** (0.031)	0.014 (0.208)	0.025** (0.013)
λ_{ij*m}	yes	yes	yes	yes	yes
λ_s	yes	yes	yes	yes	yes
Obs	1,760	550	1,210	858	352
R-sq	0.555	0.136	0.441	0.367	0.378

Notes: Wild-bootstrapped p-values in parentheses, clustered at the unique group level. *** p<0.01, ** p<0.05, * p<0.1

Table 9: Results for *players surplus*_{ij,T}

	(1)	(2)	(3)	(4)	(5)
transaction	all	not profitable	profitable	mutually profitable	collectively profitable
T_B	0.206*** (0.000)	0.038* (0.081)	0.264*** (0.000)	0.297*** (0.000)	0.236*** (0.002)
T_C	0.180*** (0.000)	0.020 (0.116)	0.243*** (0.000)	0.270*** (0.000)	0.226*** (0.010)
T_D	0.267*** (0.000)	0.037 (0.146)	0.354*** (0.000)	0.413*** (0.000)	0.268*** (0.000)
S	0.015*** (0.000)	0.012** (0.047)	0.017*** (0.010)	0.015* (0.061)	0.023** (0.015)
λ_{ij*m}	yes	yes	yes	yes	yes
λ_s	yes	yes	yes	yes	yes
Obs	1,760	550	1,210	858	352
R-sq	0.555	0.136	0.441	0.367	0.378

Notes: Wild-bootstrapped p-values in parentheses, clustered at the unique group level. *** p<0.01, ** p<0.05, * p<0.1

cuts game time in half and converges naturally, it may overall be preferable to T_B in spite of a small difference in efficiency, depending on the situation.

B.2. The Dynamics of Sequential Bargaining

Bidding up

We now focus on sequential treatments T_B and T_C and investigate how bids and gain evolve as rounds progress. We do so from two different perspectives: player level, and transaction level. We first estimate a player-level regression of the form

$$outcome_{k,r} = \beta_0 + \beta_1 T_C + \beta_2 \cdot R + \beta_3 \cdot S + \lambda_m + \lambda_s + \varepsilon_{k,r} \quad (5)$$

where $outcome_{k,r}$ is the outcome of player k at the end of round $r > 0$ for treatments T_B and T_C . Note that the number of rounds varies from 2 to 8 in the sequential treatments, and is endogenously determined to a large extent. We include a dummy for treatment T_C as well as round and game order indicators R and S . We also include matrix and session fixed effects λ_m and λ_s . Error terms are wild-bootstrapped at the group level, as before.

Results from regression (5) are shown in Table 10 for two outcomes: the total value of all bids placed by player k at the end of round r ($b_{k,r} = \sum_{ij} b_{ij,r}^k$); and the total gain of player k for all the transactions (provisionally) formed at the end of round r ($g_{k,r} = \sum_{ij} g_{ij,r}^k$). At the intensive margin (i.e., for a formed transaction), bidding more leads to lower gains. But by bidding more the player can form the transaction, which benefits her at the extensive margin. We see that, over rounds, players increase both the total value of all bids *and* the total gains. This suggests that the extensive-margin effect of increased bidding (i.e., more formed transactions) dominates the intensive-margin negative effect. This is consistent with the results in the unconstrained treatment T_D showing that most players increase their bid until a profitable transaction is revealed.

Next we take a transaction-level perspective and focus on final outcomes in treat-

Table 10: Player-level outcomes in T_B and T_C

	(1)	(4)
Dependent variable:	$b_{k,r}$	$g_{k,r}$
T_C	1.572 (0.964)	0.109 (0.770)
R	0.486** (0.196)	0.447*** (0.000)
S	0.184 (0.209)	0.174*** (0.007)
λ_m	yes	yes
λ_s	yes	yes
Constant	-13.975*** (2.113)	1.477** (0.022)
Observations	2,736	2,736
R-squared	0.049	0.202

Notes: Wild-bootstrapped p-values in parentheses, clustered at the unique group level. *** p<0.01, ** p<0.05, * p<0.1

ments T_B and T_C (N=880). We test whether bidding intensity increases the likelihood of a deal by estimating a linear regression of the form:

$$outcome_{ij,T} = \beta_0 + \beta_1 T_C + \beta_2 N_{ij,T} + \beta_3 S + \lambda_{round_{ij,T}} + \lambda_s + \varepsilon_{ij,T} \quad (6)$$

where $outcome_{ij,T}$ is the outcome for transaction ij by the end of game, T_C is a dummy for treatment T_C , and $N_{ij,T}$ represents the total number of unique bids placed on transaction ij during that game (this captures the intensity at which players bid on that transaction). S represents the order in which the game is played (from 1 to 8). Since we are dealing with sequential bargaining protocols, we include a set of dummies $\lambda_{round_{ij,T}}$ to control for the round at which the last bid on transaction ij was placed. We also include session fixed effects λ_s , and we wild-bootstrap the error term at the group level as before. Here $outcome_{ij,T}$ stands for two outcomes of interest: $b_{ij,T}$ which represents the sum of the final bids on transaction ij at the end of the game; and the dummy $deal_{ij,T}$ which takes value one if, by the end of the

Table 11: Transaction-level outcomes in T_B and T_C

Dependent variable:	$b_{ij,T}$	$deal_{ij,T}$
T_C	-1.422*	-0.092**
	(0.072)	(0.041)
$N_{ij,T}$	0.358***	0.014***
	(0.002)	(0.009)
S	0.522	0.023**
	(0.105)	(0.025)
$\lambda_{round_{ij,T}}$	yes	yes
λ_s	yes	yes
Constant	-12.78***	0.40***
	(0.001)	(0.000)
Observations	880	880
R-squared	0.051	0.040

Notes: Wild-bootstrapped p-values in parentheses, clustered at the unique group level. *** p<0.01, ** p<0.05, * p<0.1

game, $b_{ij,T} > 0$ and the transaction is formed. Results, presented in Table 6, show that, even after controlling for the round at which the last bid on the transaction is placed, more intensive bargaining is associated with an increase in bids and in the probability of a deal. Again this is consistent with experimental subjects playing a bidding-up strategy until profitable transactions are revealed.

Discovery and Appropriation

We conclude by briefly discussing how players react to the information of a transaction being formed in sequential treatments T_B and T_C . To recall, sequential treatments are organized into rounds and turns, and subjects only observe whether a transaction is formed at the end of a turn, once all players had the opportunity of revising their bids on that transaction. We notice that when players place a bid on transaction ij in round r and the transaction does not get formed in that round (4459 unique bids, 72% of total), a large fraction of players (46%) keep their bid unchanged

in the next round $r + 1$. But we also observe that 42% increase their bid and only 12% decrease it in round $r + 1$ relative to round r .⁵⁵ That is, approximately half of subjects respond to an unformed transactions by increasing their bid. A substantial fraction of the other players adopt a wait-and-see attitude and a small number of players reduce their bid in the hope the others would increase theirs.

Conversely, we notice that when players place a bid on transaction ij in round r , and the transaction is formed in that round (1714 bids, 28% of total), they all stall or decrease their bid: in the next round $r + 1$ we observe that 67% of these players place no new bid, and 33% decrease their previous bid.⁵⁶ These results are in line with evidence on discovery and appropriation in the simultaneous bargaining treatment discussed in Section 5.

⁵⁵A similar bidding pattern is observed if we take a 3-period window and restrict the analysis to the sub-sample of transactions that were not formed in round $r - 1$ but formed in round r .

⁵⁶If we analyze longer play sequences the aggressive strategy of decreasing the bid tends to be temporary in case it triggers deletion. If the transaction is dropped after players reduce their bid, they overwhelmingly revert to their previous behavior and increase their bid in response.