Dynamic Bargaining with Incomplete Information: Theory and Experimental Evidence

Dongkyu Chang* Peter Cramton† Jeongbin Kim‡ Axel Ockenfels§
Nicolas Roever¶

August 4, 2026

Abstract

Private information is often viewed as a source of bargaining power. We argue that its value depends on how bargaining translates hidden values into prices. We study this link in a model-guided, continuous-time laboratory experiment with 384 subjects and 5,393 bilateral negotiations. The experiment varies information structure and transaction costs while preserving an unstructured bargaining protocol. Bargaining reveals private information only partially. Rather than separating through delay, subjects rapidly exchange offers and narrow the gap between opening positions through successive revisions. This partial revelation generates price compression: deal prices respond to private values, but less than full revelation predicts. Consequently, private information has different distributional effects across types. High-value informed buyers earn information rents, whereas intermediate types are harmed. In our setting, private information becomes bargaining power only when bargaining translates information into prices with sufficient precision.

Keywords: Bargaining; Transaction costs; Delay; Information asymmetry; Laboratory experiment; First-mover advantage

JEL codes: C78; C91; C92; D82; D44

*City University of Hong Kong, email: donchang@cityu.edu.hk †University of Maryland, email: cramton@umd.edu ‡Florida State University, email: jkim33@fsu.edu §Adenauer School of Government and Department of Economics, University of Cologne, 50923 Cologne, Germany; Max Planck Institute for Behavioral Economics, 53113 Bonn, Germany, email: ockenfels@uni-koeln.de ¶University of Cologne, email: nicolas.roever@wiso.uni-koeln.de

Ockenfels gratefully acknowledges support by the German Research Foundation (DFG) under Germany's Excellence Strategy – EXC 2126/2 – 390838866. We gratefully thank Max R. P. Grossmann for his outstanding support in running the experiment and time management. We also thank Aenne Läufer for programming assistance and Dongjin Hwang for excellent research assistance. The data collection was pre-registered on AsPredicted (#232620). The research described in this article was approved by the ethics committee of the University of Cologne. All data and code are available at https://github.com/NicolasRoever/bargaining_experiment_analysis.

1 Introduction

Private information is commonly viewed as a source of bargaining power. Bargaining theory, however, gives a more complicated answer. Across canonical models, the informed party captures varying amounts of surplus in equilibrium. When and how does private information become bargaining power? We argue that private information does not automatically confer bargaining power. Its value depends on how the bargaining process translates hidden values into prices. When this translation is precise, informed agents can convert private information into rents. When it is coarse, prices are compressed toward a pooling level, so private information can benefit some informed types while harming others.

In the theoretical literature, different answers arise because bargaining processes differ in how they incorporate private information into prices. When the uninformed party can commit to a take-it-or-leave-it offer, private information does little to shift bargaining power toward its holder: the offer maximizes the uninformed party's expected payoff, and the informed party either accepts or walks away (Riley and Zeckhauser, 1983). In this case, prices contain private information only coarsely. At the opposite extreme, the Coase conjecture establishes that without credible commitment, frequent offers unravel the uninformed party's surplus and deliver nearly all gains to the informed side (Ausubel and Deneckere, 1989; Fudenberg and Tirole, 1983; Gul and Sonnenschein, 1988). Between these poles lie signaling models, in which the informed party reveals its value through costly delay and the parties subsequently split the surplus at type-specific prices (Admati and Perry, 1987; Ausubel et al., 2002; Cramton, 1991, 1992). In these models, delay becomes an information-transmission device: once the informed party's type is revealed, prices adjust to the revealed value.

These benchmarks clarify what is at stake empirically. If bargaining resembles commitment, prices should contain little information and the uninformed party should capture a large share of surplus. If bargaining resembles Coasian unraveling, the uninformed party should concede most of the surplus to the informed side. If bargaining resembles signaling, timing should reveal types and prices should respond to revealed information. Existing theory yields divergent predictions because each benchmark describes a different process by which private information is incorporated into prices. Existing evidence, however, has struggled to observe this process directly. Field bargaining environments rarely isolate the mechanisms that distinguish these models: parties have access to unknown outside options, face multidimensional uncertainty, and operate under institutional constraints absent from the theory (Backus et al., 2020; Larsen, 2021). As a result, it is difficult to tell whether private information affects bargaining power because it is precisely revealed, only partially inferred, or not priced at all.

We address this challenge with a model-guided laboratory experiment of continuous-time bilateral bargaining. The experiment implements the key primitives of incomplete-information bargaining theory—private values, discounting through random termination, and transaction costs—while leaving the bargaining protocol deliberately unstructured. Subjects can make offers, accept standing offers, or terminate negotiations at any time. This design lets us compare a sharp theoretical

benchmark with the process by which subjects actually negotiate: whether information is transmitted through the delay-based signaling mechanism predicted by the model, or instead through a different channel such as rapid offer exchange.

Our analytical benchmark builds on the framework of Cramton (1991, 1992), a model of bilateral bargaining under one-sided incomplete information with discounting and per-period transaction costs. In the benchmark environment, the buyer privately observes her valuation, while the seller's cost is common knowledge. Transaction costs make delay costly, thereby allowing buyers to signal through delay. The equilibrium partitions buyer valuations into three regions. Low-value buyers terminate immediately because the gains from trade are too small to justify continued negotiation. High-value buyers accept the seller's initial offer and earn information rents because the seller cannot distinguish among them. Intermediate buyers signal their values through costly delay; once the buyer's value is revealed, the continuation game becomes one of complete information and the parties split the surplus at the type-specific benchmark price.

We use this one-sided information environment as a sharp diagnostic benchmark and embed it in a broader 2×2 design that varies the information structure and the presence of transaction costs. Private information is held either only by the buyer or by both parties, and delay is either costly or not. We implement the experiment with parameters matched to the model's assumptions: uniform value distributions, known seller costs in the one-sided treatment, private buyer values, continuous-time unstructured bargaining, and random termination implementing discounting. We analyze a total of 5,393 negotiations from 384 subjects who are repeatedly matched in groups of eight participants and each play 30 bargaining games. The one-sided treatment with transaction costs provides the cleanest diagnostic test of the model's information-transmission mechanism, while the full design allows us to examine whether the same bargaining forces appear across different informational and cost environments. Because the protocol is deliberately unstructured, the benchmark does not impose an equilibrium path on subjects; rather, it disciplines the comparison between delay-based signaling and the bargaining process that subjects actually use to incorporate information into prices.

The data reveal that bargaining transmits private information only partially. The model's predictions are supported at the extremes of the buyer-value distribution: low-value buyers frequently exit, while high-value buyers trade quickly and capture on average 56% of the surplus, significantly more than the equal split. This is consistent with the model's information-rent logic: because sellers cannot distinguish the highest-value buyers from moderately high types, prices do not fully track buyer valuation. The central departure occurs in the middle of the value distribution. Instead of fully revealing their types through strategic delay and then splitting the surplus at the type-specific complete-information benchmark, intermediate buyers receive less than that benchmark: sellers—the party without private information—capture on average 58% of the surplus. This reversal does not occur because sellers learn buyer values precisely and bargain hard against them. Rather, it occurs because bargaining reveals buyer values only imperfectly.

These findings on the division of surplus are tied together by price compression. Deal prices move with buyer valuation, but less than the type-specific complete-information benchmark would

require. Sellers cannot condition their opening offers on buyer valuation, and the subsequent offer-counteroffer process only partially separates buyer types. For high-value buyers, the resulting compressed price lies below the equal-split benchmark and generates information rents. For intermediate buyers, the same compressed pricing works in the opposite direction: prices remain above the type-specific equal split, transferring surplus to the seller. Thus, private information has a type-dependent value. It benefits some informed types, harms others, and in our empirical setting reduces the informed party's surplus on average.

Price compression arises because the bargaining process maps private values into prices only imperfectly. In the equilibrium benchmark, signaling through delay fully reveals the buyer's valuation, so that the continuation game collapses to complete information and surplus is split at the type-specific Rubinstein price. In the data, bargaining departs sharply from this process. Rather than screening through a single initial offer followed by strategic silence, both parties exchange offers rapidly and narrow the gap between their opening positions through successive revisions. Buyers and sellers are roughly equally likely to make the first offer, offer timing shows no monotone relationship with buyer valuation, and deal prices respond to buyer valuation with a positive but attenuated gradient. Information enters prices, but not enough to make deal prices track the type-specific complete-information benchmark implied by full revelation.

The same qualitative pattern appears across the other treatment cells of our 2×2 design. In every cell, buyers and sellers engage in rapid back-and-forth bargaining that narrows the gap between opening offers, but each revision closes only a fraction of the remaining gap on average. Thus, the offer process provides a gradual and incomplete channel through which private information enters prices. Introducing transaction costs significantly reduces every process measure—negotiation duration, number of offers, pace of bargaining, and time to first offer—while leaving outcome measures such as trade rates, deal prices, and surplus division statistically unchanged. Transaction costs therefore accelerate bargaining, compressing time without substantially altering who trades, at what price, or how the surplus is divided. These comparative statics suggest that the limits of information transmission through offer exchange are central for understanding bargaining outcomes in our environment.

Finally, two behavioral regularities cut across the design. A visible spike of negotiations settles exactly at the equal split, and average surplus divisions deviate from 50–50 by at most about 5 percentage points in every cell, echoing the well-documented gravitational pull of fairness norms in bargaining. The party making the first offer obtains a meaningful surplus advantage, consistent with anchoring in negotiations. These forces coexist with, but do not subsume, the strategic price-compression mechanism: fairness and anchoring shape bargaining outcomes, while incomplete screening explains why the value of private information differs sharply across types.

Our paper contributes to three strands of the literature. First, we contribute to the empirical literature on bargaining under incomplete information. Field studies document rich bargaining behavior in settings such as eBay negotiations, wholesale auto markets, and consumer bargaining (Backus et al., 2020; Freyberger and Larsen, 2025; Larsen, 2021). These studies provide powerful

evidence on bargaining outcomes, but field settings make it difficult to isolate how private information is incorporated into prices. Experimental studies allow for tighter control, but early work often imposed rigid protocols (Forsythe et al., 1991; Rapoport et al., 1995), while recent work has moved toward richer bargaining environments (Bochet et al., 2024). Our experiment combines unstructured continuous-time bargaining with a fully specified structural benchmark, allowing us to study not only whether outcomes are consistent with theory, but whether the bargaining process implements the theory's information-transmission mechanism.

Second, we contribute to the long-standing question of whether private information translates into bargaining power. Canonical models give sharply different answers depending on commitment, screening, and signaling (Ausubel and Deneckere, 1992; Cramton, 1991; Fudenberg et al., 1985; Gul et al., 1986; Myerson and Satterthwaite, 1983). We show experimentally that the value of private information is type-dependent. High-value informed buyers earn information rents, but intermediate informed buyers are harmed because prices remain above their type-specific complete-information benchmark. The mechanism is price compression: bargaining incorporates private information into prices, but not with the precision required by full revelation. Thus, private information can be valuable locally but harmful on average.

Third, we contribute to the literature on fairness norms and behavioral regularities in bargaining. Equal-split norms are among the most robust findings in experimental economics (Camerer, 2003; Güth et al., 1982), and recent field evidence shows that negotiators often split the difference between recent proposals (Keniston et al., 2024). We show that equal-split norms and first-offer anchoring extend to a richer environment with continuous-time unstructured negotiation, private information, and transaction costs. At the same time, these behavioral forces coexist with systematic type-dependent deviations that fairness models alone do not predict, suggesting that fairness and strategic information transmission jointly shape bargaining outcomes.

The remainder of the paper is organized as follows. Section 2 develops the theoretical model and the diagnostic benchmark for information transmission through delay. Section 3 describes the experimental design, and Section 4 provides a descriptive overview of bargaining behavior across treatments. Section 5 studies the main one-sided treatment with transaction costs and documents the price-compression mechanism. Section 6 examines how offer exchange provides a gradual and incomplete channel for incorporating private information into prices. Section 7 analyzes the effects of information structure and transaction costs, Section 8 examines behavioral factors, and Section 9 and Section 10 discuss implications and conclude.

2 Bargaining Model

2.1 Setup

We develop a model of bilateral bargaining based on Cramton (1991, 1992). A buyer and a seller bargain over the price of a single unit of trade. At the beginning of the negotiation, the buyer's private valuation $b \in [\underline{B}, \overline{B}] \subset [0, \infty)$ and the seller's private cost (or reservation value) $s \in [\underline{S}, \overline{S}] \subset [0, \infty)$ are drawn from distributions F and G, respectively. These distributions are common knowledge and admit density functions f and g on their supports. We refer to the realized values of b and s as the parties' *types*. Each party privately observes their own type.

Bargaining Protocol Time is continuous, indexed by $t \in [0, \infty)$. After observing their types, both parties begin at $t = 0$ in an *inaction phase*. At any time t, either party may break this phase by making an initial price offer or by unilaterally terminating the negotiation. If either party terminates at any point, the negotiation ends immediately without trade. If either party makes an initial offer, the negotiation enters a *standing-offer phase* described below.[1]

Whenever a party makes an offer at time $t \geq 0$ (the *proposer*), the offer becomes the *standing offer*, and the negotiation enters a *standing-offer phase*. We impose a minimum delay $t_0 > 0$ following each offer, during which neither party may take any action.[2] After the delay, the other party (the *responder*) can take three actions: (i) accept the standing offer, in which case trade occurs at the standing price, (ii) make a counteroffer, which replaces the standing offer and switches roles so that the original proposer becomes the responder, or (iii) terminate the negotiation. The proposer may also unilaterally terminate the negotiation at any time $\tau \geq t + t_0$ while their offer is standing. The procedure repeats until agreement or termination occurs.

Payoffs and Equilibrium If the parties reach an agreement at time $t \geq 0$ at price $p \geq 0$ (paid by the buyer to the seller), then the seller's and buyer's payoffs are

$$u_S(s,t,p) := e^{-rt}(p-s) - \int_0^t e^{-rx} c \, \mathrm{d}x = e^{-rt}(p-s) - \frac{c}{r}\left(1 - e^{-rt}\right), \tag{1}$$

$$u_B(b,t,p) := e^{-rt}(b-p) - \int_0^t e^{-rx} c \, \mathrm{d}x = e^{-rt}(b-p) - \frac{c}{r}\left(1 - e^{-rt}\right). \tag{2}$$

Here $r > 0$ is the common discount rate, and $c \geq 0$ is a per-unit-time transaction cost incurred until the negotiation ends. If instead the negotiation continues up to time t and then terminates without trade, where $t = \infty$ if the negotiation continues indefinitely without trade or termination, the seller's

[1] If both parties make initial offers at the same time, a fair coin is flipped to determine which offer becomes the initial standing offer.

[2] This delay is imposed to avoid the well-known technical difficulties of continuous-time games. We are ultimately interested in the limit $t_0 \to 0$.

and buyer's payoffs are given by the discounted transaction costs:

$$-\int_0^t e^{-rx} c \, \mathrm{d}x = -\frac{c}{r}\left(1 - e^{-rt}\right), \tag{3}$$

with the expression understood as its limit when $t = \infty$.

We focus on pure-strategy Perfect Bayesian equilibria (PBEs), which we will often call simply equilibria hereafter. In a PBE, each party's strategy is optimal given that party's beliefs after every history of the negotiation, and beliefs are updated consistently with the observed history whenever Bayes' rule applies.[3]

2.2 Equilibrium Characterization and Comparative Statics

The general setup introduced above accommodates both one-sided and two-sided asymmetric information, matching the informational environments studied in our experiment. To develop the main theoretical predictions, we use the one-sided asymmetric-information case as our analytical benchmark. This case isolates the central information-revelation mechanism while allowing for a clear equilibrium characterization and comparative statics. The two-sided asymmetric-information case extends this mechanism to a richer environment in which both parties may use delay and offers to reveal private information.[4]

Complete-information benchmark As a benchmark, suppose the buyer's value b and the seller's cost s are common knowledge (equivalently, F and G are degenerate at b and s). If $s \geq b$, there are no gains from trade, and no trade occurs in equilibrium. Now suppose $s < b$. In this complete-information case, the model is essentially equivalent to the standard alternating-offer bargaining game à la Rubinstein (1982), augmented with symmetric transaction costs as in Cramton (1991). Since the transaction costs are symmetric, they affect the value of delay but not the limiting division of the gross gains from trade. In equilibrium, agreement occurs at the earliest feasible time, and the equilibrium trade price is

$$p^R(b,s) + \varepsilon(t_0), \tag{4}$$

where $p^R(b,s) := (b+s)/2$ and $\varepsilon(t_0) \to 0$ as $t_0 \to 0$.

One-sided asymmetric information We now turn to the one-sided asymmetric-information case, in which the seller's cost s is known (i.e., G degenerates at s), while the buyer privately observes her valuation $b \in [\underline{B}, \overline{B}]$. Relative to the complete-information benchmark, the key new feature is that response delay now serves as a signaling device: among buyers who do not accept or terminate immediately, the timing of the counteroffer helps the seller infer the buyer's type.

[3] See Fudenberg and Tirole (1991, Definition 8.2) for the standard definition of PBE. Fudenberg and Tirole define PBE for finite games; the extension to the present continuous-time setting is standard and omitted.

[4] Cramton (1992) provides a detailed analysis of the two-sided asymmetric-information case without transaction costs.

We begin by describing the equilibrium path in words. At the beginning of the negotiation ($t = 0$), the seller makes an initial offer that sorts buyer types into three groups based on their responses. Buyers with sufficiently low valuations terminate the negotiation immediately, because their gains from trade are too small to justify the transaction costs of continuing. Buyers with sufficiently high valuations accept the initial offer without delay, because delaying to obtain a better price is not worthwhile for them.

The remaining buyer types are in an intermediate range: they do not respond right away but instead wait before making a counteroffer. We refer to this intermediate interval as the *separating-delay region*. Within this region, delay is informative. In equilibrium, higher-value buyers wait for shorter periods, while lower-value buyers wait longer. In this sense, buyer types separate from one another through the timing of their counteroffers. Once the buyer's type is revealed, the continuation game becomes equivalent to the complete-information benchmark, and trade occurs immediately.

We now describe the equilibrium in further detail. For expositional simplicity, we state the strategies abstracting from the minimum-delay constraint $t_0 > 0$; incorporating this constraint requires only straightforward modifications and is therefore omitted. We restrict attention here to equilibrium-path play; the appendix describes off-path behavior. Finally, without loss of generality, we focus on the case $s \leq \underline{B}$.[5]

Let $b^\dagger$ and b^* denote the lower and upper cutoffs in the support of buyer types, respectively. These cutoffs will be pinned down shortly. In equilibrium, the seller makes an initial offer $p^{initial} = p^R(b^*, s)$ at time $t = 0$, where $p^R(b^*, s)$ is the equilibrium price in the complete-information benchmark in (4). In response to this initial offer, all buyer types below $b^\dagger$ terminate immediately, while buyer types in $[b^*, \overline{B}]$ accept the offer without delay. The remaining buyer types, $b \in [b^\dagger, b^*)$, are in the separating-delay region: each waits for an equilibrium delay $\Delta(b|b^*) > 0$ before making a counteroffer.

The equilibrium delay schedule is given by

$$\Delta(b|b^*) = \frac{1}{r} \ln \left[\frac{r(b^* - s) + 2c}{r(b - s) + 2c} \right], \qquad b \in [b^\dagger, b^*).$$

This expression has two useful properties. First, $\lim_{b \uparrow b^*} \Delta(b|b^*) = 0$; that is, the highest buyer types in the separating-delay region delay for an arbitrarily short time. Second, $\Delta(b|b^*)$ is strictly decreasing in b, so lower buyer types delay longer.[6] Hence, the timing of the counteroffer fully reveals the buyer's type. Once the buyer's type is revealed, the continuation game is equivalent to the complete-information benchmark, so the parties trade immediately at the complete-information price $p^R(b, s)$.

[5] If instead $\underline{B} < s$, then buyer types $b < s$ have negative gains from trade and terminate immediately in equilibrium. Since these types obtain zero and do not affect the continuation incentives of buyer types $b \geq s$, the same analysis applies after truncating the buyer-type distribution to $[s, \overline{B}]$.

[6] Appendix Sections Appendix D.2 and Appendix D.3 derive this delay schedule from the incentive-compatibility conditions of the buyer types in the separating-delay region.

On the equilibrium path, each buyer type b obtains

$$v_B(b|b^*) = \begin{cases} b - p^R(b^*, s) & \text{if } b \in [b^*, \overline{B}], \\ e^{-r\Delta(b|b^*)}\big(b - p^R(b,s)\big) - \displaystyle\int_0^{\Delta(b|b^*)} e^{-rx} c \, \mathrm{d}x & \text{if } b \in [b^\dagger, b^*), \\ 0 & \text{if } b \in [\underline{B}, b^\dagger), \end{cases} \tag{5}$$

where the three cases correspond, respectively, to immediate acceptance, delayed counteroffer, and immediate termination. The seller's equilibrium expected payoff is

$$v_S(b^*|b^\dagger) = \big(p^R(b^*,s) - s\big)\big(1 - F(b^*)\big) + \int_{b^\dagger}^{b^*} \left[e^{-r\Delta(b|b^*)}\big(p^R(b,s) - s\big) - \int_0^{\Delta(b|b^*)} e^{-rx} c \, \mathrm{d}x \right] \mathrm{d}F(b).$$

The first term is the expected-payoff contribution from buyer types in $[b^*, \overline{B}]$: with probability $1 - F(b^*)$, the buyer accepts the initial offer, and the seller earns $p^R(b^*,s) - s$. The integral is the expected-payoff contribution from buyer types in the separating-delay region $[b^\dagger, b^*)$. For each such buyer type b, trade occurs after delay $\Delta(b|b^*)$ at price $p^R(b,s)$, so the seller earns the discounted price-cost margin net of the transaction cost incurred during the delay. Specifically, the integrand equals $u_S(s,t,p)$ evaluated at time $t = \Delta(b|b^*)$ and price $p = p^R(b,s)$. Buyer types below $b^\dagger$ terminate immediately, so they contribute zero to the seller's expected payoff.

Finally, we complete the description of equilibrium play on the path by determining the two cutoffs b^* and $b^\dagger$. The lower cutoff $b^\dagger$ is the lowest buyer type willing to enter the separating-delay region rather than terminate. Equivalently, for any given b^*, $b^\dagger$ is determined by

$$b^\dagger = \min\{b \in [\underline{B}, b^*] : v_B(b|b^*) \geq 0\}. \tag{6}$$

When the cutoff is interior, this condition binds, so $v_B(b^\dagger|b^*) = 0$; otherwise, $b^\dagger = \underline{B}$. Hence, once b^* is fixed, this participation condition determines $b^\dagger$.

The remaining free variable is b^*, which is pinned down by the seller's choice of the initial offer $p^{initial} = p^R(b^*, s)$. If the seller deviates to a different initial price $\tilde{p}^{initial} \neq p^{initial}$, equilibrium play in the continuation game has the same qualitative form but with a different acceptance cutoff $\tilde{b}^*$ such that $\tilde{p}^{initial} = p^R(\tilde{b}^*, s)$. Therefore, b^* solves the seller's optimization problem

$$b^* \in \arg\max_{\tilde{b}^* \in [\underline{B}, \overline{B}]} v_S\big(\tilde{b}^* | b^\dagger(\tilde{b}^*)\big), \tag{7}$$

where the lower cutoff $b^\dagger(\tilde{b}^*)$ is determined by (6). If multiple maximizers exist, we select the smallest one.[7]

Several features of the equilibrium are worth noting. First, the equilibrium combines screening

[7] A unique maximizer exists if the buyer's type is uniformly distributed over $[\underline{B}, \overline{B}]$.

and signaling. The seller screens buyer types through the initial offer: high types accept immediately, low types terminate, and intermediate types enter the separating-delay region. The intermediate buyer types then signal their private valuations through response delay. Once a counteroffer is made, its timing fully reveals the buyer's valuation, so the continuation game has no remaining asymmetric information and the parties split the gains from trade equally without further delay.

Second, the presence of transaction costs affects the equilibrium outcome both qualitatively and quantitatively. (a) Some low buyer types terminate the negotiation immediately even though the expected gains from trade are positive, because their net expected payoff from continuing the negotiation is negative. (b) For any fixed cutoff b^*, the duration of delay $\Delta(b|b^*)$ decreases as the transaction cost increases, since delay becomes more costly.

However, observation (b) does not immediately imply that buyers delay less as c increases, because it does not account for the effect of a change in c on b^* and $b^\dagger$. The following proposition characterizes how the equilibrium cutoff b^* and the delay function jointly respond to changes in the transaction cost. It shows that higher transaction costs shorten negotiation on the equilibrium path in two ways: the seller's initial offer becomes less aggressive, inducing a larger set of buyer types to accept immediately, and buyer types who belong to the separating-delay region under both cost levels choose shorter delays. To emphasize the dependence of the cutoffs on c, we write them as $b^*(c)$ and $b^\dagger(c)$ in the proposition.

Proposition 1. *Suppose that $f(b)$ is nondecreasing on $[\underline{B},\overline{B}]$.*[8]

(i) The cutoff $b^*(c)$ decreases in c.

(ii) Consider c_H and c_L such that $c_H > c_L > 0$. Then

$$\Delta(b|b^*(c_H)) < \Delta(b|b^*(c_L))$$

for any

$$b \in [b^\dagger(c_H), b^*(c_H)) \cap [b^\dagger(c_L), b^*(c_L)).$$

Statement (i) implies that the cutoff $b^*(c)$ decreases in c. Thus, as c increases, the seller's initial offer $p^{initial} = p^R(b^*(c), s)$ weakly decreases, so a weakly larger set of buyer types accept the initial offer immediately. Statement (ii) can be interpreted as follows. Consider any buyer type b who chooses to delay rather than accept $p^{initial}$ or terminate the negotiation immediately under both $c = c_H$ and $c = c_L$. Then this buyer type delays for a shorter duration when $c = c_H$. Overall, the proposition implies that the negotiation concludes earlier with a higher transaction cost, conditional on the event that neither party terminates the negotiation.

[8]The nondecreasing-density assumption is a sufficient condition for the monotone comparative statics in the proposition and is satisfied by the uniform distribution used in our laboratory environment.

3 Experimental Design

We design a laboratory environment that implements the key primitives of the model—asymmetric information, discounting through random termination, and transaction costs—while allowing the bargaining protocol itself to remain unstructured. The theoretical model therefore serves as a diagnostic benchmark: it identifies the outcome patterns and information-transmission process implied by the equilibrium path characterized in Section 2, in which the seller's initial offer partitions buyer types into termination, immediate acceptance, and separating-delay regions. The experiment asks whether this benchmark organizes behavior when subjects can instead rely on the broader process of unstructured offer exchange.

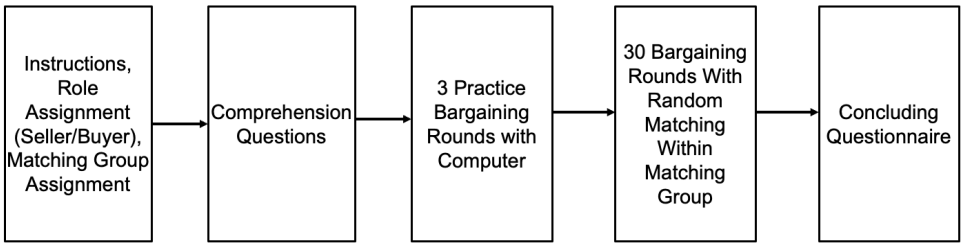
The primary treatment cell, *OneSidedCost*, implements one-sided incomplete information with transaction costs and discounting. To test the model's comparative-statics predictions, we embed this cell in a 2×2 between-subjects design that independently varies the information structure (one-sided vs. two-sided incomplete information) and the presence of transaction costs (zero vs. positive).

Table 1: 2×2 experimental design

	Transaction Cost	
Information Structure	0 cts/sec	5 cts/sec
Two-sided Incomplete Information	*TwoSidedNoCost*	*TwoSidedCost*
One-sided Incomplete Information (Buyer only)	*OneSidedNoCost*	*OneSidedCost*

Basic Setup At the beginning of the experiment, we communicate subjects their role (either buyer or seller), which remains fixed throughout. After reading the instructions and answering comprehension questions, subjects engage in three practice bargaining rounds with the computer. They are informed beforehand of the computer's strategy in these practice rounds. Then, subjects are matched within a matching group of 8 people (4 buyers and 4 sellers) for 30 rounds of bargaining.[9] At the end of the experiment, subjects answer a concluding questionnaire.

Figure 1: Flow of experiment

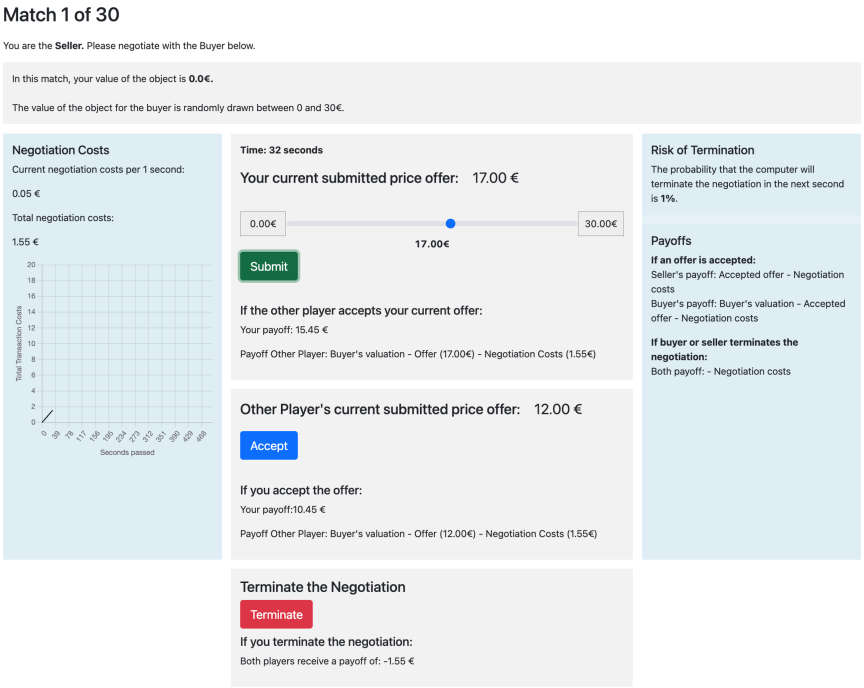


Note: The figure illustrates the timeline of a single experimental session, from the initial role assignment and instructions through the practice rounds, the 30 bargaining rounds, and the concluding questionnaire.

[9] Subjects are not aware of the matching group procedure.

Bargaining Interface We developed our own bargaining environment using the software package oTree (Chen et al., 2016).[10] Figure 2 shows the bargaining environment for a seller in treatment *OneSidedCost*. In the experimental conditions where there are no transaction costs, the left panel is not shown and participants see an empty blue surface instead. Players can at any time either submit a new offer, accept the other player's offer or terminate the negotiation. To ease cognitive burden, we automatically calculate payoffs for the participant for her current submitted offer, in case she accepts the other player's offer, or in case of termination. We also help participants to reduce mistakes when submitting an offer: In case buyers attempt to make an offer above their valuation (or sellers make an offer below their valuation), they are shown a pop-up window, asking them if they are sure to make this offer.

Figure 2: Bargaining environment for seller in *OneSidedCost*



Note: The left panel displays the accruing transaction costs in real time. Players can submit a new offer, accept the counterpart's current offer, or terminate the negotiation at any time. The interface automatically computes the payoff associated with each available action.

Implementing Time Discounting To incorporate time discounting (model parameter r), we implement a random termination procedure, following Roth and Murnighan (1978).[11] At the end of each second, the continuation of bargaining is determined by a geometric distribution with a 1% probability of termination. This implies that bargaining lasts 100 seconds on average before being terminated by the computer. The same sequence of termination draws is applied across all individual matching groups.

[10] The code is available at `https://github.com/NicolasRoever/bargaining_experiment`.

[11] For an application of random termination in repeated game experiments, see Dal Bó and Fréchette (2018).

Valuation and Costs For the one-sided incomplete information cases, the seller always has a valuation of 0, which is common knowledge. The buyer's valuation is drawn from a uniform distribution between 0 and 30 Euros (boundaries included). While the buyer observes their own valuation, the seller only knows the distribution from which it is drawn. For the two-sided incomplete information cases, valuations of buyer and seller are independently drawn from a uniform distribution between 0 and 60. This range is chosen to ensure that the expected gains from trade are comparable to those in the one-sided information treatments. In the transaction cost treatments, all players incur a cost of five cents per second during bargaining phase. This cost is common knowledge and is prominently displayed in the bargaining interface.

Logistics We conduct the experiment in summer and fall 2025 at the Cologne Laboratory for Experimental Research (CLER), which recruits participants from the student population in Cologne, Germany. In each session, subjects were randomly assigned to a matching group of 8 participants—4 buyers and 4 sellers—who were continuously re-matched over 30 rounds of bargaining. Our data-collection is pre-registered under AsPredicted (#232620) and we recruit $N = 96$ subjects per treatment arm. Table B.2 presents balance tests. During the experiment, two subjects became unresponsive midway through the bargaining rounds. In these instances, a researcher assumed their roles to allow the session to continue, and all rounds involving researcher participation were subsequently excluded from the analysis; we additionally exclude negotiations where one player accepted and the other terminated at the same time. Figure A.6 shows how our full sample is constructed.

Ensuring Data Quality With growing concerns about subject confusion and data quality in economics, we take several measures to ensure high data quality. First, we incentivise subjects: We randomly select one of the 30 bargaining rounds to determine the experiment payment, ensuring that subjects remain concentrated and mitigating concerns that actions are distorted by social desirability bias (Bursztyn et al., 2025) and experimenter demand effects (De Quidt et al., 2018). Second, subjects need to answer several comprehension questions after they have read the instructions. Third, our experiment environment encourages learning. After every bargaining round, subjects are informed of their payoff if the round is selected for payment; thus subjects are saliently informed if they have made mistakes in an individual negotiation. Fourth, before bargaining with other human subjects, there are three practice rounds. In these rounds, participants interact with a computer who follows a strategy they are informed about before. During these practice rounds subjects are informed about the payoff that would have resulted from the rounds—as in the real experiment. Fifth, we introduce a nudge when sellers propose an offer lower than their valuation and buyers propose an offer higher than their valuation: We show a screen asking them if they really intend to make this offer. This procedure is effective: we observe no such suboptimal offers in the experiment. Sixth, we designed our experiment to be rather short to not strain respondent's attention throughout. Our average full session length is only 37.52 minutes, much lower than in similar experiments

(Bochet et al., 2024, e.g.). Seventh, we pre-registered to exclude the first real round of bargaining, as well as negotiations in which one party makes an obvious mistake by accepting an offer that results in large losses; the threshold is a loss of 15 Euros or more from acceptance. These exclusions are designed to remove the lowest-quality observations—those most likely driven by initial unfamiliarity with the interface or clear decision errors—from the estimation sample. Eighth, we intentionally allow subjects to accept offers which are dominated by termination so that we can test our data quality. Across our final analysis sample of 5,393 negotiations, this occurs only 81 times, providing additional reassurance about data quality.[12]

4 Descriptive Overview of the Experiment

Table 2 presents summary statistics of bargaining behavior and outcomes across the four experimental conditions. Bargaining unfolds quickly: average negotiation duration is below 30 seconds in all conditions, and negotiations involve between 2.5 and 5.6 offers on average across treatments. The pace of play suggests that subjects rely on rapid judgment rather than long, deliberate thinking. The experimental environment therefore mirrors bargaining settings such as used-goods transactions, informal trade, and everyday commercial exchange, in which decision-making is rather fast.

Mechanically, introducing transaction costs decreases the average payoff players receive. Transaction costs also significantly accelerate the bargaining process, which systematically alters the nature of negotiation failures. Without transaction costs, subjects frequently exhaust the available time, and computer-enforced terminations account for 35–36% of all outcomes. With transaction costs, computer terminations fall sharply, but endogenous player terminations rise correspondingly—reaching 55% under two-sided incomplete information.

Table 2: Mean Bargaining Actions by Treatment

	OneSidedCost	*OneSidedNoCost*	*TwoSidedNoCost*	*TwoSidedCost*
Time per Negotiation (sec)	11.86 (11.75)	26.44 (25.99)	25.07 (27.32)	10.67 (12.99)
Number of Offers	3.51 (3.38)	5.61 (6.40)	4.71 (4.99)	2.54 (2.85)
Payoff	4.77 (6.25)	4.93 (6.58)	3.57 (8.06)	3.29 (8.43)
Avg. Computer Terminations	0.16 (0.36)	0.33 (0.47)	0.34 (0.47)	0.15 (0.35)
Avg. Player Terminations	0.31 (0.46)	0.15 (0.36)	0.40 (0.49)	0.57 (0.49)
Avg. Agreements	0.54 (0.50)	0.52 (0.50)	0.26 (0.44)	0.28 (0.45)

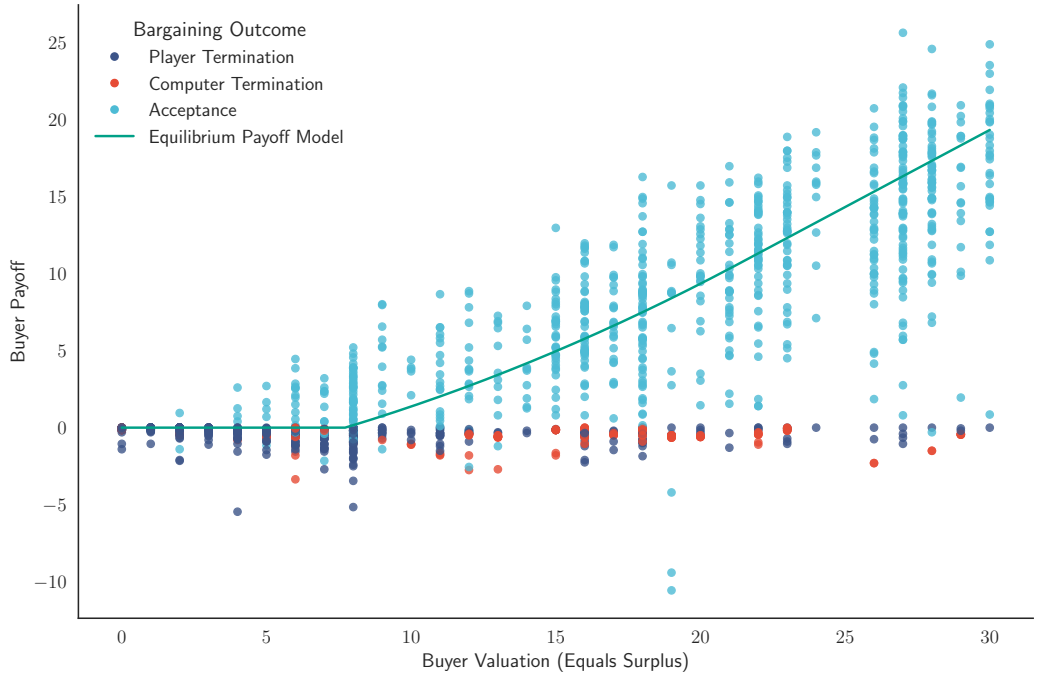
Notes: This table presents summary statistics of bargaining outcomes and actions across the four experimental treatments. Standard deviations are reported in parentheses. Terminations and agreements are expressed as sample proportions.

[12]Figure A.6 shows that also mistakes due to our pre-registered exclusion criterion are rare.

5 Main Results: Bargaining under One-Sided Incomplete Information

We now analyze our main treatment cell—one-sided incomplete information with transaction costs and discounting—in detail. Figure 3 shows the equilibrium buyer payoff from our model calibrated to the parameters of the experimental environment as well as the empirical payoffs of buyers from our experiment.[13]

Figure 3: Buyer payoff vs. valuation under buyer-only uncertainty with transaction costs (*OneSided-Cost*)



Note: The solid line shows the equilibrium prediction calibrated to our experimental environment with cutoffs at $b^{\dagger} = 7.72$ and $b^{*} = 21.40$.

The average residual between predicted and observed buyer payoffs is -1.71, which is 11.0% of the available surplus. In Table B.6, we compare the empirical fit of our model against several natural benchmarks: a simple linear regression of buyer payoff on buyer valuation, the equal-split benchmark, the Coase conjecture (informed party extracts entire surplus), and a model in which the seller makes a take-it-or-leave-it offer at the seller-optimal posted price of 15 (TILO). The first column of Table B.6 reports the R^2 from a regression of observed buyer payoffs on the payoffs predicted by each model. Our model achieves an R^2 of 58.0%, improving on the linear regression benchmark of 53%. This gain is meaningful: the linear benchmark exhausts all variation in buyer payoffs that can be attributed to buyer valuation through any affine transformation, so any improvement must come from non-linear structure. In our model, that structure comes from the equilibrium's three-region partition. Most R^2 improvement over the linear regression model comes

[13] See Appendix E for details on the calibration.

from the termination kink at $b^\dagger = 7.72$, below which buyers exit and receive nothing—which fits the data well. The Equal Split and Coase benchmarks, both linear in b, replicate the linear regression R^2 and offer no such improvement. The comparison against TILO ($R^2 = 57\%$) is closer. TILO also features a kink at the seller-optimal posted price of 15 and captures much of the same non-linearity as our model. Distinguishing our model from TILO on aggregate fit alone is therefore difficult; the two predictors agree on the broad shape of the payoff function and disagree primarily on where the kink sits and what happens above it. We thus turn to the three-region equilibrium of our model in detail, examining bargaining dynamics to test our model's predictions against the alternatives.

5.1 Buyers of Low Types ($b \leq b^\dagger = 7.72$)

In this region, our model predicts immediate termination by buyers. The data broadly confirm this prediction: 74% of negotiations end in player-initiated termination, compared to only 16% when $b \geq b^\dagger$. Among terminated negotiations, 89% are terminated by the buyer, consistent with the model's mechanism: upon learning that their valuation is low, buyers recognize that the gains from trade are too small to justify the costs of continued bargaining.

However, the data reveal a departure from our theory's stark prediction of immediate exit. Many buyers in this region attempt to bargain, particularly those with valuations between 5 and 7 (see Figure A.4). These buyers appear to understand that their position is weak but recognize that a quick agreement could still yield a positive payoff. Negotiations are shorter than in the upper regions, lasting 9.24 seconds on average versus 12.77 seconds for $b \geq b^\dagger$, but they are not trivially short; parties typically exchange a few quick offers before terminating or reaching an agreement.

5.2 Buyers of High Types $b \geq b^* = 21.40$

Our theory predicts that buyers of high types accept the seller's initial offer of $p^{initial} = 10.70$, so that trade occurs at this common price invariant to buyer valuation.

The bargaining *outcome* prediction matches the empirical data rather well. The trade rate is high (87%), and the mean trade price of 12 is close to the theoretical prediction ($p = 0.059$ for a t-test against $p^{initial} = 10.70$). Importantly, buyers capture on average 56% of the surplus, significantly more than the equal split ($p = 0.002$).

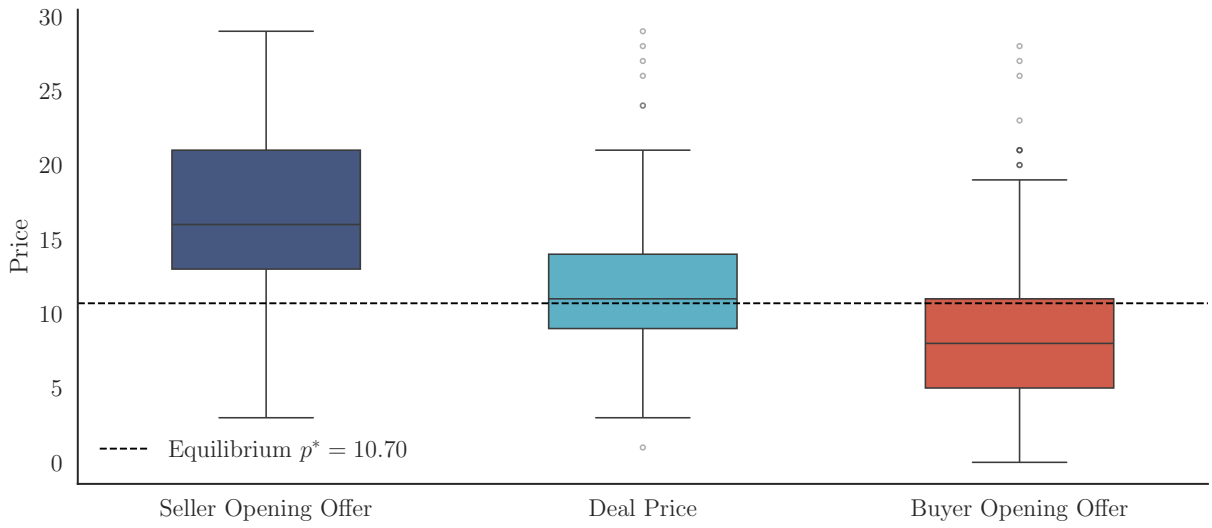
However, the *process* through which these bargaining outcomes are reached is very different from the model's prediction. The model predicts that the seller screens buyer types through a single initial offer that high-value buyers accept immediately. In our data, the first offer is roughly equally likely to come from either side: sellers move first in 50% of negotiations, buyers in 48%.[14] Moreover, the seller's opening offer averages 17 ($\sigma = 6$), well above the predicted price of $p^{initial} = 10.70$ and parties exchange on average 3.65 offers before reaching agreement. Figure 4 illustrates the

[14]These do not sum to one because some negotiations end—e.g., through computer termination—before either party makes an offer.

bargaining pattern in our data: sellers open high, buyers open low, and the agreed price falls in the middle of this gap, close to the model's predicted equilibrium price. Rather than screening and immediate acceptance, the data reveal a process of narrowing the gap through successive offers.

Despite the deviation from the process, the mechanism generating information rents for buyers operates much as the theory suggests. Since sellers are uncertain about the true buyer's valuation, their first offers are uncorrelated with the buyer's valuation. The empirical correlation coefficient is even slightly negative ($\rho = -0.11, p = 0.04$). The resulting deal price after bargaining does respond to valuation, but with a slope of $\hat{\beta} = 0.35$, statistically lower than 0.5 ($p < 0.001$); as a result, the absolute gap between a buyer's true valuation and the deal price widens as the valuation increases, allowing higher-valuation buyers to retain a larger overall share of the surplus. This is empirical evidence for information rents: because sellers cannot distinguish high-value buyers from moderately high types, prices do not fully track valuation, allowing high-value buyers to retain a larger share of surplus.

Figure 4: First offers and deal prices for buyer types $b \geq b^*$ in treatment *OneSidedCost*



Note: We restrict to negotiations where the buyer type is ≥ 21.40 and which end in agreement. We show the distribution for the first buyer offer, the first seller offer and the agreed price.

5.3 Buyers of Intermediate Types $b^\dagger < b < b^*$

In this region, our model predicts that buyers signal their type through strategic delay: lower types wait longer before making a counteroffer, and the resulting deal price tracks $b/2$.

As with high buyer types, this predicted *process* finds little support in our data. Sellers make the first offer in only 53% of negotiations, and buyers' offer timing shows no evidence of strategic delay. The correlation between buyer valuation and the time of the buyer's first offer is small and even slightly positive ($\rho = 0.09$), $p = 0.04$).

Bargaining *outcomes* reveal a striking pattern, which is also at odds with the model's prediction. While high-value buyers convert their informational advantage into a disproportionate share of the surplus, the pattern reverses for intermediate buyer types. There, it is the *sellers*—the party without private information—who capture the larger share, extracting on average 58% of the surplus, significantly above the equal-split ($p < 0.001$).

The mechanism behind this reversal is the same price compression that generates information rents in negotiations with high buyer types. Similarly, sellers' opening offers are independent of the buyer's valuation (correlation coefficient $\rho = 0.08, p = 0.11$), and the subsequent narrowing-the-gap process converges toward an equal split of the surplus: regressing the deal price on buyer valuation yields a slope of $\hat{\beta} = 0.48$ which is statistically indistinguishable from 0.5 ($p = 0.30$). But this convergence is incomplete. Deal prices systematically exceed $b/2$, a gap that is approximately constant across buyer valuations (see Table B.4). Because sellers cannot observe the buyer's type, their opening offers do not adjust to the available surplus. As a result, bargaining converges not to the exact type-specific equal split predicted by complete information, but to a compressed price schedule that retains a pooled premium reflecting the seller's initial expectation. For high-value buyers, this pooling price lies *below* the actual equal-split, generating the information rents documented above. For intermediate buyers, the same pooling price lies *above* the actual equal-split, transferring surplus to the seller.[15]

5.4 Within-Treatment Evidence on the Price Compression Mechanism

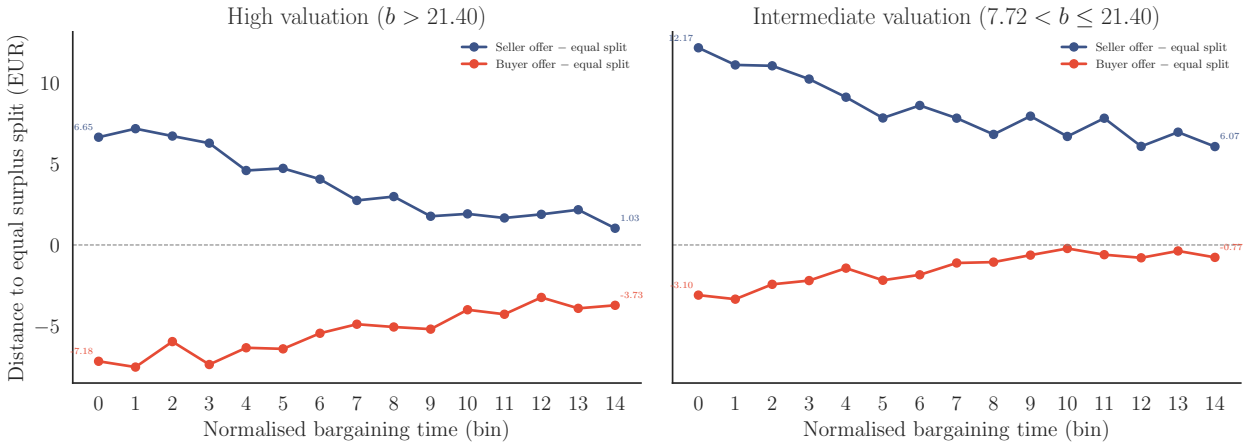
We now examine the bargaining process in more detail to show how price compression arises within negotiations. We focus on multi-offer rounds, that is, negotiations in which at least two offers are submitted.[16] To compare negotiations of different lengths, we normalize time within each match to the unit interval $[0, 1]$, where 0 denotes the first bilateral offer and 1 the last, divide this interval into 15 equal-width bins, and compute mean standing buyer and seller offers relative to the equal-split benchmark $b/2$. Figure 5 plots these distances for the high-valuation and intermediate-valuation regions.

Figure 5 makes the asymmetric-information structure visible, but it is important to interpret the normalization carefully. The figure plots offers relative to the type-specific equal-split benchmark, $b/2$. This normalization is useful for showing the distributional consequences of bargaining, but it also means that differences across valuation regions partly reflect changes in the benchmark itself. In absolute euro terms, sellers' offers are quite similar across the high- and intermediate-valuation regions. The mean opening seller offer is 17.32 in the high region and 16.95 in the intermediate region; the mean final seller offer is 13.97 in the high region and 13.42 in the intermediate region. Both at the opening and at the close, sellers post a near-identical nominal price schedule despite the

[15] Figure A.1 provides additional evidence, showing that for buyers with $b > b^{\dagger}$, the share of the surplus they extract increases in their type.

[16] One-offer rounds typically end immediately after the opening proposal and contain little information about within-round adjustment.

Figure 5: Standing offers relative to the equal-split benchmark over the bargaining clock in *OneSidedCost*



Note: Each panel plots the mean distance of the standing buyer offer (red) and standing seller offer (blue) from the equal-split benchmark $b/2$, where b denotes the buyer's valuation. Distances are computed at each offer event once both players have submitted at least one offer, using a normalized within-match clock divided into 15 equal-width bins (bin 0 = earliest bilateral offer, bin 14 = latest offer in the match). The left panel restricts to negotiations with a high buyer valuation ($b > 21.40$); the right panel covers the intermediate region ($7.72 < b \leq 21.40$). The dashed horizontal line marks zero. Sample: *OneSidedCost* negotiations with at least two offers submitted.

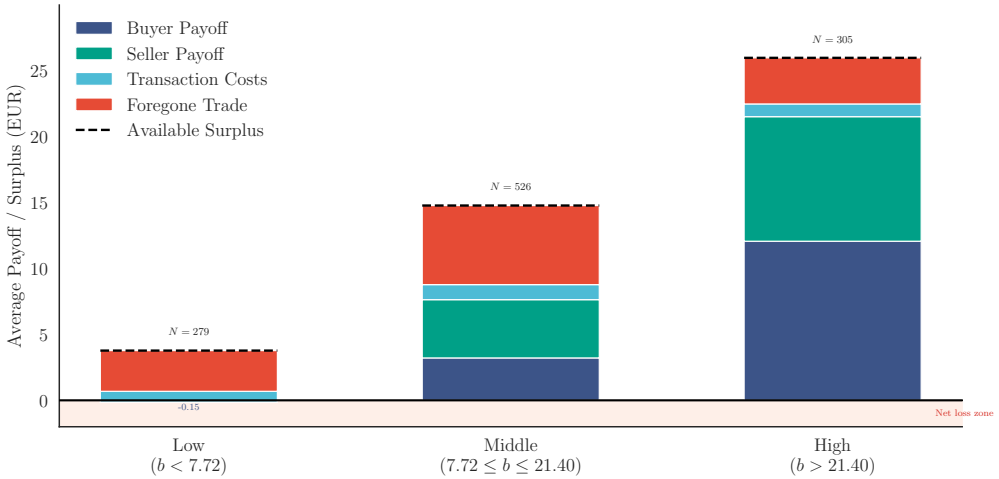
underlying surplus differing substantially across the two cells. Thus, the larger distance of seller offers from $b/2$ in the intermediate region does not arise because sellers condition strongly on buyer values. It arises because similar nominal seller offers are evaluated against a lower type-specific equal-split benchmark.

Buyers' offers, by contrast, respond to their privately observed valuations. The mean opening buyer offer is 7.81 in the high region and 5.38 in the intermediate region; the mean final buyer offer is 9.36 in the high region and 6.47 in the intermediate region. Private information therefore enters the bargaining process primarily through buyer offers rather than through seller screening. At the same time, convergence remains incomplete. In both regions, buyer and seller standing offers move toward each other over the bargaining clock, but they do not converge to the type-specific equal split. In the intermediate region, the resulting price remains above $b/2$, so sellers capture the larger share of the surplus. In the high region, the same compression works in the opposite direction, leaving buyers with information rents. This incomplete convergence is the within-treatment manifestation of price compression.

5.5 Efficiency

Figure 6 decomposes the available surplus into its empirical components: buyer payoff, seller payoff, transaction costs, and foregone trade due to terminations. We measure efficiency as the ratio of realized surplus to available surplus; the model's predictions are derived from the expectations given

Figure 6: Efficiency decomposition



Note: The figure decomposes the available surplus in the *OneSidedCost* treatment into its empirical components: buyer payoff, seller payoff, transaction costs, and foregone trade due to terminations. Efficiency is measured as the ratio of realized joint surplus to available surplus. Predictions are derived from the theoretical equilibrium calibrated to the experimental parameters.

our experimental parameters, i.e. Predicted Efficiency $= \mathbb{E}[\text{joint realized surplus}]/\mathbb{E}[b]$.

For low buyer types, nearly all surplus is foregone, consistent with the model's prediction that these buyer types should not trade. A small amount of transaction costs accumulates from failed negotiations, pushing average empirical efficiency slightly below zero (-3%) relative to the model's prediction of 0%.

In the high region, the model predicts full efficiency. Empirically, efficiency is 83%, meaningfully below this benchmark. This efficiency loss is driven by two sources of surplus destruction: terminations and transaction costs. Because our laboratory implementation of discounting relies on computer-imposed termination times, negotiations occasionally end after just a few seconds (Figure A.2), mechanically precluding some trades. In contrast, player-initiated terminations are rare, occurring in just 5% of interactions. The remaining shortfall is attributable to minor transaction costs that accumulate during brief delays prior to agreement.

The intermediate region is where the model and the data diverge most sharply. Empirical efficiency is 52%, below the model's prediction of 67%. The model attributes all inefficiency to costly delay: intermediate buyer types separate by waiting, with predicted delays averaging 25.34 seconds. Empirically, however, subjects reach agreement much faster (12.92 seconds on average), so delay costs are substantially lower than predicted. Instead, surplus is primarily destroyed by terminations that the model does not predict for this region. Figure A.3 shows that these terminations are not uniformly distributed across the intermediate region. In the lowest valuation bin $(7.72, 11]$, buyer-initiated terminations account for 38% of negotiations, and only 47% of negotiations end in agreement. As buyer valuations increase, agreement rates rise sharply and buyer terminations largely disappear. This gradient suggests that intermediate buyers near the lower cutoff face a calculus similar to that of low-region buyers: the available surplus is too thin to justify the cost of negotiation,

and they walk away. Subjects in the experiment appear to have a higher cutoff for engaging in bargaining and finding agreement than the model predicts.

Taken together, the model correctly predicts that low-valuation buyers terminate rather than engage in bargaining: resolving the information asymmetry through costly delay is not worthwhile given the thin available surplus. However, the model places this boundary too low—empirically, buyer types well into the intermediate region exhibit the same behavior, so that more surplus is destroyed through foregone trade than the theory predicts. At the same time, the model overestimates delay among those who do bargain, meaning that conditional on reaching agreement, the process is more efficient than predicted.

6 Bargaining Dynamics Across Treatments

Section 5.4 showed that, within the main treatment, standing offers converge gradually but incompletely toward the equal-split benchmark. We now step back from the type-specific benchmark $b/2$ and ask whether the same qualitative pattern appears more broadly across the full 2×2 design. Our aim is not to recover type-specific pricing in each treatment, but to characterize how information is incorporated into offers and prices. We do so using two reduced-form summaries of within-round behavior: the evolution of the buyer–seller proposal gap and the way current offers respond to the most recent proposals on both sides.

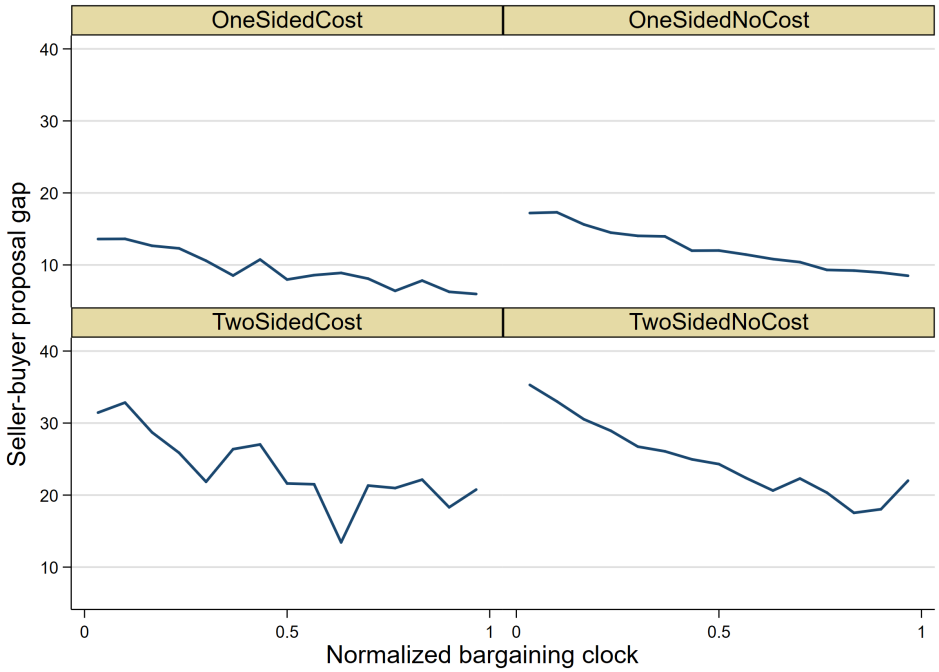
We again restrict attention to multi-offer rounds. For Figure 7, we summarize bargaining at the level of submitted proposals. We normalize time at the match level, divide the normalized clock into 15 equal-width bins, and compute mean submitted offers separately by treatment, role, and bin. The figure plots the resulting convergence gap, defined as the bin-level mean seller offer minus the bin-level mean buyer offer.

Because valuation support differs across treatments, we do not interpret the level of this nominal gap as directly comparable across information structures. Buyer values are drawn from $[0,30]$ in the one-sided treatments, whereas buyer and seller values are drawn from $[0,60]$ in the two-sided treatments. The figure is therefore meant to document a within-treatment process regularity: in each treatment, the buyer–seller proposal gap narrows over the bargaining clock.

Figure 7 shows a common pattern across treatments: conditional on bargaining continuing beyond the opening move, the buyer–seller proposal gap narrows over the bargaining clock. This within-round convergence differs from the delay-based signaling mechanism in the model. The model assigns the informational role to response timing: intermediate buyers separate by waiting different lengths of time before making a counteroffer. In the data, by contrast, subjects actively exchange proposals, and information appears to enter bargaining through successive offer revisions. The convergence pattern is present both with and without transaction costs, although the costly treatments display noisier paths because they contain fewer within-round revisions.

This figure documents proposal convergence, but it does not show how far each revision moves

Figure 7: Convergence gap over the normalized bargaining clock



Note: The figure plots the bin-level mean seller offer minus the bin-level mean buyer offer over the normalized bargaining clock. Offer times are normalized to the unit interval within each match and divided into 15 equal-width bins. The figure is intended to document within-treatment convergence in nominal proposal gaps. Because valuation support differs across information structures, nominal gap levels are not directly comparable across one-sided and two-sided treatments.

toward the opponent's position. To characterize the updating process more directly, we estimate how much of the remaining disagreement gap is closed by each new offer. The dependent variable is the movement of the current offer toward the opponent's most recent proposal, and the independent variable is the gap between the player's own most recent offer and the opponent's most recent offer. Both variables are direction-adjusted so that positive values correspond to movement toward the opponent. We estimate this relationship at the offer-event level, excluding initial offers for which the opponent has not yet made a proposal. A coefficient of zero corresponds to no adjustment, while a coefficient of 0.5 corresponds to moving exactly to the midpoint between the two most recent proposals.

Table 3 reports the results. In the pooled specification, the estimated coefficient is 0.075, implying that each revision closes about 7.5% of the remaining disagreement gap. Thus, subjects do move toward their counterpart's latest proposal, but the adjustment is far below the split-the-difference benchmark of one-half. The same pattern holds within each information structure and for both roles. In the one-sided treatments, buyers close only 3.3% of the remaining gap, while sellers close 17.9%. In the two-sided treatments, buyers close 5.8% and sellers close 8.1%. In every case, we strongly reject the midpoint benchmark. Within-round convergence therefore occurs through incremental concessions rather than immediate midpoint matching.

This regression does not by itself identify beliefs or posterior updating. Rather, it documents the

Table 3: Partial adjustment toward the opponent's most recent offer

	Pooled	One-sided Buyer	One-sided Seller	Two-sided Buyer	Two-sided Seller
Remaining gap	0.075***	0.033***	0.179***	0.058***	0.081***
	(0.003)	(0.006)	(0.006)	(0.004)	(0.005)
Observations	11,944	5,015	3,012	2,134	1,783
Test: $\lambda = 0.5$	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$

Notes: The table reports offer-level regressions of movement toward the opponent's most recent offer on the remaining disagreement gap. All regressions are estimated without a constant. Both variables are direction-adjusted so that positive values correspond to movement toward the opponent. A coefficient of 0.5 corresponds to moving exactly to the midpoint between the player's own most recent offer and the opponent's most recent offer. Standard errors clustered at the match level are reported in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

behavioral channel through which information can enter prices: offers respond to previous proposals, but each response incorporates only a small fraction of the remaining disagreement.

Taken together, Figure 7 and Table 3 provide a process-level explanation for why screening is incomplete in the data. Conditional on negotiations continuing beyond the opening move, bargaining proceeds through active offer revision rather than delay-based signaling alone. The buyer–seller proposal gap narrows over the bargaining clock, but each revision closes only a small fraction of the remaining disagreement gap. The offer process therefore provides a gradual and incomplete channel through which information enters prices. This helps explain why deal prices respond to buyer valuation yet remain compressed relative to the type-specific complete-information benchmark.

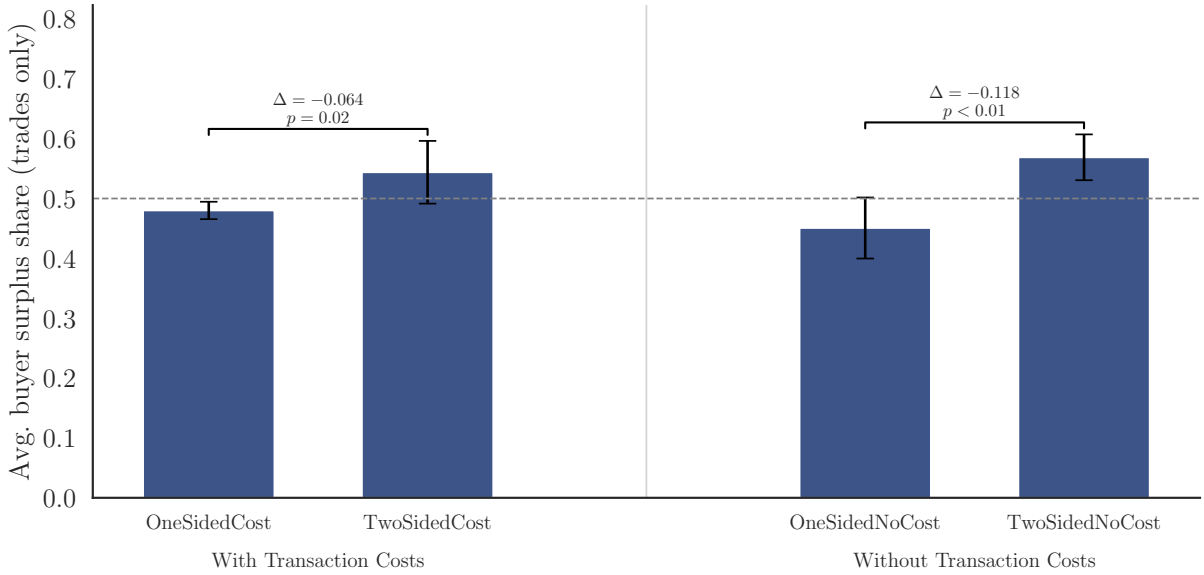
7 Effects of Information Structure and Transaction Costs

7.1 The Empirical Effect of Information Structure

Our main bargaining environment features one-sided asymmetric information, where the buyer holds private information about her valuation while the seller's cost of $s = 0$ is common knowledge. The preceding analysis showed that this private information can actually *hurt* the informed party through the price compression mechanism. Removing the buyer's informational advantage by making both parties' types private indeed improves the buyer's relative position. In both treatments with and without transaction costs, buyers extract a larger share of the surplus under two-sided than under one-sided asymmetric information (see Figure 8, and Table B.7).[17]

[17] In the two-sided asymmetric treatments, buyers extract on average more than half of the surplus. This advantage traces to a behavioral asymmetry in opening offers. Buyers' first offers average 17.5, which falls below the seller's average valuation of 19.1; sellers' first offers average 38.1, just below the buyer's average valuation of 41.7. Buyers thus anchor the negotiation more aggressively relative to their counterpart's valuation than sellers do. Because the remainder of the bargaining process is strikingly symmetric—buyers and sellers accept, terminate, and initiate offers at similar rates (Table B.3)—this opening asymmetry carries through to the final division of surplus.

Figure 8: Average surplus extracted by buyers by treatment



Note: To ensure comparability across treatments, we restrict to negotiations where the available gains from trade are positive, Table B.7 shows regressions for further analysis.

Contrasting the bargaining processes under the two information regimes shows that one-sided asymmetric information leads both parties to make complex, iterated strategic adaptations. Consider the seller's position. Under one-sided asymmetric information, the seller's reservation value of $s = 0$ is common knowledge. This transparency undermines the seller's ability to hold firm, and the seller concedes substantially from her initial offer. We measure the seller concession rate as the fraction of the seller's initial asking price conceded to reach agreement, $(p^{\text{ask}} - p^{\text{deal}})/p^{\text{ask}}$. The seller concession rate is 0.41 under one-sided information, compared to only 0.15 under two-sided uncertainty ($p < 0.001$). Yet the seller's greater concessions under one-sided information do not translate into higher buyer surplus, because buyers simultaneously adjust their own behavior. Under one-sided uncertainty, 58.81% of trades conclude with the buyer accepting the seller's standing offer, compared to only 47.42% under two-sided uncertainty. Termination patterns shift in parallel: 78.7% of breakdowns are buyer-initiated under one-sided information versus 51.0% under two-sided uncertainty. The buyer's greater willingness to accept is the rational complement to the seller's larger concessions: because the seller's weak position is common knowledge, her offers already incorporate larger concessions, leaving less surplus for the buyer to extract through continued bargaining. In the majority of one-sided negotiations, the buyer either accepts the seller's offer or walks away, rather than engaging in extended counteroffering. The net effect favors the sellers in our setting: the concessions they make in price are more than offset by the buyer's willingness to settle early, producing the surplus advantage documented above.

7.2 The Empirical Effect of Transaction Costs

We now turn to analyze the effect of transaction costs. We estimate the cost effect $\hat{\beta}$ from a pooled OLS regression of each outcome variable on an indicator for transaction costs and an indicator for the information structure; $\hat{\beta}$ is the coefficient on the cost indicator and thus $\hat{\beta}$ identifies the average main effect of introducing transaction costs across our experimental settings.

Table 4: Effect of transaction costs: process vs. outcomes

Outcome	**Mean (no cost)**	**Mean (cost)**	**Cost effect $\hat{\beta}$**
A: Process Variables			
Negotiation duration (sec)	27.925	12.817	-15.101 (1.408)***
Number of offers	5.491	3.342	-2.151 (0.303)***
Time per offer (sec)	5.284	3.640	-1.647 (0.289)***
Time to first offer (sec)	3.378	2.231	-1.153 (0.155)***
B: Outcome Variables			
Trade rate	0.536	0.559	+0.023 (0.019)
Buyer surplus share	0.489	0.499	+0.010 (0.037)
Deal price	15.056	15.447	+0.320 (0.631)

Notes: Cost effect $\hat{\beta}$ from OLS regression of the outcome variable on an indicator for whether the negotiation is from a treatment with or without transaction costs, and with one-sided or two-sided private information. We report the coefficient on the cost dummy; SEs clustered by participant in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4 shows that introducing transaction costs significantly reduces every process measure—negotiation duration, number of offers, pace of bargaining, and time to first offer. This acceleration is directly consistent with the comparative statics of our model, which predicts shorter delays when bargaining is costly. Interestingly, we find empirically that these costs act almost entirely as an acceleration mechanism without substantially altering who trades, at what price, or how the surplus is divided; outcome measures such as trade rates, deal prices, and surplus division remain statistically unchanged.

8 Behavioral Factors

We now turn to whether individual behavioral attributes commonly studied in behavioral economics systematically influence bargaining outcomes. Table 5 reports regressions in which the dependent variable is either the share of available gains from trade captured by a player or the player's final payoff. Fairness preferences and time preferences have no statistically significant effect on bargaining outcomes; however, we document large bunching at exactly 50–50 splits across all treatment cells (see Figure A.5). Moreover, the average surplus division deviates from 50–50 by at most roughly 5 percentage points. The equal-split norm appears to impose a strong gravitational force on bargaining outcomes. This robustness echoes a large body of evidence from ultimatum and unstructured bargaining experiments documenting the pervasiveness of equal-split norms (Camerer,

2003; Güth et al., 1982), and it is consistent with models of fairness-driven behavior such as Fehr and Schmidt (1999) and Bolton and Ockenfels (2000).

Risk aversion, by contrast, is associated with lower payoffs and a smaller share of the surplus. The magnitude of the effect is economically modest, but the direction is consistent with the interpretation that more risk-averse individuals settle earlier at less favorable terms rather than risk prolonged negotiation and the possibility of breakdown—for instance, through exogenous termination by the computer. Table B.5 provides further evidence in line with this interpretation.[18] We also find a significant and sizable effect of posting the first offer, which is robust to the inclusion of individual fixed effects. This is consistent with the well-documented anchoring effect in negotiations (Galinsky and Mussweiler, 2001; Thorsteinson et al., 2011): the party who moves first exerts disproportionate influence over the final price.

Table 5: Impact of behavioral factors

	Split Gains from Trade					Payoff				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Ultimatum Offer	0.00					-0.02				
	(0.00)					(0.02)				
Risk Aversion		-0.03***					-0.10			
		(0.01)					(0.09)			
Time Preference			0.01					0.02		
			(0.01)					(0.07)		
First Offer				0.07***	0.05***				0.79***	0.45**
				(0.02)	(0.02)				(0.21)	(0.22)
Individual FE					✓					✓
Baseline Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Obs	3493	3493	3493	3501	3501	6348	6348	6348	5999	5999
Adj. R^2	-0.00	0.00	0.00	0.02	0.15	0.39	0.39	0.39	0.41	0.44

Notes: The table reports OLS estimates of the effects of behavioral factors on bargaining outcomes. The sample is restricted to negotiations with strictly positive gains from trade. The dependent variable in columns (1)–(5) is the *split of gains from trade*, defined as the share of the total surplus accruing to the participant: $(p - v_s)/g$ for sellers and $(v_b - p)/g$ for buyers, where p is the deal price, v_s (v_b) is the seller's (buyer's) valuation, and $g = v_b - v_s$ is the total gains from trade. The dependent variable in columns (6)–(10) is the participant's *payoff* in euros. Each column includes a single behavioral variable of interest: *Ultimatum Offer* is an indicator equal to one if the participant's offer claimed more than 50% in the endline task; *Risk Aversion* is an ordinal variable depending on the lottery choice, higher values indicate higher risk-aversion; *Time Preference* is the switching point in a multiple price list eliciting discount rates; *First Offer* is a dummy equal to one if the person made the first offer. Baseline controls include treatment fixed effects and total gains from trade. Columns (5) and (10) additionally include individual fixed effects. Standard errors are clustered at the individual level and reported in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

9 Discussion

Our experimental environment implements key ingredients of canonical bargaining models—continuous time, bilateral incomplete information, discounting, and transaction costs—while al-

[18] This is consistent with the classical intuition that risk aversion weakens a bargainer's position by making the outside option of disagreement relatively more costly (see, e.g., Binmore et al., 1986; Roth, 1979) and also empirically documented in (Bochet et al., 2024).

lowing the bargaining process itself to unfold in an unstructured way. This combination lets us distinguish two questions that are often bundled together in theory: whether the model identifies the relevant economic forces, and whether subjects implement the specific equilibrium path through which those forces operate. We now discuss what the experiment reveals about the bargaining process, the distributional consequences of private information, and the implications for modeling.

Bargaining Process: Narrowing the Gap versus Screening and Signaling. The experiment shows that the model overstates how precisely bargaining reveals private information. In the model, delay is the channel of information transmission: intermediate buyers wait different lengths of time, and the timing of the counteroffer reveals their valuation. In the data, information appears to enter bargaining through a different channel. Subjects do not typically screen through a single offer followed by strategic silence. Instead, bargaining unfolds as a rapid back-and-forth exchange in which the two parties gradually narrow the gap between their opening positions. The process therefore looks less like full separation through delay and more like partial revelation through offer revision.

This pattern is consistent with a large body of experimental evidence showing that subjects in bargaining games often fail to implement the backward-induction strategies predicted by sequential models, even when outcomes approximate equilibrium predictions (Binmore et al., 1985; Roth et al., 1991). In our environment, the costs of delay are especially salient: subjects observe a prominent, continuously updating display of accruing costs, and agreements are reached in seconds. Under these conditions, rapid offer exchange appears to replace delay-based signaling as the main channel through which information enters prices. The data therefore suggest that the channel of information transmission depends on the time scale and protocol of bargaining. In longer-horizon negotiations—such as labor disputes, mergers and acquisitions, or legislative bargaining—waiting may be a more natural way to signal private information. In our setting, however, information enters prices primarily through offers and counteroffers.

This distinction refines, rather than overturns, the theoretical mechanism. The model identifies the key economic force—price compression arising from incomplete screening—but overstates the precision with which bargaining transmits information into prices.

Bargaining Outcomes: Information Rents and Price Compression. A central message of incomplete-information bargaining theory is that private information can create bargaining power. In Coasian models, the uninformed party's commitment power erodes as offers become frequent (Ausubel and Deneckere, 1992; Fudenberg et al., 1985; Gul et al., 1986); in screening models, the informed party earns information rents (Cramton, 1991). The common intuition is that knowing more can mean getting more.

Our results qualify this intuition in two ways. First, high-value buyers earn economically meaningful information rents, but these rents are far from complete surplus extraction. Second, and more importantly, information rents are not uniform across types. The same price compression that

benefits high-value buyers harms intermediate buyers.

The reason is incomplete screening. The model predicts that once an intermediate buyer's type is revealed through strategic delay, the continuation game collapses to complete information and surplus is split at the type-specific Rubinstein price $p^R(b,s)$. In the data, prices do respond to buyer valuation, but the response is attenuated. Deal prices compress toward a common level rather than tracking the type-specific equal split. This compression organizes the distributional findings. High-value buyers benefit because the compressed price lies below their type-specific equal split. Intermediate buyers lose because the same compressed price lies above it. Low-value buyers mostly avoid the mechanism altogether by terminating.

Implications for Modeling. For modeling, the implication is that the value of private information depends on how precisely bargaining translates private values into prices. When screening is precise, informed agents can convert private information into rents. When screening is incomplete, prices are compressed toward a pooling level, and the distributional consequences depend on where this pooling price lies relative to the type-specific complete-information benchmark. Private information can then benefit some informed types, harm others, and reduce the informed side's average surplus.

More broadly, our findings suggest a natural extension of canonical bargaining models. The model's three-region structure—termination, delay, and immediate acceptance—finds empirical support as an organizing principle: low-value types exit, high-value types trade quickly at favorable prices, and intermediate types face the most friction. However, the model overstates the precision of information transmission. The theory assumes that signaling fully reveals the buyer's intermediate type, so that the continuation game is one of complete information. The data are more consistent with a process in which information transmission is partial—enough to generate the qualitative pattern of price compression, but not enough to support the type-specific pricing that full revelation would imply. A model with imperfect or noisy signaling, allowing for residual uncertainty at the point of agreement, could therefore generate the kind of price compression we observe while preserving the core insight that private information shapes bargaining power through the process by which it is revealed.

10 Conclusion

We study bilateral bargaining with private information and transaction costs in a model-guided laboratory experiment. The model delivers sharp predictions for the one-sided information environment and generates a three-region equilibrium structure—termination, strategic delay, and immediate acceptance. This structure organizes important features of the data: low-value buyers often exit, high-value buyers trade quickly and earn information rents, and intermediate buyers face the most friction. At the same time, the experiment reveals a central departure from the equilibrium path: screening of intermediate buyer types is incomplete. Rather than fully revealing buyer types through

strategic delay, bargaining unfolds as rapid offer exchange that transmits information only partially into prices.

This partial information transmission generates price compression, the mechanism that ties our distributional findings together. Prices respond to buyer valuation, but less than the type-specific complete-information benchmark would imply. As a result, high-value buyers benefit because compressed prices lie below their type-specific equal-split benchmark, while intermediate buyers are harmed because the same compressed prices lie above their type-specific benchmark and transfer surplus to sellers. Transaction costs, by contrast, primarily affect the bargaining process: they accelerate negotiation and compress time without substantially altering who trades, at what price, or how the surplus is divided. In our setting, an exclusive information advantage therefore makes the informed party worse off overall.

The broader implication is that the value of private information depends on how precisely bargaining translates private values into prices. When screening is precise, informed agents can convert private information into rents. When screening is incomplete, prices are compressed toward a pooling level, and private information can benefit some informed types while harming others. This distinction refines, rather than overturns, the canonical theory. The model identifies the key economic force—price compression arising from incomplete screening—but overstates the precision with which bargaining transmits information into prices.

Our findings also suggest directions for future theory and experiments. A model with imperfect or noisy signaling, allowing for residual uncertainty at the point of agreement, could generate the kind of price compression we observe while preserving the core insight that private information shapes bargaining power through the bargaining process itself. Our experimental environment involves short negotiations in which the costs of delay are salient and continuously displayed. In longer-horizon settings—such as labor disputes, mergers and acquisitions, or legislative bargaining—waiting may play a more important informational role. Understanding when private information helps rather than hurts, and how this depends on the information structure, type distribution, bargaining protocol, and time scale of negotiation, remains an important question for future work.

References

Admati, Anat R. and Motty Perry, "Strategic Delay in Bargaining," *Review of Economic Studies*, 1987, *54* (3).

Andersen, Steffen, Glenn W Harrison, Morten I Lau, and E Elisabet Rutström, "Eliciting risk and time preferences," *Econometrica*, 2008, *76* (3), 583–618.

Ausubel, Lawrence M. and Raymond J. Deneckere, "Reputation in Bargaining and Durable Goods Monopoly," *Econometrica*, 1989, *57* (3), 511–531.

Ausubel, Lawrence M and Raymond J Deneckere, "Bargaining and the right to remain silent," *Econometrica: Journal of the Econometric Society*, 1992, pp. 597–625.

Ausubel, Lawrence M., Peter Cramton, and Raymond J. Deneckere, "Bargaining with Incomplete Information," in Robert J. Aumann and Sergiu Hart, eds., *Handbook of Game Theory with Economic Applications*, Vol. 3, Elsevier, 2002, chapter 50, pp. 1897–1945.

Backus, Matthew, Thomas Blake, and Steven Tadelis, "Sequential Bargaining in the Field: Evidence from Millions of Online Bargaining Interactions," *Quarterly Journal of Economics*, 2020, *135* (3), 1319–1362.

Binmore, Ken, Ariel Rubinstein, and Asher Wolinsky, "The Nash Bargaining Solution in Economic Modelling," *The RAND Journal of Economics*, 1986, *17* (2), 176–188.

_ , Avner Shaked, and John Sutton, "Testing Noncooperative Bargaining Theory: A Preliminary Study," *American Economic Review*, 1985, *75* (5), 1178–1180.

Bochet, Olivier, Manshu Khanna, and Simon Siegenthaler, "Beyond Dividing the Pie: Multi-Issue Bargaining in the Laboratory," *Review of Economic Studies*, 2024, *91* (1), 163–191.

Bolton, Gary E. and Axel Ockenfels, "ERC: A Theory of Equity, Reciprocity, and Competition," *American Economic Review*, 2000, *90* (1), 166–193.

Bursztyn, Leonardo, Ingar K Haaland, Nicolas Röver, and Christopher Roth, "The Social Desirability Atlas," Technical Report, National Bureau of Economic Research 2025.

Camerer, Colin F., *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton: Princeton University Press, 2003.

Chen, Daniel L., Martin Schonger, and Chris Wickens, "oTree—An open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 2016, *9* (C), 88–97.

Cramton, Peter C., "Dynamic Bargaining with Transaction Costs," *Management Science*, 1991, *37* (10), 1221–1233.

_ , "Strategic Delay in Bargaining with Two-Sided Uncertainty," *The Review of Economic Studies*, 1992, *59* (1), 205.

Dal Bó, Pedro and Guillaume R. Fréchette, "On the determinants of cooperation in infinitely repeated games: A survey," *Journal of Economic Literature*, 2018, *56* (1), 60–114.

De Quidt, Jonathan, Johannes Haushofer, and Christopher Roth, "Measuring and bounding experimenter demand," *American Economic Review*, 2018, *108* (11), 3266–3302.

Fehr, Ernst and Klaus M. Schmidt, "A Theory of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics*, 1999, *114* (3), 817–868.

Forsythe, Robert, John Kennan, and Barry Sopher, "An Experimental Analysis of Strikes in Bargaining Games with One-Sided Private Information," *American Economic Review*, 1991, *81* (1), 253–278.

Freyberger, Joachim and Bradley J. Larsen, "How Well Does Bargaining Work in Consumer Markets? A Robust Bounds Approach," *Econometrica*, 2025, *93* (1), 161–194.

Fudenberg, Drew and Jean Tirole, "Sequential Bargaining with Incomplete Information," *Review of Economic Studies*, 1983, *50* (2).

_ **and** _ , *Game Theory*, Cambridge, MA: MIT Press, 1991.

_ **, David Levine, and Jean Tirole**, "Infinite-Horizon Models of Bargaining with One-Sided Incomplete Information," in Alvin E. Roth, ed., *Game-Theoretic Models of Bargaining*, Cambridge: Cambridge University Press, 1985, pp. 73–98.

Galinsky, Adam D and Thomas Mussweiler, "First Offers as Anchors: The Role of Perspective-Taking and Negotiator Focus," *Journal of Personality and Social Psychology*, 2001, *81* (4), 657–669.

Gul, Faruk and Hugo Sonnenschein, "On Delay in Bargaining with One-Sided Uncertainty," *Econometrica*, 1988, *56* (3), 601–611.

_ **,** _ **, and Robert Wilson**, "Foundations of Dynamic Monopoly and the Coase Conjecture," *Journal of Economic Theory*, 1986, *39* (1), 155–190.

Güth, Werner, Rolf Schmittberger, and Bernd Schwarze, "An Experimental Analysis of Ultimatum Bargaining," *Journal of Economic Behavior and Organization*, 1982, *3* (4), 367–388.

Keniston, Daniel, Bradley J. Larsen, Shengwu Li, J.J. Prescott, Bernardo S. Silveira, and Chuan Yu, "Splitting the Difference in Incomplete-Information Bargaining: Theory and Widespread Evidence from the Field," *International Economic Review*, 2024, *65* (4), 1911–1939.

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

Myerson, Roger B. and Mark A. Satterthwaite, "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 1983, *29* (2), 265–281.

Rapoport, Amnon, Ido Erev, and Rami Zwick, "An Experimental Study of Buyer-Seller Negotiation with One-Sided Incomplete Information and Time Discounting," *Management Science*, 1995, *41* (3), 377–394.

Riley, John G. and Richard Zeckhauser, "Optimal Selling Strategies: When to Haggle, When to Hold Firm," *Quarterly Journal of Economics*, May 1983, *98* (2), 267–289.

Roth, Alvin E, "Axiomatic Models of Bargaining," *Lecture Notes in Economics and Mathematical Systems*, 1979, *170.*

_ **and J Keith Murnighan**, "Equilibrium behavior and repeated play of the prisoner's dilemma," *Journal of Mathematical psychology*, 1978, *17* (2), 189–198.

Roth, Alvin E., Vesna Prasnikar, Masahiro Okuno-Fujiwara, and Shmuel Zamir, "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study," *American Economic Review*, 1991, *81* (5), 1068–1095.

Rubinstein, Ariel, "Perfect equilibrium in a bargaining model," *Econometrica: Journal of the Econometric Society*, 1982, pp. 97–109.

Thorsteinson, Todd J, Jessica Breier, Allison Atwell, Crystal Hamilton, and Monica Privette, "Anchoring Effects on Performance Judgments," *Organizational Behavior and Human Decision Processes*, 2011, *115* (1), 109–117.

For online publication only:

Dynamic Bargaining with Incomplete Information: Theory and Experimental Evidence

Dongkyu Chang Peter Cramton Jeongbin Kim Axel Ockenfels Nicolas Roever

August 4, 2026

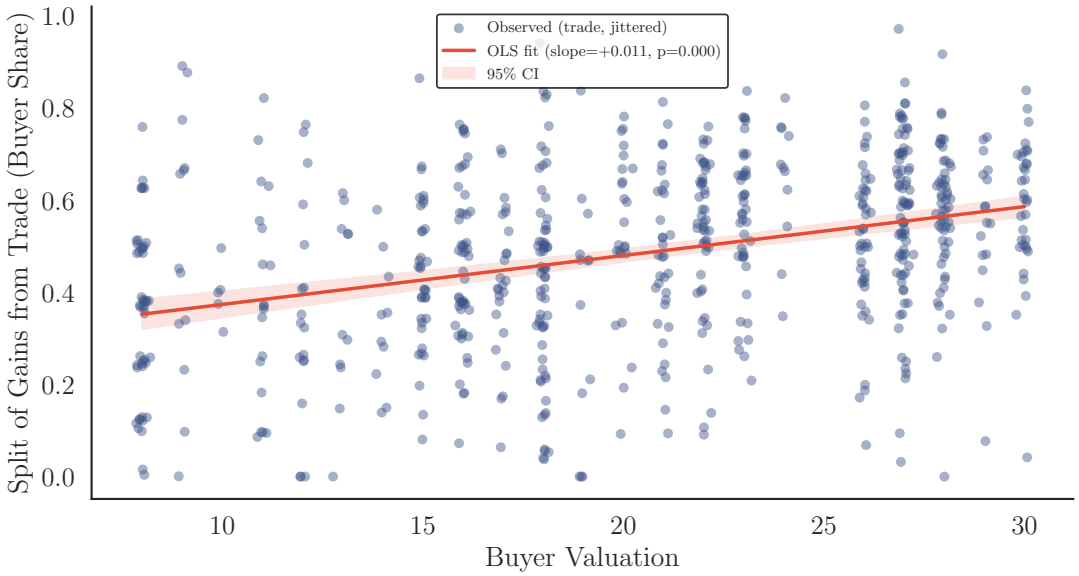
Summary of the Online Appendix

The Online Appendix provides supplementary descriptive evidence, robustness analyses, and implementation details.

- Appendix A presents additional figures.
- Appendix B contains additional tables.
- Appendix C describes the elicitation of behavioral preferences.
- Appendix D provides omitted theoretical details and proofs.
- Appendix E describes the calibration of the theoretical prediction.

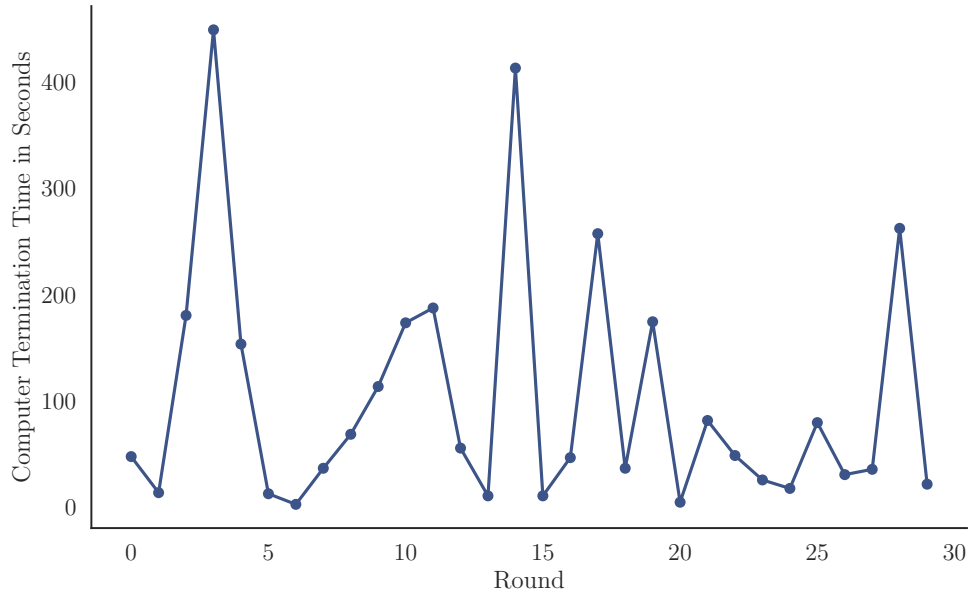
A Additional Figures

Figure A.1: Surplus for buyers in the *OneSidedCost* treatment



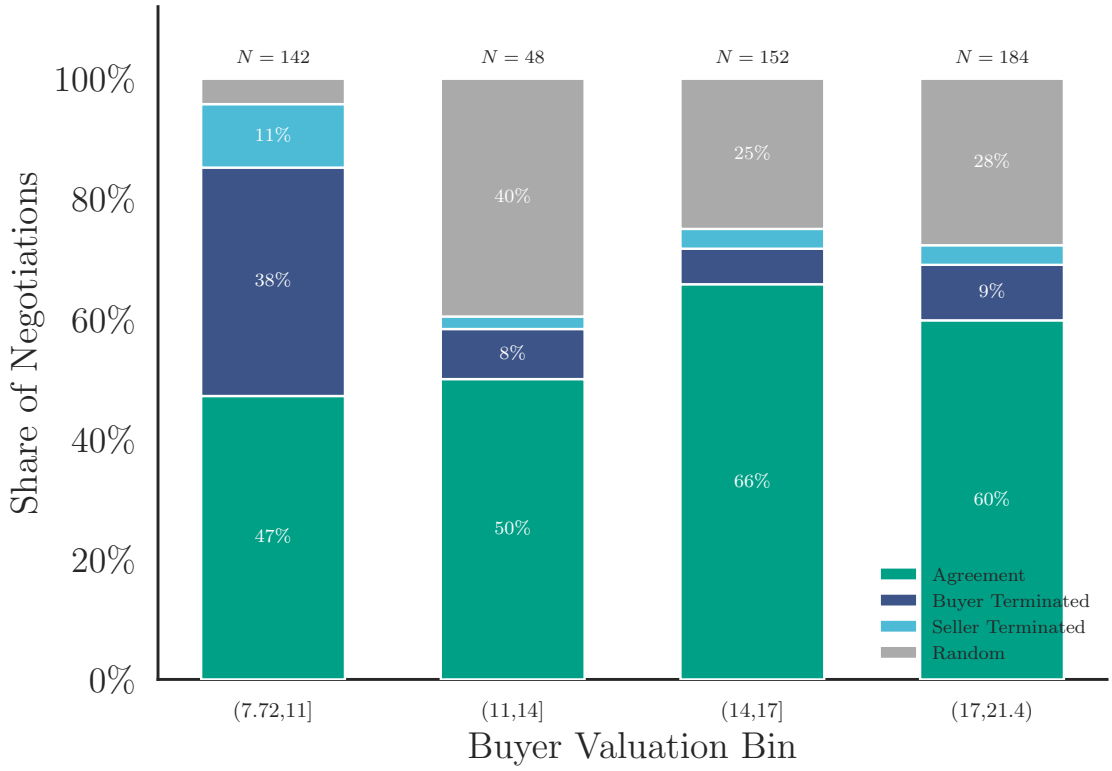
Note: The figure shows the share of the surplus extracted by buyers as a function of buyer valuation in the *OneSidedCost* treatment. The sample is restricted to negotiations that end in agreement with strictly positive gains from trade.

Figure A.2: Empirical draws of the computer termination times used in all treatments



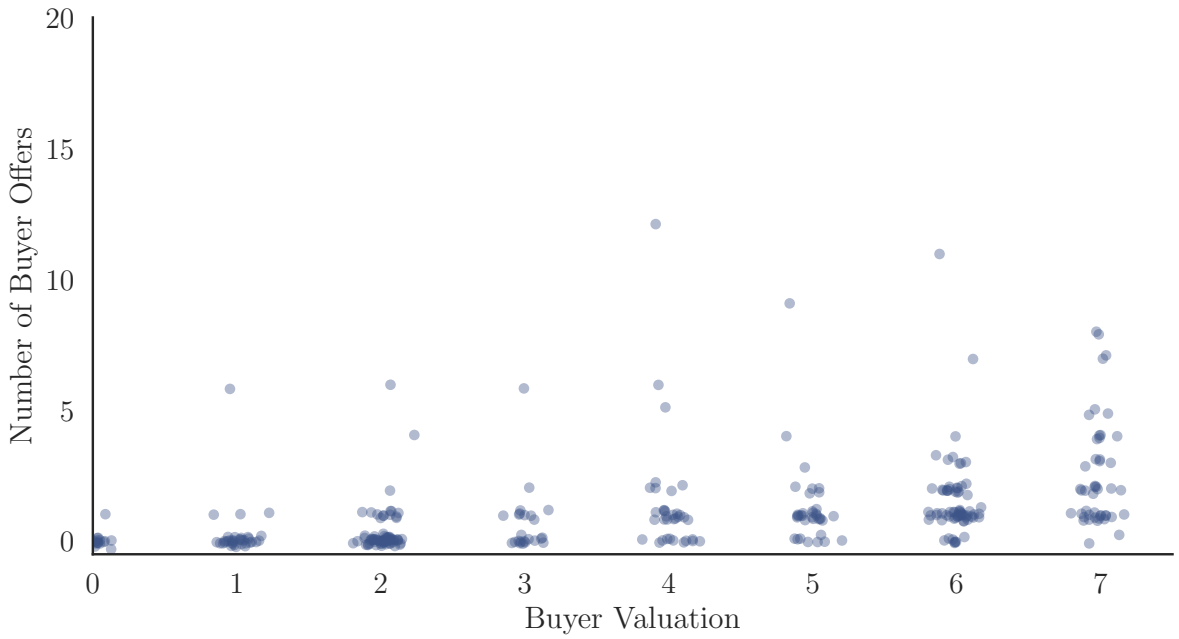
Note: The figure shows the distribution of computer-imposed termination times across all sessions and treatments. Termination times are drawn from the same geometric distribution with a 1% per-second termination probability applied to all matching groups.

Figure A.3: Bargaining outcomes for intermediate buyer types in the treatment with one-sided incomplete information and transaction costs



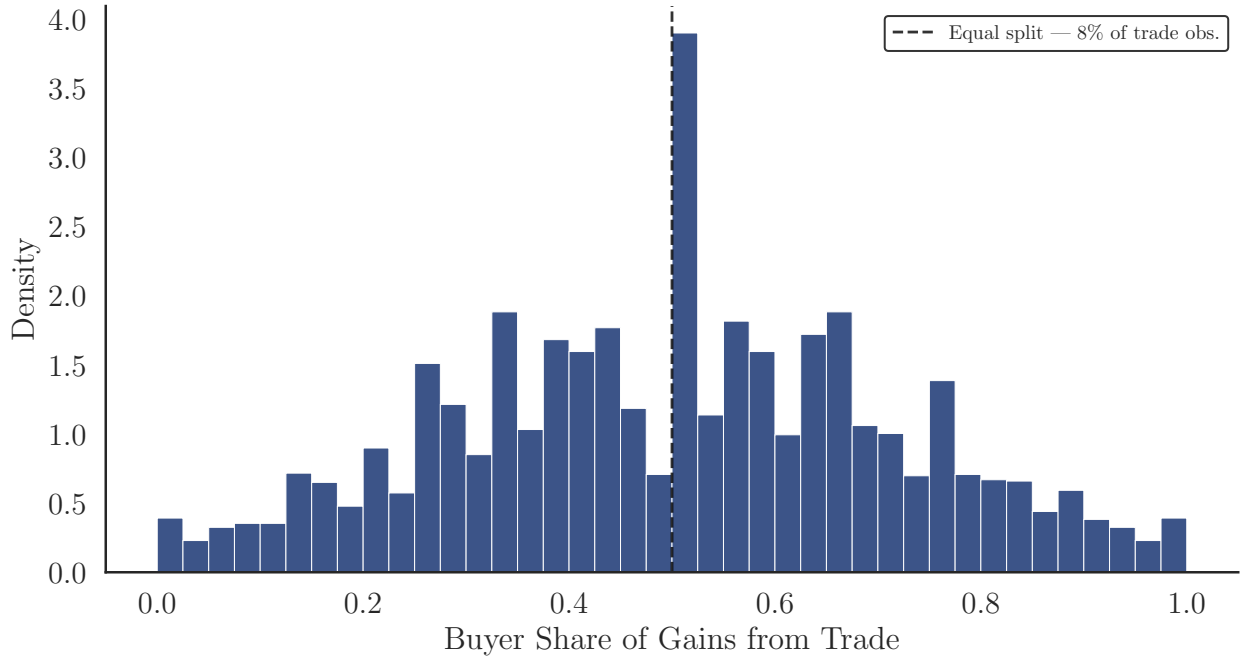
Note: The figure shows the composition of bargaining outcomes (agreement, buyer-initiated termination, seller-initiated termination, and computer termination) across buyer valuation bins in the intermediate region ($b^{\dagger} < b < b^{*}$) of the *OneSidedCost* treatment.

Figure A.4: Number of offers from buyers vs. valuation of buyers in treatment *OneSidedCost*



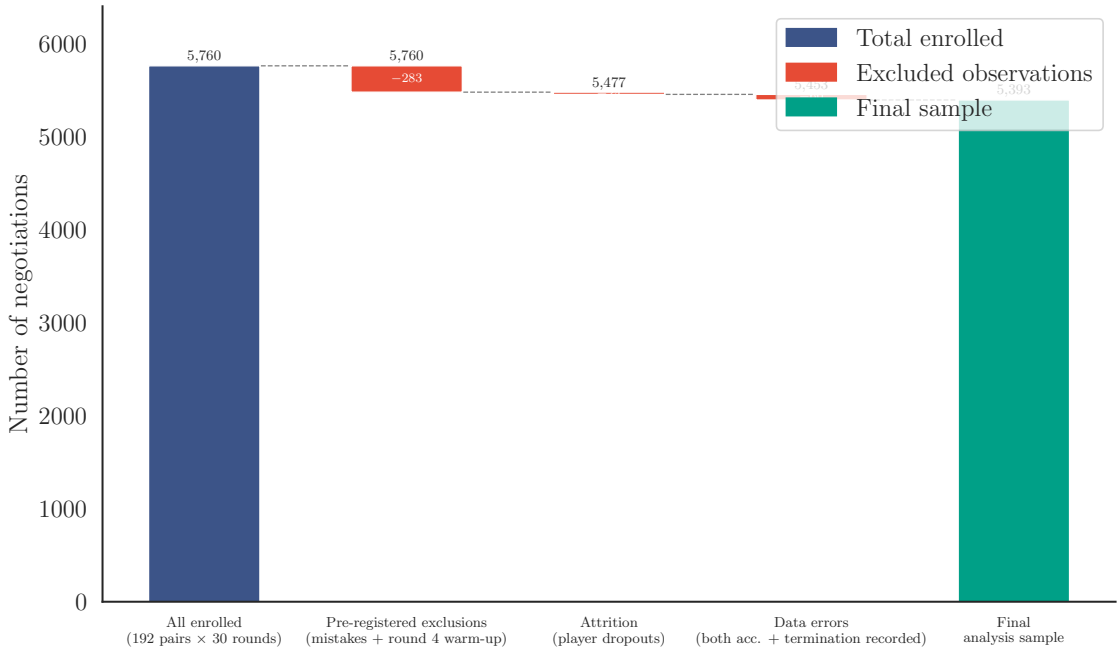
Note: The figure shows the number of offers made by buyers as a function of buyer valuation in the *OneSidedCost* treatment. Small random noise is added to facilitate visualization.

Figure A.5: Share of gains from trade for buyers pooled over all treatments



Note: The figure shows the distribution of the buyer's share of the gains from trade pooled across all four treatment cells. The sample is restricted to negotiations ending in agreement with strictly positive gains from trade.

Figure A.6: Sample construction



Note: The figure shows how the final analysis sample of 5,393 negotiations is derived from the 5,760 negotiations in the raw data (192 matched pairs × 30 rounds). Pre-registered exclusions remove negotiations in which at least one player's payoff fell below −15 and all round-1 warm-up negotiations. Attrition removes negotiations involving players who dropped out mid-session. We remove further negotiations in which both an acceptance and a termination were recorded simultaneously ("data errors"); this can happen if both players made their choice at roughly the same time.

B Additional Tables

Table B.1: Descriptive Statistics by Treatment

	T1 *No Cost,* *Symmetric Uncertainty*	T2 *Costly,* *Symmetric Uncertainty*	T3 *No Cost,* *Buyer-only Uncertainty*	T4 *Costly,* *Buyer-only Uncertainty*
Mean session duration (mins)	42.84	31.52	41.64	32.33
Mean age	28.42	26.64	26.70	25.68
Share females	0.51	0.60	0.54	0.59
Number participants	96	96	96	96
Number groups	12	12	12	12

Table B.2: Balance test: pre-treatment characteristics and treatment assignment

	TwoSidedNoCost	TwoSidedCost	OneSidedNoCost	OneSidedCost
Age	0.009**	-0.001	-0.001	-0.007***
	(0.004)	(0.004)	(0.004)	(0.003)
Share Female	-0.048	0.043	-0.019	0.024
	(0.046)	(0.046)	(0.046)	(0.045)
Risk Aversion	0.016	0.001	-0.018	0.002
	(0.019)	(0.019)	(0.018)	(0.018)
Ultimatum Offer	-0.002	0.002	0.000	-0.001
	(0.004)	(0.004)	(0.003)	(0.003)
Time Preference	0.008	0.008	-0.017	0.001
	(0.012)	(0.012)	(0.012)	(0.012)
F p-val	0.118	0.874	0.677	0.164
Obs	382	382	382	382
F-stat	1.770	0.362	0.630	1.581

Notes: The table reports OLS estimates of regressions of a binary treatment-indicator on given characteristics. Each column corresponds to one treatment: *TwoSidedNoCost* (two-sided private information, no transaction costs), *TwoSidedCost* (two-sided private information, transaction costs), *OneSidedNoCost* (buyer-only private information, no transaction costs), and *OneSidedCost* (buyer-only private information, transaction costs). The dependent variable equals one if the participant is assigned to the respective treatment and zero otherwise. *Age* is the participant's age in years. *Share Female* is an indicator equal to one for female participants. *Risk Aversion* is the switching point in a lottery task. *Ultimatum Offer* is the participant's offer in an ultimatum game. *Time Preference* is the switching point in a multiple price list eliciting discount rates. The F-statistic and its p-value test the joint null that no pre-treatment characteristic predicts treatment assignment. Heteroskedasticity-robust standard errors (HC3) in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.3: Comparison of actions in one- and two-sided incomplete information treatments

	One-sided		Two-sided	
	Buyer	Seller	Buyer	Seller
Who accepts	0.581***	0.419***	0.469	0.531
Who terminates	0.799***	0.201***	0.523	0.477
Who makes first offer	0.461***	0.539***	0.473**	0.527**

Stars: one-sample t-test against 0.5 ($^{*}p < .10$, $^{**}p < .05$, $^{***}p < .01$)

Table B.4: Mean deal price by buyer valuation bin (*OneSidedCost*, intermediate region $b \in (7.72, 21.40)$)

Buyer Valuation	**N**	**Mean Deal Price**	**Mean b/2**	**Diff**
$(7.72, 11]$	84	5.46	4.48	+0.99
$(11, 14]$	39	8.38	6.40	+1.99
$(14, 17]$	110	8.85	7.95	+0.90
$(17, 21.4)$	134	10.68	9.61	+1.07

Notes: This table shows all negotiations in the treatment with transaction costs where the buyer type is between 7.72 and 21.40 and where trade occurred.

Table B.5: Influence of risk aversion on bargaining behavior

	Accepted Other's Offer		Bargaining Time (sec)
	(1)	(2)	(3)
Risk Aversion	0.014**	-0.009	-0.681*
	(0.006)	(0.009)	(0.378)
Individual FE		✓	
Baseline Controls	✓	✓	✓
Obs	8717	8717	5723
Adj. R^2	0.124	0.187	0.115

Notes: The table reports OLS estimates of the relationship between risk aversion and bargaining behavior. The unit of observation is a participant-round. In columns (1) and (2), the dependent variable is *Accepted Other's Offer*, an indicator equal to one if the participant accepted the counterpart's offer in that round, and zero otherwise. In column (3), the dependent variable is *Bargaining Time*, measured in seconds from the start of the round until agreement or termination. The key independent variable, *Risk Aversion*, is from our endline survey, where higher values indicate greater risk aversion (see Appendix C). Baseline controls include treatment fixed effects, total gains from trade, and group-session fixed effects. Column (2) additionally includes individual fixed effects. Column (3) excludes the first four experimental sessions (June 11, 13, 16, and 17, 2025) due to a technical issue that affected the recording of timing data. Standard errors are clustered at the individual level and reported in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.6: Model fit comparison: buyer payoffs in treatment *OneSidedCost*

	Our Model	Linear Regression	Equal Split	Coase Conjecture	TILO
R^2	0.58	0.53	0.53	0.53	0.57
Avg. Residual	-1.71	0.00	-2.66	-10.14	1.02
Slope	0.80***	1.00***	1.09***	0.55***	1.03***

Notes: The table evaluates five models of buyer payoffs on the *OneSidedCost* subsample (buyer-only uncertainty with transaction costs). For each model, *actual payoff* is regressed on *predicted payoff* via OLS; the table reports the resulting R^2, the average residual (actual minus predicted), and the slope coefficient. Standard errors are clustered at the participant level. **Our Model** predicts buyer payoff using the theoretical equilibrium derived in this paper, which maps buyer valuation to payoff via a piecewise function reflecting the bargaining dynamics under transaction costs. **Linear Regression** is an OLS fit of actual payoff on buyer valuation; by construction it yields a slope of one and a zero average residual. **Equal Split** assigns the buyer exactly 50 % of the gains from trade. **Coase Conjecture** assigns the buyer the entire gains from trade, reflecting the benchmark in which a seller with market power cannot extract any surplus. **TILO** (Take-It-or-Leave-It, seller optimum) assigns the buyer payoff $\max(\text{valuation} - 15, 0)$, i.e. the payoff under a seller-optimal posted price of 15.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.7: Effect of information on buyer's realized surplus share

	No TA costs		TA costs	
	(1)	(2)	(3)	(4)
Two-sided Private Info	0.118**	0.043	0.095**	0.034
	(0.053)	(0.046)	(0.047)	(0.059)
Gains from trade		0.008**		0.008***
		(0.003)		(0.002)
GFT control	No	Yes	No	Yes
Obs	1056	1056	1097	1097
Adj. R^2	0.007	0.026	0.013	0.052

Notes: OLS regressions of the buyer's share of realised surplus for trades with positive gains from trade. The sample is split by transaction-cost condition, with columns 1 and 2 estimated on the no-TA-cost subsample and columns 3 and 4 on the TA-cost subsample. In each subsample, the first specification includes an indicator for two-sided private information (meaning that one-sided private information is the omitted category), and the second specification additionally controls for gains from trade. Standard errors are clustered at the participant level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C Elicitation of Behavioral Preferences

Our experiment is conducted in German; here we present the key instructions translated to English.[1] At the end of the experiment, we elicit the most important behavioral primitives for bargaining: risk preferences, time preferences, and fairness attitudes. We use a standard ultimatum game question where subjects are asked to split 50 Euros to elicit fairness preferences (Bochet et al., 2024). In the analysis, we include a dummy equal to 1 if the person decides to keep more than 25 Euros as an indicator of *reduced* fairness preference. We follow Bochet et al. (2024) to elicit risk attitudes and ask subjects to choose one among six lotteries with increasing risk. We use an elicitation following Andersen et al. (2008) to understand individuals' time preferences. Subjects are presented with a choice between receiving 100 Euros now and 100+X Euros in 6 months. The 6-month interest rate varies from 5% to 30%. The "switching point" at which subjects choose the later payment is indicative of how patient the subject is—more patient subjects switch earlier.

Risk Preferences For each of the following options, tell us which you would prefer.

Option	Lottery
1	80% chance of winning €2
2	70% chance of winning €3
3	60% chance of winning €4
4	50% chance of winning €5
5	40% chance of winning €6
6	30% chance of winning €7

Time Preferences Imagine, you can choose between receiving money now or in 6 months. For each of the following options, tell us which you would prefer to receive.

Table C.1: Time preference elicitation choices

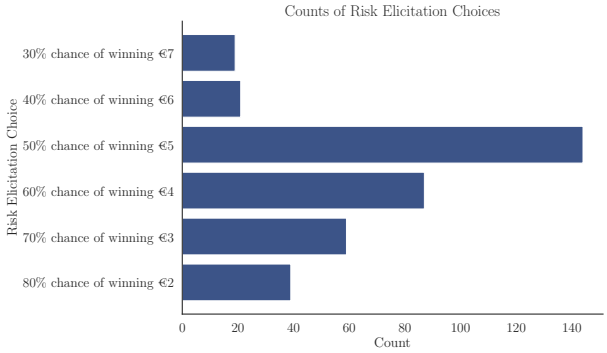
	Receive Now	Receive in 6 months
1)	€100	€105
2)	€100	€110
3)	€100	€115
4)	€100	€120
5)	€100	€125
6)	€100	€130

Fairness Preferences Imagine you receive 50 Euros, and you have to divide it between you and another randomly selected participant in the experiment. After you submit your offer, the other

[1] The full oTree program including the German instructions is available at `https://github.com/NicolasRoever/bargaining_experiment`.

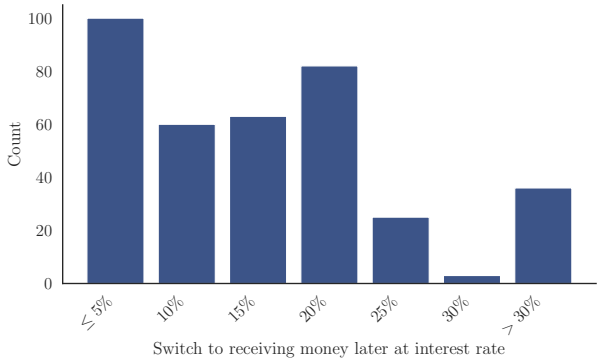
participant can accept or reject it. If they reject it, you and the other participant receive 0 euros. If they accept it, you both receive the money. What offer would you make?

Figure C.1: Risk elicitation choice distribution through entire sample



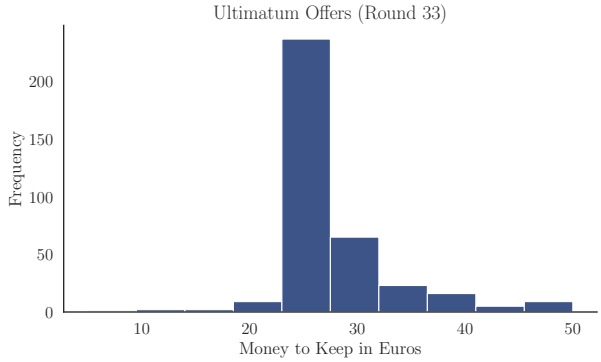
Note: The figure shows the distribution of lottery choices across all participants.

Figure C.2: Time preference distribution



Note: The x-axis values show the 6 month interest rate where the subject first reports preferring to receive the money later. Subjects who never switch are summarized under the $> 30\%$ group.

Figure C.3: Distribution of money kept out of 50 Euros in our entire sample



Note: The figure shows the distribution of the amount kept by the proposer in the ultimatum game elicitation task. A value of 25 corresponds to an equal split.

D Omitted details in Section 2

D.1 Proof of Proposition 1

Proof of (i). In this proof, we first consider the case $s = 0$. The general case $0 \leq s \leq \underline{B}$ will be discussed at the end of the proof. Let $\widetilde{v}_S(b^*, c) := v_S(b^* \mid b^\dagger(b^*, c))$ denote the seller's reduced objective, where $b^\dagger(b^*, c)$ denotes the cutoff type defined by (6). Under $s = 0$,

$$\widetilde{v}_S(b^*, c) = \frac{b^*}{2}\big(1 - F(b^*)\big) + \int_{b^\dagger(b^*,c)}^{b^*} \underbrace{\left[e^{-r\Delta(b|b^*)}\frac{b}{2} - \frac{c}{r}\big(1 - e^{-r\Delta(b|b^*)}\big)\right]}_{H(b;b^*,c)} f(b)\,\mathrm{d}b.$$

The delay function $\Delta(b|b^*)$ is given by

$$\Delta(b \mid b^*) := \int_b^{b^*} \frac{\partial p^R(z,s)/\partial z}{r[z - p^R(z,s)] + c}\,\mathrm{d}z = \frac{1}{r}\ln\left[\frac{r(b^* - s) + 2c}{r(b - s) + 2c}\right] \qquad b \in [b^\dagger, b^*). \tag{8}$$

Since the integrand is positive, $\Delta(b \mid b^*)$ is continuous and strictly decreasing in b. In Appendix D.2 we will prove that this delay function is indeed incentive compatible for the buyer types in $[b^\dagger, b^*)$; that is, these buyer types find it optimal to delay exactly until $t = \Delta(b|b^*)$.

Using (8), we can rewrite the integrand $H(b; b^*, c)$ as

$$H(b; b^*, c) = \frac{rb + 2c}{rb^* + 2c} \cdot \frac{b}{2} - \frac{c}{r} \cdot \frac{r(b^* - b)}{rb^* + 2c} = \frac{rb^2 + 4cb - 2cb^*}{2(rb^* + 2c)}.$$

Note that $H(b; b^*, c)$ coincides with the buyer's payoff $v_B(b \mid b^*)$ in (5). Define $\widehat{b}^\dagger(b^*, c)$ as the unique solution to $H(b; b^*, c) = 0$. Then the cutoff $b^\dagger(b^*, c)$ is $b^\dagger(b^*, c) = \max\{\underline{B}, \widehat{b}^\dagger(b^*, c)\}$. Equivalently,

$$b^\dagger(b^*, c) = \underline{B} \qquad \text{whenever} \qquad H(\underline{B}; b^*, c) \geq 0,$$

and otherwise $b^\dagger(b^*,c) > \underline{B}$ is characterized by

$$H(b^\dagger(b^*,c);b^*,c) = 0.$$

Next, we show that

$$\frac{\partial^2 \widetilde{v}_S(b^*,c)}{\partial c\,\partial b^*} \leq 0$$

in both cases: $H(\underline{B};b^*,c) \geq 0$ and $H(\underline{B};b^*,c) < 0$. In what follows, we suppress the arguments (b^*,c) from $b^\dagger(b^*,c)$ to simplify notation.

Case 1: $H(\underline{B};b^*,c) \geq 0$. Then $b^\dagger(b^*,c) = \underline{B}$, so the lower limit of integration is fixed. Differentiating $\widetilde{v}_S$ with respect to b^* and applying Leibniz's rule, the upper-boundary term $H(b^*;b^*,c) = b^*/2$ cancels with the derivative of the first term. Thus,

$$\frac{\partial \widetilde{v}_S(b^*,c)}{\partial b^*} = \frac{1-F(b^*)}{2} - \int_{\underline{B}}^{b^*} K(b;b^*,c) f(b)\,\mathrm{d}b,$$

where

$$K(b;b^*,c) := -\frac{\partial H(b;b^*,c)}{\partial b^*} = \frac{(rb+2c)^2}{2(rb^*+2c)^2}.$$

Differentiating once more with respect to c gives

$$\frac{\partial^2 \widetilde{v}_S(b^*,c)}{\partial c\,\partial b^*} = -\int_{\underline{B}}^{b^*} K_c(b;b^*,c) f(b)\,\mathrm{d}b \leq 0,$$

since

$$K_c(b;b^*,c) := \frac{\partial K(b;b^*,c)}{\partial c} = \frac{2r(b^*-b)(rb+2c)}{(rb^*+2c)^3} > 0 \qquad \forall b \in [\underline{B},b^*].$$

Case 2: $H(\underline{B};b^*,c) < 0$. Then $b^\dagger(b^*,c) > \underline{B}$ and is characterized by

$$H(b^\dagger(b^*,c);b^*,c) = 0.$$

Differentiating $\widetilde{v}_S$ with respect to b^* and applying Leibniz's rule, the lower-boundary term vanishes because $H(b^\dagger;b^*,c) = 0$, while the upper-boundary term $H(b^*;b^*,c) = b^*/2$ again cancels with the derivative of the first term. Hence

$$\frac{\partial \widetilde{v}_S(b^*,c)}{\partial b^*} = \frac{1-F(b^*)}{2} - \int_{b^\dagger}^{b^*} K(b;b^*,c) f(b)\,\mathrm{d}b,$$

where

$$K(b;b^*,c) = -\frac{\partial H(b;b^*,c)}{\partial b^*} = \frac{(rb+2c)^2}{2(rb^*+2c)^2}.$$

Differentiating once more with respect to c gives

$$\frac{\partial^2 \widetilde{v}_S(b^*,c)}{\partial c \partial b^*} = -\int_{b^\dagger}^{b^*} K_c(b;b^*,c) f(b)\,\mathrm{d}b + K(b^\dagger;b^*,c)\, f(b^\dagger)\,\frac{\partial b^\dagger}{\partial c},$$

where

$$K_c(b;b^*,c) = \frac{\partial K(b;b^*,c)}{\partial c} = \frac{2r(b^*-b)(rb+2c)}{(rb^*+2c)^3} > 0 \qquad \forall b \in [b^\dagger, b^*].$$

Adding and subtracting $f(b^\dagger)$ inside the integral yields

$$\int_{b^\dagger}^{b^*} f(b)K_c(b;b^*,c)\,\mathrm{d}b = \int_{b^\dagger}^{b^*} \big(f(b)-f(b^\dagger)\big)K_c(b;b^*,c)\,\mathrm{d}b + f(b^\dagger)\int_{b^\dagger}^{b^*} K_c(b;b^*,c)\,\mathrm{d}b.$$

Therefore

$$\begin{aligned}\frac{\partial^2 \widetilde{v}_S(b^*,c)}{\partial c \partial b^*} = &-\int_{b^\dagger}^{b^*} \big(f(b)-f(b^\dagger)\big)K_c(b;b^*,c)\,\mathrm{d}b \\ &- f(b^\dagger)\left[\int_{b^\dagger}^{b^*} K_c(b;b^*,c)\,\mathrm{d}b - K(b^\dagger;b^*,c)\frac{\partial b^\dagger}{\partial c}\right]. \end{aligned} \tag{9}$$

Here the first term on the right-hand side of (9) is weakly negative whenever f is nondecreasing, and $\partial b^\dagger/\partial c$ denotes the partial derivative with respect to c, holding b^* fixed.

It therefore remains to show that the bracketed term in (9) is strictly positive. First,

$$\begin{aligned}\int_{b^\dagger}^{b^*} K_c(b;b^*,c)\,\mathrm{d}b &= \frac{2r}{(rb^*+2c)^3}\int_{b^\dagger}^{b^*}(b^*-b)(rb+2c)\,\mathrm{d}b \\ &= \frac{2r}{(rb^*+2c)^3}\left[\frac{r(b^*)^3}{6} + c(b^*)^2 - \frac{rb^*(b^\dagger)^2}{2} - 2cb^*b^\dagger + \frac{r(b^\dagger)^3}{3} + c(b^\dagger)^2\right]. \end{aligned} \tag{10}$$

Using

$$H(b^\dagger;b^*,c) = 0 \quad \Longrightarrow \quad b^* = \frac{r(b^\dagger)^2+4cb^\dagger}{2c}, \quad rb^*+2c = \frac{(rb^\dagger+2c)^2}{2c}, \tag{11}$$

we can simplify (10) to

$$\int_{b^\dagger}^{b^*} K_c(b;b^*,c)\,\mathrm{d}b = \frac{r(b^\dagger)^2(rb^\dagger+6c)}{3(rb^\dagger+2c)^3}. \tag{12}$$

Moreover, since

$$H(b^\dagger;b^*,c) = 0 \quad \Longleftrightarrow \quad r(b^\dagger)^2 + 4cb^\dagger - 2cb^* = 0,$$

implicit differentiation with respect to c, while holding b^* fixed, gives

$$\frac{\partial b^\dagger}{\partial c} = \frac{r(b^\dagger)^2}{2c(rb^\dagger+2c)}.$$

Also, from (11),

$$rb^*+2c = \frac{(rb^\dagger+2c)^2}{2c},$$

and hence

$$K(b^\dagger;b^*,c)\frac{\partial b^\dagger}{\partial c} = \frac{(rb^\dagger+2c)^2}{2(rb^*+2c)^2}\cdot\frac{r(b^\dagger)^2}{2c(rb^\dagger+2c)} = \frac{cr(b^\dagger)^2}{(rb^\dagger+2c)^3}. \tag{13}$$

Hence, combining (12) and (13), we obtain

$$\int_{b^\dagger}^{b^*} K_c(b;b^*,c)\,\mathrm{d}b - K(b^\dagger;b^*,c)\frac{\partial b^\dagger}{\partial c} = \frac{r(b^\dagger)^2(rb^\dagger+6c)}{3(rb^\dagger+2c)^3} - \frac{cr(b^\dagger)^2}{(rb^\dagger+2c)^3}$$
$$= \frac{r(b^\dagger)^2(rb^\dagger+3c)}{3(rb^\dagger+2c)^3} > 0.$$

Therefore, the bracketed term in (9) is positive, and thus

$$\frac{\partial^2 \widetilde{v}_S(b^*,c)}{\partial c \partial b^*} \leq 0.$$

Combining Cases 1 and 2, we conclude that $\widetilde{v}_S$ has decreasing differences in (b^*,c). Therefore, by standard monotone comparative statics results such as Topkis's theorem, the seller's optimal-cutoff correspondence is weakly decreasing in c. If the maximizer is unique, then $b^*(c)$ is weakly decreasing in c.

The same argument extends to any $s \leq \underline{B}$ by replacing buyer types with surplus types $x := b - s$. Under symmetric costs,

$$p^R(b,s) = \frac{b+s}{2},$$

so the continuation payoff depends on (b,b^*,s) only through $(b-s,b^*-s)$. Defining

$$x := b-s, \qquad x^* := b^*-s, \qquad x^\dagger := b^\dagger - s,$$

we obtain

$$H(b;b^*,c,s) = \frac{rx^2+4cx-2cx^*}{2(rx^*+2c)},$$

which is exactly the same expression as in the case $s=0$. Since the density of x is $f(x+s)$, the assumption that f is nondecreasing in b implies that the density of x is nondecreasing in x. Therefore, the decreasing-differences argument applies unchanged.

Proof of (ii). Under $s=0$ and $c_B = c_S = c$, we have

$$p^R(b,0) = \frac{b}{2}, \qquad \frac{\partial p^R(z,0)}{\partial z} = \frac{1}{2}, \qquad r\big[z - p^R(z,0)\big] + c = \frac{rz}{2} + c.$$

Hence (8) becomes

$$\Delta(b \mid b^*(c)) = \int_b^{b^*(c)} \frac{\mathrm{d}z}{rz+2c}.$$

Fix $c_H > c_L > 0$, and let

$$b \in [b^\dagger(c_H), b^*(c_H)) \cap [b^\dagger(c_L), b^*(c_L)).$$

By part (i), the cutoff $b^*(c)$ is weakly decreasing in c, so

$$b^*(c_H) \leq b^*(c_L).$$

Therefore,

$$\Delta(b|b^*(c_H)) = \int_b^{b^*(c_H)} \frac{\mathrm{d}z}{rz+2c_H} < \int_b^{b^*(c_H)} \frac{\mathrm{d}z}{rz+2c_L} \leq \int_b^{b^*(c_L)} \frac{\mathrm{d}z}{rz+2c_L} = \Delta(b|b^*(c_L)).$$

The strict inequality in the second line follows from $c_H > c_L$ and the fact that the interval of integration has positive length, since $b < b^*(c_H)$ by assumption. This proves part (ii).

D.2 Equilibrium Strategy Profile and Beliefs

In this section, we complete the equilibrium description for the case with one-sided asymmetric information by specifying the strategy profile and beliefs after arbitrary histories. We state the strategies for the case without a minimum delay t_0; incorporating the delay constraint requires only straightforward modifications and is therefore omitted.

In what follows, the seller's posterior belief at time $\tau \geq 0$ refers to the posterior belief conditional on the history observed strictly before τ. For any $\underline{b}, \overline{b} \in [\underline{B}, \overline{B}]$ such that $\underline{b} \leq \overline{b}$, let $F(\cdot|\underline{b}, \overline{b})$ denote the truncation of the prior distribution F to $[\underline{b}, \overline{b}]$. If $\underline{b} < \overline{b}$, then

$$F(x|\underline{b}, \overline{b}) = \frac{F(x) - F(\underline{b})}{F(\overline{b}) - F(\underline{b})} \quad \forall x \in [\underline{b}, \overline{b}].$$

If $\underline{b} = \overline{b}$, $F(\cdot|\underline{b}, \overline{b})$ denotes the degenerate distribution that assigns probability one to the common value $\underline{b} = \overline{b}$. We also extend the delay function as follows. For any $\beta, b \in [\underline{B}, \overline{B}]$, define

$$\Delta(b|\beta) := \begin{cases} \displaystyle\int_b^{\beta} \frac{\partial p^R(z,s)/\partial z}{r[z - p^R(z,s)] + c} \mathrm{d}z = \frac{1}{r} \ln \left[\frac{r(\beta - s) + 2c}{r(b - s) + 2c} \right] & \text{if } b \in [\underline{B}, \beta], \\ 0 & \text{if } b \in (\beta, \overline{B}]. \end{cases} \tag{14}$$

Finally, throughout the strategy description and verification below, we focus on seller initial offers satisfying $p^{initial} \geq s$. This entails no loss if $s = 0$, since prices are nonnegative. If $s > 0$, we complete the off-path strategy profile as follows. After any seller initial offer $p^{initial} < s$, every buyer type accepts the offer immediately. If a history is nevertheless reached in which such an offer remains outstanding, the seller terminates the negotiation immediately. Since any accepted offer $p^{initial} < s$ gives the seller payoff $p^{initial} - s < 0$, while termination gives payoff zero, such an initial offer is never profitable for the seller.

We organize the equilibrium strategy profile and beliefs into three cases.

Case I: Inaction Phase. Consider a history in which the negotiation is still in the inaction phase at time t. That is, either $t = 0$, or neither player has made an offer or terminated the negotiation strictly before t.

- The buyer's behavioral strategy: All buyer types $b < b^\dagger$ terminate the negotiation immediately at time t; all buyer types $b \geq b^\dagger$ take no action at time t.
- The seller's behavioral strategy: The seller immediately offers $p = p^R(b^*, s)$ at time t.
- The seller's posterior belief: If $t = 0$, the seller's posterior belief coincides with the prior distribution F. If $t > 0$, the seller's posterior belief is $F(\cdot|b^\dagger, \overline{B})$.

Case II-(i): Standing-Offer Phase after the Seller's Initial Offer. This case is analogous to the equilibrium-path strategy described in the main text, except that the seller's initial offer is now arbitrary and need not coincide with the equilibrium offer $p^R(b^*, s)$. Consider a history in which the seller has made an initial offer $p^{initial}$ at time τ, and the buyer has not responded to this offer by time $t \geq \tau$.

Notation. We first define the cutoff type associated with the seller's initial offer:

$$\hat{b}^* := \begin{cases} \overline{B} & \text{if } p^{initial} > p^R(\overline{B}, s), \\ 2p^{initial} - s & \text{if } p^R(\underline{B}, s) \leq p^{initial} \leq p^R(\overline{B}, s), \\ \underline{B} & \text{if } p^{initial} < p^R(\underline{B}, s). \end{cases}$$

Thus, in the middle case, $p^{initial} = p^R(\hat{b}^*, s)$. Next, define the lowest type willing to enter the separating-delay region after the seller's initial offer by

$$\tilde{b}^\dagger := \inf \left\{ b \in [\underline{B}, \hat{b}^*] : e^{-r\Delta(b|\hat{b}^*)} \big(b - p^R(b, s)\big) - \int_0^{\Delta(b|\hat{b}^*)} e^{-rx} c \, \mathrm{d}x \geq 0 \right\}.$$

This set is nonempty because $\hat{b}^*$ itself belongs to it. The effective lower cutoff after accounting for both the initial inaction phase and the seller's initial offer is

$$\hat{b}^\dagger := \min\{\hat{b}^*, \max\{b^\dagger, \tilde{b}^\dagger\}\}.$$

Let $T^\dagger := \tau + \Delta(\hat{b}^\dagger|\hat{b}^*)$. For each $t \in [\tau, T^\dagger]$, define $b(t)$ as the unique buyer type in $[\hat{b}^\dagger, \hat{b}^*]$ satisfying $t = \tau + \Delta(b(t)|\hat{b}^*)$. Note that $b(\tau) = \hat{b}^*$ and $b(T^\dagger) = \hat{b}^\dagger$. We also need one auxiliary payoff for off-path histories. Define

$$W_t^\dagger(b) := \begin{cases} e^{-r(T^\dagger - t)} \big(b - p^R(\hat{b}^\dagger, s)\big) - \displaystyle\int_0^{T^\dagger - t} e^{-rx} c \, \mathrm{d}x & \text{if } t \leq T^\dagger, \\ b - p^R(\hat{b}^\dagger, s) & \text{if } t > T^\dagger. \end{cases}$$

The term $W_t^\dagger(b)$ is the continuation payoff of buyer type b from mimicking type $\hat{b}^\dagger$ from time t onward. For any fixed t, it is strictly increasing in b.

The Buyer's Strategy. We first specify the strategies of buyer types $b \geq \hat{b}^\dagger$.

- At time $t = \tau$, each such buyer type satisfying $p^{initial} \leq p^R(b,s)$ accepts the seller's initial offer $p^{initial}$. Buyer types who do not accept follow the continuation rules below.
- At any time $t \in [\tau, T^\dagger]$ before the buyer has responded to the seller's initial offer, each buyer type $b \geq \hat{b}^\dagger$ who has not yet responded behaves as follows. If $b \geq b(t)$, this buyer type makes the counteroffer $p = p^R(b(t),s)$. If $b < b(t)$, this buyer type takes no action at time t.
- At any time $t > T^\dagger$ before the buyer has responded to the seller's initial offer, each buyer type $b \geq \hat{b}^\dagger$ compares the payoff from making the counteroffer $p = p^R(\hat{b}^\dagger,s)$ with the payoff from accepting the seller's standing offer $p^{initial}$. These payoffs are $W_t^\dagger(b) = b - p^R(\hat{b}^\dagger,s)$ and $b - p^{initial}$, respectively. If $W_t^\dagger(b) > b - p^{initial}$, this buyer type makes the counteroffer $p = p^R(\hat{b}^\dagger,s)$. If $W_t^\dagger(b) \leq b - p^{initial}$, she accepts the seller's standing offer $p^{initial}$.

Now consider any buyer type $b < \hat{b}^\dagger$ who has not yet responded by time $t \geq \tau$. If $W_t^\dagger(b) < 0$, this buyer type terminates the negotiation immediately. If $W_t^\dagger(b) \geq 0$, this buyer type mimics type $\hat{b}^\dagger$: she takes no action when $t < T^\dagger$ and makes the counteroffer $p = p^R(\hat{b}^\dagger,s)$ when $t \geq T^\dagger$.

The Seller's Strategy and Beliefs. The seller takes no action while her initial offer is outstanding. Conditional on no buyer response strictly before time t, if $t \in [\tau, T^\dagger]$, the seller holds the posterior belief $F(\cdot|\hat{b}^\dagger, b(t))$. If $t > T^\dagger$, the seller believes that the buyer is of type $\hat{b}^\dagger$ with probability one.

Case II-(ii): Continuation Games after the Buyer's Counteroffer. Consider a history in which the negotiation was previously in Case II-(i), and the buyer makes a counteroffer q at time $t \geq \tau$. There are two subcases.

- Subcase A: Suppose $t \in [\tau, T^\dagger]$. In this case, the seller believes that the buyer is of type $b = b(t)$.
- Subcase B: Suppose $t > T^\dagger$. In this case, the seller assigns probability one to buyer type $\hat{b}^\dagger$.

In both subcases, the seller's belief reduces the continuation game to the standard Rubinstein bargaining game with complete information. Let $\theta = b(t)$ in Subcase A and $\theta = \hat{b}^\dagger$ in Subcase B. Since the buyer's counteroffer q is the standing offer, the seller accepts it whenever $q \geq p^R(\theta,s)$. If $q < p^R(\theta,s)$, the seller does not accept the counteroffer, and the continuation game yields trade at $p^R(\theta,s)$. Thus, along the prescribed continuation path, where the buyer offers $q = p^R(\theta,s)$, the outcome in Subcase A is immediate trade at price $p^R(b(t),s)$, and the outcome in Subcase B is immediate trade at price $p^R(\hat{b}^\dagger,s)$. We omit the remaining details of the complete-information continuation strategies.

Case III: Buyer-Initiated Initial Offer. Consider a history in which the buyer makes the initial offer q at time $t \geq 0$. By Case I, the seller makes the initial offer on the equilibrium path, so any history in which the buyer makes the initial offer is off the equilibrium path. Since Bayes' rule does not restrict the seller's belief after such a history, we specify that the seller assigns probability one to buyer type $\overline{B}$.

Let $\bar{p} := p^R(\overline{B}, s)$. Since the buyer's offer q is the standing offer, the seller accepts it if $q \geq \bar{p}$. If $q < \bar{p}$, the seller does not accept the offer and instead follows the complete-information continuation strategy corresponding to the belief that the buyer's type is $\overline{B}$, which yields the continuation price $\bar{p}$. A buyer of actual type b accepts this continuation price if $b \geq \bar{p}$ and terminates otherwise.

D.3 Verification of Equilibrium

We verify that the strategy profile and beliefs described in Section D.2 constitute a PBE.

The Buyer's Sequential Rationality in Case II-(ii). Let θ denote the buyer type inferred by the seller from the timing of the counteroffer: $\theta = b(t)$ in Subcase A and $\theta = \hat{b}^\dagger$ in Subcase B. Suppose the buyer makes a counteroffer q. Under the continuation strategies specified in Case II-(ii), if $q \geq p^R(\theta, s)$, the seller accepts the counteroffer; if $q < p^R(\theta, s)$, the seller does not accept and play continues according to the complete-information equilibrium, yielding trade at $p^R(\theta, s)$. Therefore, conditional on the seller's belief θ, choosing $q > p^R(\theta, s)$ gives a strictly lower payoff than choosing $q = p^R(\theta, s)$, while choosing $q < p^R(\theta, s)$ cannot improve on the complete-information continuation outcome. Hence, the prescribed counteroffer $q = p^R(\theta, s)$ is optimal.

The Buyer's Sequential Rationality in Case II-(i). Next, we verify the buyer's sequential rationality in Case II-(i). Fix a seller initial offer $p^{initial}$ made at time τ, and let $\hat{b}^*$, $\tilde{b}^\dagger$, and $\hat{b}^\dagger$ be defined as above. We focus on the case in which $\underline{B} < \hat{b}^\dagger < \hat{b}^* < \overline{B}$. That is, both $\hat{b}^\dagger$ and $\hat{b}^*$ are interior, and the separating-delay region does not degenerate. The proof for the remaining cases is similar and is therefore omitted.

First, focus on buyer types $b \geq \hat{b}^\dagger$; the sequential rationality of buyer types $b < \hat{b}^\dagger$ will be discussed later. Since $\hat{b}^*$ is interior, $p^{initial} = p^R(\hat{b}^*, s)$. Let $\hat{v}_B(b)$ denote the buyer's continuation payoff from time τ, gross of any costs incurred before τ, when buyer type b follows the prescribed strategy after the seller's initial offer. Using the extended definition of $\Delta(\cdot|\hat{b}^*)$ in (14), this payoff can be written as

$$\hat{v}_B(b) = e^{-r\Delta(b|\hat{b}^*)}\big(b - p^R(\min\{b, \hat{b}^*\}, s)\big) - \int_0^{\Delta(b|\hat{b}^*)} e^{-rx} c \, \mathrm{d}x \quad \forall b \in [\hat{b}^\dagger, \overline{B}]. \tag{15}$$

For $b \in [\hat{b}^\dagger, \hat{b}^*)$, this is the payoff from waiting until time $\tau + \Delta(b|\hat{b}^*)$ and then making the counteroffer $p = p^R(b, s)$. For $b \in [\hat{b}^*, \overline{B}]$, since $\Delta(b|\hat{b}^*) = 0$, this reduces to the payoff from accepting the seller's initial offer immediately, $b - p^R(\hat{b}^*, s)$.

The payoff function satisfies the single-crossing property, as in Cramton (1991). Moreover, by direct calculation using the definition of $\Delta(\cdot|\hat{b}^*)$, $\hat{v}_B$ is continuous and satisfies the envelope condition

$$\frac{\mathrm{d}\hat{v}_B(b)}{\mathrm{d}b} = e^{-r\Delta(b|\hat{b}^*)}$$

for almost every $b \in [\hat{b}^\dagger, \overline{B}]$. Also, the boundary participation condition $\hat{v}_B(\hat{b}^\dagger) \geq 0$ holds because $\hat{b}^\dagger \geq \tilde{b}^\dagger$, and $\tilde{b}^\dagger$ is the lowest type willing to enter the separating-delay region after the seller's initial offer. Together with single crossing and the monotonicity of $\Delta(\cdot|\hat{b}^*)$, this establishes incentive compatibility and participation for all buyer types $b \geq \hat{b}^\dagger$ at time τ.

Now consider any later history in Case II-(i), at time $t \in (\tau, T^\dagger]$, at which the buyer has not yet responded. Conditional on reaching such a history, the continuation problem is the same as the one above with the current upper remaining type $b(t)$ in place of $\hat{b}^*$. The seller's standing offer $p^{initial} = p^R(\hat{b}^*, s)$ is weakly less attractive to the buyer than the corresponding current cutoff price $p^R(b(t), s)$, because $b(t) \leq \hat{b}^*$. Hence, by the same single-crossing and envelope argument, each buyer type $b \geq \hat{b}^\dagger$ optimally responds immediately if $b \geq b(t)$ and optimally waits otherwise.

Next, consider a history in Case II-(i) with $t > T^\dagger$. At such a history, the seller's belief is degenerate at $\hat{b}^\dagger$. If a buyer type $b \geq \hat{b}^\dagger$ makes the counteroffer $p = p^R(\hat{b}^\dagger, s)$, the seller accepts it, and the buyer obtains $W_t^\dagger(b) = b - p^R(\hat{b}^\dagger, s)$. Alternatively, the buyer can accept the seller's standing offer and obtain $b - p^{initial}$. Termination is weakly dominated by making the counteroffer because $W_t^\dagger(b) \geq 0$ for all $b \geq \hat{b}^\dagger$. No further delay is useful because the seller's belief is already degenerate. Thus, the prescribed comparison between $W_t^\dagger(b)$ and $b - p^{initial}$ is sequentially rational.

Finally, consider any buyer type $b < \hat{b}^\dagger$ at a history in Case II-(i) at which she has not yet responded. By the same single-crossing argument, among all non-terminating deviations, including accepting the seller's standing offer, mimicking $\hat{b}^\dagger$ is optimal for such a type. The payoff from mimicking $\hat{b}^\dagger$ is precisely $W_t^\dagger(b)$. Hence, termination is optimal if $W_t^\dagger(b) < 0$, while mimicking $\hat{b}^\dagger$ is optimal if $W_t^\dagger(b) \geq 0$. This verifies the prescribed behavior of buyer types $b < \hat{b}^\dagger$ in Case II-(i).

The Buyer's Sequential Rationality in Cases I and III. The remaining cases are straightforward. Again, we focus on the case in which both b^* and $b^\dagger$ are interior, with $b^\dagger < b^*$; the proof for the other cases is essentially the same.

First, consider Case III. Let $\bar{p} := p^R(\overline{B}, s)$. If the buyer makes an initial offer $q \geq \bar{p}$, the seller accepts it immediately, so the buyer has no further decision to make. If the buyer makes an initial offer $q < \bar{p}$, the seller rejects it and follows the complete-information continuation strategy corresponding to buyer type $\overline{B}$, which yields the continuation price $\bar{p}$. Hence, a buyer of actual type b optimally accepts this continuation price if $b \geq \bar{p}$ and terminates otherwise. This verifies the buyer's sequential rationality after histories in Case III.

Next, consider Case I. At any inaction history at time t, if the buyer does not terminate or make the initial offer, she correctly expects that the seller will immediately offer $p^R(b^*, s)$. The equilibrium

continuation payoff is the same as the one obtained by substituting $\tau = 0$, $\hat{b}^* = b^*$, and $\hat{b}^\dagger = b^\dagger$ in (15) and can be written as follows:

$$\tilde{v}_B(b) := \begin{cases} e^{-r\Delta(b|b^*)}\big(b - p^R(\min\{b, b^*\}, s)\big) - \displaystyle\int_0^{\Delta(b|b^*)} e^{-rx} c \, \mathrm{d}x & \forall b \in [b^\dagger, \overline{B}], \\ 0 & \forall b \in [\underline{B}, b^\dagger). \end{cases}$$

By the same argument used above for Case II-(i), buyer types $b \geq b^\dagger$ optimally wait for the seller's initial offer, while buyer types $b < b^\dagger$ optimally terminate.

A buyer also cannot profit by making the initial offer. Under the continuation strategies specified in Case III, the best payoff a buyer of type b can obtain after making the initial offer is $\max\{b - \bar{p}, 0\}$. Indeed, offers $q \geq \bar{p}$ are accepted by the seller immediately and yield the buyer payoff at most $b - \bar{p}$; offers $q < \bar{p}$ are not accepted by the seller and lead either to trade at $\bar{p}$ or to termination.

We compare this payoff with the buyer's equilibrium payoff $\tilde{v}_B$ as defined above. Note that $\tilde{v}_B(b) \geq 0$ by construction. First, consider buyer types below $b^\dagger$. By direct calculation,

$$\tilde{v}_B(b^\dagger) = \frac{r(b^\dagger - s)^2 + 4c(b^\dagger - s) - 2c(b^* - s)}{2(r(b^* - s) + 2c)} = 0,$$

which implies

$$4c(b^\dagger - s) = 2c(b^* - s) - r(b^\dagger - s)^2 < 2c(b^* - s) \quad \Longrightarrow \quad b^\dagger - s < \frac{b^* - s}{2} \leq \frac{\overline{B} - s}{2} = \bar{p} - s.$$

Therefore, $b < \bar{p}$ for all $b < b^\dagger$, so $\max\{b - \bar{p}, 0\} = 0 = \tilde{v}_B(b)$ for all $b < b^\dagger$. For buyer types $b \geq b^\dagger$ with $b \leq \bar{p}$, we have $\max\{b - \bar{p}, 0\} = 0$, while $\tilde{v}_B(b) \geq 0$. Finally, for buyer types $b \geq b^\dagger$ with $b > \bar{p}$, the equilibrium payoff $\tilde{v}_B(b)$ must be weakly larger than the payoff from accepting $p^R(b^*, s)$ immediately. Thus, $\tilde{v}_B(b) \geq b - p^R(b^*, s) \geq b - p^R(\overline{B}, s) = b - \bar{p} = \max\{b - \bar{p}, 0\}$. This completes the proof that no buyer type can obtain a higher payoff by making the initial offer.

The Seller's Sequential Rationality in Case II-(i). We next verify the seller's sequential rationality. First, consider Case II-(i). Fix a seller initial offer $p^{initial}$ made at time τ. While the seller's initial offer is outstanding, the seller can either wait or terminate the negotiation. Termination gives continuation payoff zero. We show that waiting gives a weakly nonnegative continuation payoff.

First suppose $t \in [\tau, T^\dagger]$. Given the seller's posterior belief $F(\cdot|\hat{b}^\dagger, b(t))$, any buyer type still possible at time t lies in $[\hat{b}^\dagger, b(t)]$. For such a type b, the buyer is prescribed to respond at time $\tau + \Delta(b|\hat{b}^*)$. Hence, conditional on buyer type b, the seller's continuation payoff from waiting is

$$e^{-r[\Delta(b|\hat{b}^*) - (t-\tau)]}\big(p^R(b, s) - s\big) - \int_0^{\Delta(b|\hat{b}^*) - (t-\tau)} e^{-rx} c \, \mathrm{d}x.$$

To see that this payoff is weakly positive, note that each buyer type $b \in [\hat{b}^\dagger, b(t)]$ obtains a weakly

nonnegative continuation payoff from following the separating-delay path. Since costs incurred before t are sunk, the buyer's continuation payoff from time t onward is also weakly nonnegative. At the response time, trade occurs at $p^R(b,s)$, so the buyer's and seller's gross gains from trade are equal: $p^R(b,s)-s = b-p^R(b,s)$. Because both parties face the same transaction cost over the same remaining delay, the seller's continuation payoff from waiting is also weakly nonnegative. Therefore, taking expectations over $F(\cdot|\hat{b}^\dagger, b(t))$, the seller's expected continuation payoff from waiting is weakly nonnegative.

Now suppose $t > T^\dagger$. In this case, the seller's belief is degenerate at $\hat{b}^\dagger$. Under the prescribed buyer strategy, a buyer of type $\hat{b}^\dagger$ either accepts the seller's standing offer $p^{initial}$ or makes the counteroffer $p = p^R(\hat{b}^\dagger, s)$ immediately. In either case, using the maintained restriction $p^{initial} \geq s$, the seller obtains a weakly nonnegative continuation payoff. Hence, waiting is again weakly optimal relative to termination.

The Seller's Sequential Rationality in Case II-(ii). Next, consider Case II-(ii). Let θ denote the buyer type inferred by the seller from the timing of the counteroffer: $\theta = b(t)$ in Subcase A and $\theta = \hat{b}^\dagger$ in Subcase B. The buyer's counteroffer q is the standing offer. If the seller accepts q, her payoff is $q-s$. If instead she does not accept and continues according to the complete-information equilibrium, her payoff is $p^R(\theta,s)-s$. Since $p^R(\theta,s)-s \geq 0$, continuing also weakly dominates termination. Hence, accepting the counteroffer is optimal whenever $q \geq p^R(\theta,s)$, while not accepting and continuing according to the complete-information equilibrium is optimal whenever $q < p^R(\theta,s)$. This verifies the seller's sequential rationality in Case II-(ii).

The Seller's Sequential Rationality in Case I. Next, consider Case I. At time $t=0$, the seller's posterior belief coincides with the prior distribution F. Initial offers below $p^R(\underline{B},s)$ are weakly dominated by $p^R(\underline{B},s)$, and initial offers above $p^R(\overline{B},s)$ are weakly dominated by $p^R(\overline{B},s)$. Therefore, it is without loss to restrict attention to offers of the form $p^R(\tilde{b}^*,s)$ with $\tilde{b}^* \in [\underline{B},\overline{B}]$. Given such a cutoff, the continuation strategies described in Case II-(i) induce the seller's expected payoff $v_S(\tilde{b}^*|b^\dagger(\tilde{b}^*))$, as defined in the main text. By (7), b^* maximizes this objective over $\tilde{b}^* \in [\underline{B},\overline{B}]$. Hence, the prescribed offer $p^R(b^*,s)$ is optimal at $t=0$.

Now consider an inaction history at time $t>0$. The seller's posterior belief is $F(\cdot|b^\dagger,\overline{B})$. We claim that $p^R(b^*,s)$ remains optimal under this posterior. Let $\pi_S(b|\tilde{b}^*)$ denote the seller's continuation payoff conditional on buyer type b when the seller chooses cutoff $\tilde{b}^*$ and continuation play follows the strategies specified in Case II-(i). For any $a \in [\underline{B},\overline{B}]$, define

$$\Pi(a,\tilde{b}^*) := \int_a^{\overline{B}} \pi_S(b|\tilde{b}^*)\,\mathrm{d}F(b).$$

Under the posterior belief $F(\cdot|b^\dagger,\overline{B})$, the seller's expected payoff from choosing cutoff $\tilde{b}^*$ is $\Pi(b^\dagger,\tilde{b}^*)/[F(\overline{B})-F(b^\dagger)]$. Since $F(\overline{B})-F(b^\dagger)$ is independent of $\tilde{b}^*$, it is enough to compare the Π terms. For any $\tilde{b}^*$, $\Pi(b^\dagger,\tilde{b}^*) \leq \Pi(\underline{B},\tilde{b}^*)$, because buyer types below $b^\dagger$ generate weakly

nonnegative seller payoff under the continuation strategies as shown in Case II-(i). For the cutoff b^*, buyer types below $b^\dagger$ terminate and generate zero seller payoff, so $\Pi(b^\dagger, b^*) = \Pi(\underline{B}, b^*)$. Since b^* maximizes the seller's ex-ante expected payoff under the prior, it follows that, for every $\tilde{b}^*$, $\Pi(b^\dagger, \tilde{b}^*) \leq \Pi(\underline{B}, \tilde{b}^*) \leq \Pi(\underline{B}, b^*) = \Pi(b^\dagger, b^*)$. Thus, b^* also maximizes the seller's payoff under the posterior $F(\cdot|b^\dagger, \overline{B})$, and the prescribed offer $p^R(b^*, s)$ is optimal at any inaction history with $t > 0$.

Finally, the seller cannot profit by delaying her initial offer. Suppose the seller waits until some later time and then makes an initial offer p. The seller could instead make the same offer immediately. Doing so induces the same continuation behavior from all buyer types who would still be present later, but it reaches them earlier and avoids the discounting and transaction costs incurred during the additional delay. Hence delaying the initial offer cannot yield a higher payoff than making an initial offer immediately.

The Seller's Sequential Rationality in Case III. In Case III, after a buyer-initiated offer q, the seller assigns probability one to buyer type $\overline{B}$. Let $\bar{p} := p^R(\overline{B}, s)$. Accepting the standing offer q gives the seller payoff $q - s$, while not accepting and following the complete-information continuation strategy corresponding to buyer type $\overline{B}$ gives payoff $\bar{p} - s$. Since $\bar{p} - s \geq 0$, continuing also weakly dominates termination. Hence, accepting is optimal whenever $q \geq \bar{p}$, and not accepting is optimal whenever $q < \bar{p}$. This verifies the seller's sequential rationality in Case III.

Consistency of the Seller's Beliefs. In Case I, if $t = 0$, the seller's posterior belief coincides with the prior distribution F. If $t > 0$, then observing that no buyer has terminated before time t eliminates all buyer types below $b^\dagger$, while all types $b \geq b^\dagger$ remain inactive under the prescribed strategy. Hence Bayes' rule yields the posterior belief $F(\cdot|b^\dagger, \overline{B})$.

Next, consider Case II-(i), after the seller makes an initial offer and the buyer has not yet responded. Consider a history with $t \in [\tau, T^\dagger]$ at which the seller's initial offer is still outstanding. Under the prescribed strategies leading to such a history, buyer types below $\hat{b}^\dagger$ have been eliminated, buyer types satisfying $p^{initial} \leq p^R(b, s)$ accept the seller's initial offer, and the remaining non-accepting types in $[\hat{b}^\dagger, \hat{b}^*)$ separate through the strictly decreasing delay schedule $\Delta(b|\hat{b}^*)$. Thus, observing no buyer response strictly before time t eliminates all remaining types above $b(t)$. Hence, the seller's posterior support is $[\hat{b}^\dagger, b(t)]$, yielding the posterior belief $F(\cdot|\hat{b}^\dagger, b(t))$.

For histories in which the seller's initial offer remains outstanding after $T^\dagger$, the prescribed strategy assigns probability zero to such histories. Hence, Bayes' rule imposes no restriction, and the degenerate belief on buyer type $\hat{b}^\dagger$ specified in Case II-(i) is consistent. All other specified beliefs either follow immediately from Bayes' rule on the equilibrium path or are assigned after off-path histories where Bayes' rule imposes no restriction.

E Calibration of the Theoretical Prediction

This section describes how we compute the theoretical prediction used in the empirical analysis. We focus on the one-sided asymmetric-information environment in Section 2, where the seller's cost s is known and the buyer's valuation is uniformly distributed on $[\underline{B}, \overline{B}]$. We assume $s \leq \underline{B}$, so all buyer types have nonnegative gross gains from trade. The formulas below are stated for $c > 0$; the equilibrium cutoffs and payoffs for $c = 0$ are obtained by taking the limit as $c \downarrow 0$.

Characterization of Equilibrium Cutoffs. For any candidate upper cutoff b^*, the buyer's equilibrium payoff in (5) can be calculated as follows:

$$v_B(b \mid b^*) = \begin{cases} b - \dfrac{b^* + s}{2} & \forall b \in [b^*, \overline{B}], \\ \dfrac{r(b-s)^2 + 4c(b-s) - 2c(b^*-s)}{2(r(b^*-s)+2c)} & \forall b \in [b^\dagger, b^*), \\ 0 & \forall b \in [\underline{B}, b^\dagger). \end{cases} \tag{16}$$

The lower cutoff $b^\dagger$ is the lowest buyer type willing to enter the separating-delay region. For an interior lower cutoff, which is the case in our experimental calibration, it solves the quadratic equation

$$v_B(b^\dagger | b^*) = 0 \quad \Longleftrightarrow \quad r(b^\dagger - s)^2 + 4c(b^\dagger - s) - 2c(b^* - s) = 0. \tag{17}$$

Thus,

$$b^\dagger(b^*) = s + \frac{\sqrt{4c^2 + 2rc(b^*-s)} - 2c}{r}.$$

If the lowest type already obtains a nonnegative payoff from entering the separating-delay region, then $b^\dagger = \underline{B}$. Hence, in general,

$$b^\dagger(b^*) = \max\left\{ \underline{B},\, s + \frac{\sqrt{4c^2 + 2rc(b^*-s)} - 2c}{r} \right\}.$$

Given b^* and $b^\dagger(b^*)$, the seller's expected payoff under the uniform distribution is

$$v_S(b^* \mid b^\dagger(b^*)) = \frac{b^* - s}{2} \cdot \frac{\overline{B} - b^*}{\overline{B} - \underline{B}} + \frac{1}{\overline{B} - \underline{B}} \int_{b^\dagger(b^*)}^{b^*} \frac{r(b-s)^2 + 4c(b-s) - 2c(b^*-s)}{2(r(b^*-s)+2c)}\, db.$$

The equilibrium upper cutoff b^* maximizes $v_S(\tilde{b}^* | b^\dagger(\tilde{b}^*))$. For an interior solution, which is the case in our experimental calibration, the first-order condition can be written as

$$\overline{B} - b^* = \int_{b^\dagger(b^*)}^{b^*} \frac{(r(b-s)+2c)^2}{(r(b^*-s)+2c)^2}\, db = \frac{(r(b^*-s)+2c)^3 - (r(b^\dagger(b^*)-s)+2c)^3}{3r(r(b^*-s)+2c)^2}. \tag{18}$$

Together with the definition of $b^\dagger(b^*)$ above, condition (18) characterizes the equilibrium cutoffs b^* and $b^\dagger$.

Laboratory Calibration. In the main one-sided treatment, $s = 0$, $\underline{B} = 0$, and $\overline{B} = 30$. The experiment implements random termination at the end of each second with survival probability 0.99. We normalize one unit of time in the model to correspond to one second in the experiment. The exact continuous-time discount rate therefore satisfies

$$e^{-r} = 0.99 \quad \Longleftrightarrow \quad r = -\log(0.99) \simeq 0.01005.$$

In the calibration used in the main text, we use the approximation $r = 0.01$. The transaction cost is 5 cents per second, so we set $c = 0.05$.

Substituting $s = 0$, $\underline{B} = 0$, $\overline{B} = 30$, $r = 0.01$, and $c = 0.05$ into (17) and (18) gives $b^\dagger \simeq 7.72$ and $b^* \simeq 21.40$. The seller's initial offer $p^{initial}$, the equilibrium delay $\Delta(b|b^*)$, and the buyer's equilibrium payoff $v_B(b|b^*)$ are calibrated as follows:

$$p^{initial} = p^R(b^*, 0) = \frac{b^*}{2} \simeq 10.70,$$

$$\Delta(b \mid b^*) = \frac{1}{r}\log\left(\frac{r \cdot 21.40 + 2c}{rb + 2c}\right) = 100 \cdot \log\left(\frac{31.40}{b + 10}\right) \quad \forall b \in [7.72, 21.40),$$

$$v_B(b \mid b^*) = \begin{cases} b - 10.70 & \forall b \in [21.40, 30], \\ \dfrac{0.01b^2 + 0.2b - 2.14}{0.628} & \forall b \in [7.72, 21.40), \\ 0 & \forall b \in [0, 7.72]. \end{cases}$$

The calibrated seller payoff conditional on buyer type b, which we will denote by $\pi_S(b|b^*)$, is

$$\pi_S(b|b^*) = \begin{cases} 10.70 & \forall b \in [21.40, 30], \\ \dfrac{0.01b^2 + 0.2b - 2.14}{0.628} & \forall b \in [7.72, 21.40), \\ 0 & \forall b \in [0, 7.72]. \end{cases}$$

Note that $\pi_S(b \mid b^*) = v_B(b \mid b^*)$ on the separating-delay region $[b^\dagger, b^*)$ because trade occurs at the complete-information price $p^R(b,s)$, and both parties incur the same transaction cost over the same delay. The seller's ex-ante expected payoff under the calibrated equilibrium is

$$\begin{aligned} v_S(b^* \mid b^\dagger(b^*)) = \mathbb{E}[\pi_S(b|b^*)] &= \frac{1}{\overline{B} - \underline{B}}\left[\int_{b^\dagger(b^*)}^{b^*} \pi_S(b|b^*)\, db + (\overline{B} - b^*) p^{initial}\right] \\ &= \frac{1}{30}\left[\int_{7.72}^{21.40} \frac{0.01b^2 + 0.2b - 2.14}{0.628}\, db + (30 - 21.40) \cdot 10.70\right] \\ &\simeq 5.28. \end{aligned}$$