Positive and Negative Selection in Bargaining*

Dongkyu Chang† Duk Gyoo Kim‡ Wooyoung Lim§

May 11, 2026

Abstract

In the standard dynamic screening problem between an uninformed seller and a privately informed buyer, theory suggests that the presence (absence) of the buyer's outside option leads to a substantial surplus for the seller (buyer). This outcome arises from contrasting unraveling processes that theory predicts: negative selection occurs in the absence of an outside option, while positive selection occurs in the presence of it. We examine the validity of these contrasting unraveling processes and report laboratory data that qualitatively deviate from theoretical predictions. We found that the seller's profit ranking was reversed between the two environments. In particular, in the presence of an outside option, the buyer frequently rejected current-round offers, leading to pervasive delays; and the seller's reported beliefs about the buyer's type were qualitatively more consistent with the negative selection than with the theoretically predicted positive selection.

Keywords: Positive Selection, Outside Options, Laboratory Experiments

***JEL* classification numbers**: C78, C91, D03

*We are grateful to Jisun Baek, Simon Board, Emiliano Catonini, Ilwoo Hwang, Kyungmin (Teddy) Kim, and Youngwoo Koh for valuable comments and suggestions. We thank participants at the 2021 KAFE-SKKU International Conference, 2022 Asian Meeting of Econometric Society, 2022 SKKU International Conference on Digital Economy and Trade, 2022 KEA International Conference, 2022 KBER Summer Workshop on Behavioral and Experimental Economics, and 2022 Korea Academic Society of Industrial Organization Conference for useful feedback and seminar participants at NYU Shanghai and Korea University for their helpful discussions and comments. Shiyu (Jake) Zhang provided excellent research assistance. This study is supported by a grant from the Research Grants Council of Hong Kong (Grant No. GRF16506019).

†Department of Economics and Finance, City University of Hong Kong. *Email:* `donchang@cityu.edu.hk`

‡School of Economics, Yonsei University. *Email:* `kim.dukgyoo@yonsei.ac.kr`

§Department of Economics, HKUST. *Email:* `wooyoung@ust.hk`

1 Introduction

The Coase conjecture, one of the most fundamental ideas in bargaining theory, durable-good monopoly, and dynamic screening problem, proposes that the uninformed seller eventually benefits not at all from inter-temporal price discrimination among different buyer types. Consider a bargaining game between a buyer whose willingness to pay (type) is privately informed and an uninformed seller who only knows the prior distribution of buyer types. The remaining buyers at any offered price are more likely to be of low types, which leads to a *negative selection* in the demand pool. In the absence of commitment power, the seller then responds to cut the offering price over time. Anticipating such a price cut, even a high-type buyer tends to delay her purchase, which pushes the seller to lower the price even in the early stages to induce any purchase. As a result, the seller will charge, in effect, the lowest price all the time and earn the lowest possible expected profit in equilibrium. The idea of negative selection has been theoretically examined and confirmed by, among others, Fudenberg et al. (1985) and Gul et al. (1986).[1]

Board and Pycia (2014) (henceforth, BP) extend the bargaining game by considering an outside option available to buyers. One can think of the outside option as an alternative product from a third party. Low-type buyers tend to exercise the outside option and exit the market more quickly rather than haggling with the seller, so the remaining demand pool consists of high-type buyers. BP show that introducing the buyer's outside option overturns the Coase conjecture: The *positive selection* in the remaining demand pool drives the seller to charge a constant price equal to the monopoly price against the prior distribution of buyer's types. As a result, the seller earns a profit substantially higher than what the Coase conjecture predicts, without any delay. This result is surprisingly robust in the sense that the buyer's value of the outside option being positive (however close to zero) suffices to result in the (qualitatively) same bargaining outcome.[2]

Positive selection provides a channel through which the monopoly seller in the market overcomes a lack of commitment power and turns a substantial portion of consumer surplus into part of its profit. Thus, BP's result has a significant implication for the regulation and design of various markets with asymmetric information, including durable-good monopoly and sequential auctions, in which dynamic screening is at the core. If the market designer's goal is to protect consumer surplus, then BP's result, together with the Coase conjecture, suggests that it is sufficient for the

[1] The Coase conjecture was initially discussed in the context of durable-good monopoly (Coase, 1972). However, a durable-good monopoly is mathematically equivalent to the bargaining game between an uninformed seller and an informed buyer.

[2] The stark contrast led by the absence/presence of an outside option has a solid theoretical ground. BP show that the Coase conjecture fails in the unique equilibrium if the buyer's outside option is bounded away from zero. Catonini (2022) strengthens the BP's result by proving that the strategy profile in this equilibrium consists of the unique strongly rationalizable strategy profile. For the case without an outside option, Gul et al. (1986) prove that the Coase conjecture holds in the unique equilibrium whenever the buyer's lowest possible value is strictly positive. Cho (1994) shows that the Coase conjecture holds in all rationalizable strategy profiles with the restriction that the buyer's acceptance rule is weakly stationary.

designer to prevent buyers from accessing any outside options. However, this policy implication seems contrary to the conventional wisdom that restricting monopoly power usually makes the market more competitive and increases consumer surplus. Hence, it is crucial to obtain empirical validity of the positive and negative selection before discussing the policy implications, which justifies our approach of using controlled laboratory experiments.

The sharp contrast in quantitative (seller's profit ranking) and qualitative (contrasting directions of unraveling) theoretical predictions and their importance on the practical market design inspires our research. The primary purpose of our laboratory experiment is to investigate the impact of an outside option on the unraveling process, as well as on the distribution of surplus and overall market efficiency. More specifically, we consider the random-termination bargaining game between an uninformed seller and an informed buyer and compare two market conditions, those with and without an outside option. This comparison enables us to understand the *causal effect* of an outside option in the market as all other factors that could potentially affect the outcome such as the complexity of the game and other-regarding preferences remain constant. Since the literature has provided mixed evidence for negative selection (e.g., Güth et al., 1995; Rapoport et al., 1995; Reynolds, 2000; Srivastava, 2001; Cason and Reynolds, 2005), it is crucial to note that our main objective is neither to obtain an empirical validity of the Coase conjecture nor to test the predictions of Board and Pycia (2014) per se, but to explore the role of an outside option.

We consider three experimental treatments, one without an outside option and two with outside options of different sizes.[3] Our main interests are on (1) the bargaining lengths, (2) seller and buyer profits, (3) the frequency of rejections, and (4) the seller's beliefs about the matched buyer's valuation. In particular, (3) and (4) jointly allow us to understand if unraveling occurs in each market condition as the theory predicts. If the empirically observed differences between those with and without the outside option are consistent with the differences in theoretical predictions, we could further leverage our findings into the design of market policies, as it implies that the mere existence/absence of the outside option can determine who will take the lion's share of the gains from trade. Otherwise, however, we would want the theoretical predictions and the policy implications thereof to be considered with caveats.

We found that the overall behaviors observed in the two experimental treatments with an outside option are qualitatively the same.[4] Thus, we combine our data from these two treatments (jointly called OutYes) and compare them with the data from the treatment without an outside

[3] As discussed earlier, the qualitative theoretical prediction does not depend on the size of the (positive) outside option, provided that all players are rational. However, with some players' rationality being potentially bounded, one may worry that such theoretical robustness is not valid in laboratory settings (see Section 6 for further discussions). Our experimental design with two slightly different sizes of outside options aims to address the concern, serving as a placebo check. As we shall make clear soon, we do not find any evidence that bargaining outcomes in our lab experiment differ across the two treatments with different values of outside options.

[4] What we mean by the qualitatively same results is that in both treatments frequent rejections were observed and the seller's reported beliefs about the minimum of the remaining buyer's value declined over time.

option (OutNo). We have four main observations. First, buyers with low values in OutYes opted for the outside option and exited the market immediately, indicating an initial layer of positive selection. Second, the average number of bargaining rounds was significantly higher in OutNo than in OutYes, while a substantial degree of delay was found in OutYes. It was evident that some fraction (> 15%) of the buyers remained by rejecting the offer, contributing this delay in OutYes. This observation of a bargaining delay is inconsistent with the theoretical predictions. Third, we found that the seller's initial price offer was, on average, significantly higher in OutNo than OutYes, which also contradicts theoretical predictions. Furthermore, the average seller's initial price offers in both environments were significantly higher than the respective equilibrium price offer. Fourth, in both OutNo and OutYes, the seller's price cuts after the first round were much larger than the theoretical bounds. The pervasive rejections made by buyers in OutYes could be seen as optimal given the actual price cuts observed. Thus, it is evident that the initial layer of positive selection observed does not extend to subsequent layers of positive selection in our laboratory experiment.[5]

Given that one of our main interests is to understand whether positive selection occurs in the presence of an outside option, we elicited sellers' beliefs about whether lower-type buyers exit the market earlier or not. Precisely, we asked our seller participants to report their beliefs about the minimum buyer type after each round when the price offer was rejected. Our data reveal that the average minimum belief that the seller participants reported in OutYes was marginally larger than that in OutNo. However, the individual-level reports on the minimum buyer type in OutYes and OutNo were, by and large, the same. Moreover, we found that a substantial fraction of the reported values were even below the maximum buyer type whose dominant strategy is to take the outside option immediately and exit the market. That is, these seller participants appeared to be unsure about whether the lower-type buyers would leave the market earlier or not, which leads them to undercut the price too aggressively in the subsequent rounds, thereby breaking the theory-predicted linkage between the initial layer of positive selection and the subsequent layers of positive selection in OutYes. Anticipating the price cut, buyers optimally reject the first round offer, which explains the frequent rejections observed in OutYes.

The standard theory predicts that the absence of an outside option results in a negative selection in the demand pool, yielding the smallest profit for the seller, while the presence of an outside option leads to a positive selection with a substantially larger profit for the seller. In our experiment, the seller's average profit turned out to be substantially and significantly higher in OutNo than in OutYes, which sharply contradicts the main prediction from the positive selection. Furthermore, we observed pervasive rejections in OutYes, the observation that is never predicted by the positive selection in the demand pool, nor by the inequity aversion (Fehr and Schmidt, 1999).[6]

[5] In our companion paper (Chang, Kim and Lim, 2024), we experimentally examine how positive selection fails in the lab. The paper presents an experimental design that enables us to observe individual heterogeneity in the levels of positive selection reasoning. See Section 1.1 for a detailed discussion.

[6] A formal model with an inequity averse buyer and its predictions are presented in Online Appendix F.

The observed seller's profit ranking reversal between OutNo and OutYes is not only due to the failure of positive selection in OutYes but also due to the incomplete execution of negative selection in OutNo. On one hand, our data from OutNo exhibit several features qualitatively consistent with negative selection; first, the average price offers declined over rounds within a match in OutNo; second, the maximum belief that the sellers reported after the price offer was rejected in OutNo was steadily decreasing, implying that the posterior belief of the seller was a right-truncation of the prior in each round. On the other hand, both the average initial price offer and the reported maximum belief in OutNo were substantially higher than what negative selection predicts. In aggregate, we observed that the average seller profit in OutYes was significantly lower than the theoretical prediction from positive selection while that in OutNo was not statistically different from the theoretical prediction from negative selection. The empirical seller's profit in OutNo being not statistically different from the theoretical level was not a consequence of the seller fully exercising the inter-temporal price discrimination but more of a combination between higher price offers and higher rejection rates than the theoretical predictions.

One may claim that another proper way of experimentally examining theoretical predictions with an outside option is to compare the findings with those in another experimental environment where the seller makes a take-it-or-leave-it (TIOLI) offer and the buyer decides whether to accept the offer, exercise the outside option, or reject the offer which leaves nothing for both. Because the theoretical predictions under this TIOLI bargaining game with an outside option coincide with those with the BP model, comparing them enables us to examine whether the discrepancies between the theoretical predictions and experimental findings are the failure of positive selection (that is inevitably associated with existence of next bargaining rounds) or those are due to the experiment participants' limited understanding about the optimal pricing strategies. Regarding this concern, we conducted another experiment, the ultimatum bargaining with an outside option.[7] We found that the average TIOLI offer is statistically indifferent from the commitment price. This supplementary finding suggests that the sellers on average know the profit-maximizing offer in the TIOLI setting, so the reason why our experimental findings are inconsistent with BP's predictions is more likely due to the dynamics of the bargaining setup, not the sellers' limited understanding of the pricing strategy. Also, most buyers accepted the price offer whenever accepting it yields a larger monetary payoff, which leaves fairness concern out from the potential explanations for the observations.

The main driving force behind the positive selection is that the market *unravels* with the low-type buyers leaving earlier (by taking the outside option) than the high-type ones.[8] However,

[7]This treatment was not part of the original experimental design, but it was prompted at the recommendation of a reviewer. More detailed explanation will follow in Section 4.

[8]More precisely, unraveling in BP begins with the observation that the lowest buyer type (with the lowest gains from trade) never receives an offer strictly more favorable than the outside option *in any equilibrium*; hence, the lowest type takes the outside option immediately. Given this, the same argument applies to the next lowest type, whose gain from trade is the smallest among the remaining types.

unraveling may not take place perfectly if (a small fraction of) players 1) lack first-order rationality such that some low-type buyers do not leave the market early, 2) lack higher-order rationality such that the seller is unsure about whether the lower-type buyers leave the market early, or both. Any of these scenarios leads the buyer to believe (either wrongly or correctly) that the seller's price in the subsequent rounds may be lower than that in the current round. We explore this idea formally and present an alternative model in Online Appendix E in which the buyer lacks first-order rationality and may choose to delay sub-optimally only if the payoff-consequence of the sub-optimal decision is not larger than $\epsilon > 0$.[9] Another parsimonious way to rationalize our data is to consider the bargaining game with a small fraction of buyers who *optimistically* believe that the seller would occasionally offer a low price in a subsequent round.[10] Introducing an obstinate buyer type generates equilibria with an inefficient delay where rational buyers mimic obstinate types to increase their payoff (Myerson, 1991; Abreu and Gul, 2000). Insufficient skepticism, frequently observed in the information disclosure experiments (Jin et al., 2021) and in the field setup (Brown et al., 2012), will also lead to a failure of unraveling in our environment.

The rest of this paper is organized as follows. In the following subsection, we discuss the closely related literature. Section 2 describes the theoretical environment and characterizes the equilibrium outcomes. Section 3 presents an example in which buyers' types are drawn from a uniform distribution, which will serve as the foundation for our experiment. Section 4 describes the experimental design and procedure. The results are reported in Section 5. Section 6 discusses various mechanisms that can explain our experimental findings. Section 7 concludes.

1.1 Literature Review

The Coase conjecture is originally proposed by Coase (1972), and more formal theoretical treatments of the conjecture are provided by Fudenberg et al. (1985), Gul et al. (1986), Ausubel and Deneckere (1989), among others. The idea of negative selection has played an important role from the very beginning of the development of this literature; see, for example, Fudenberg et al. (1985, Lemma 1) and Ausubel et al. (2002, Lemma 1). The idea of negative selection is also the basis of not only the Coase conjecture but also dynamic screening problems in other contexts, including dynamic lemon markets (Evans, 1989; Vincent, 1989; Deneckere and Liang, 2006) and sequential auctions (McAfee and Vincent, 1997; Liu et al., 2019).

[9] We show that the ϵ-irrationality results in the outcome qualitatively identical to that from our main model with buyer's optimism even when ϵ tends to 0.

[10] In Online Appendix D, we show that even when the fraction of the optimistic buyers is arbitrarily close to zero, the unraveling fails and the positive selection breaks down in an epsilon-Perfect Bayesian equilibrium that allows for a small mistake in sellers' best responses. The equilibrium predicts that the negotiation takes multiple periods with positive probability, and the price declines over time. It exhibits some features reflecting both positive and negative selection: Any rational low-type buyers exercise the outside option immediately, and any rational higher-type buyers trade with the seller, possibly after a delay. Among these buyer types, one with a higher valuation trades earlier than others.

Recently, positive selection has attracted serious attention from the literature. Board and Pycia (2014) show that the introduction of the buyer's outside option, even if the value of the outside option is arbitrarily close to zero, overturns the Coase conjecture in the sense that a seller can earn profit substantially higher than what the Coase conjecture predicts. In a similar vein, Tirole (2016) shows that a principal can implement the outcome of a profit-maximizing mechanism even without commitment power for a large class of dynamic screening problems, while the Coase conjecture implies the least profit to the principal. These results demonstrate the contrasting effect of positive selection from negative selection. While negative selection is generally harmful to the interest of the principal, positive selection leads to the best outcome for the principal.

Empirical and experimental evidence for the Coasean dynamics is, by and large, mixed. A vast majority of the previous experiments have reported evidence that contradicts the Coasean dynamics in various environments (e.g., Güth et al., 1995; Rapoport et al., 1995; Reynolds, 2000; Srivastava, 2001; Cason and Reynolds, 2005). They commonly found that initial prices are increasing in the discount factor and substantially above the static monopoly price level. Cason and Sharma (2001) and Güth et al. (2004) are two exceptions that obtain partial support for the Coase conjecture. More recently, Fanning and Kloosterman (2022) propose an experimental design that relies on subjects' private information about preferences for fairness to test the Coase conjecture. Considering two settings, an infinite horizon bargaining game and an ultimatum game, they find that, consistent with the Coase conjecture, initial offers, minimum acceptable offers, responder payoffs, and efficiency are significantly larger in the infinite horizon environment.

To the best of our knowledge, no empirical or experimental evidence has been provided for positive selection. Our study, along with our companion paper (Chang, Kim and Lim, 2024), aims to address this gap by conducting a systematic experimental investigation into positive selection. In Chang et al. (2024), we conducted a simple two-round bargaining experiment with finite price alternatives to examine the sellers' belief updates and their adherence to different levels of positive selection reasoning. The main objective was to examine whether participants could engage in different levels of positive selection reasoning, rather than directly comparing bargaining outcomes with and without outside options. Consistent with our findings in this paper, Chang et al. (2024) demonstrated that a significant proportion of price offers in the initial round were rejected, leading to delays in the negotiation process. Furthermore, only a small fraction of sellers have their posterior beliefs in accordance with positive selection reasoning. These experimental results directly challenge the theoretical predictions and reveal a failure of positive selection, casting doubt on the notion that the introduction of an outside option leads to substantial profits for the seller. Our study complements Chang et al. (2024) by providing a direct comparison between situations with and without outside options, further enhancing our understanding of the impact of outside options in bargaining.

2 Theoretical Background

2.1 Model

Environment Consider a price negotiation between a seller (she) and a buyer (he) over infinite-horizon discrete time (period) $n = 0, 1, 2, \ldots$. The seller holds an indivisible good for sale, whose value to herself is normalized to zero. The buyer's value of the good (i.e., the buyer's type) $v \in V = \{v_1, v_2, \ldots, v_N\}$ is his private information, and the seller holds the prior belief that $v = v_j$ with probability $f(v_j) > 0$. Let $F(v) := \sum_{v' \le v} f(v')$ denote the cumulative distribution function. We order the buyer types in V by $0 < v_1 < v_2 < \ldots < v_N$ and denote the lowest and the highest possible buyer type as $\underline{v} \equiv v_1$ and $\overline{v} \equiv v_N$, respectively.

The buyer also has an outside option that he can exercise anytime during the negotiation. Throughout this paper, we assume that the value of the outside option is *type-independent* and worth w to all buyer types. Further, we focus on the two special cases such that $w > 0$ and $w = -\infty$.

In each period $n \ge 0$, the seller offers a price $p_n \ge 0$, and then the buyer decides to accept p_n (trade), exercise the outside option (opt-out), or delay. If the buyer accepts p_n, the negotiation ends with final payoffs $e^{-rn\Delta} p_n$ and $e^{-rn\Delta}(v - p_n)$ for the seller and the buyer respectively, where $r > 0$ and $\Delta > 0$ are the common discounting rate and the time duration between consecutive periods. If the buyer exercises the outside option in period n, the negotiation ends with final payoffs $e^{-rn\Delta} w$ for the buyer and zero for the seller. If the buyer chooses to delay, the negotiation moves on to the next period, and then the two parties repeat the same bargaining protocol. Both parties obtain zero payoffs if they fail to reach any agreement forever.

If $w = -\infty$, the buyer will never exercise the outside option as he can guarantee himself at least zero payoffs by delaying the negotiation indefinitely; hence, the assumption $w = -\infty$ is equivalent to the assumption that the buyer has no outside option at all. Define each buyer type's net-value, denoted by $u(v)$, as the difference between v and the autarky payoff that the buyer can guarantee himself regardless of which strategy the seller employs:

$$u(v) := \begin{cases} v & \text{if } w = -\infty \\ v - w & \text{if } w > 0. \end{cases} \tag{2.1}$$

We may interpret $u(v)$ as the gains from trade. Note that the buyer will never accept p_n higher than her net-value $u(v)$. We assume $u(v) \ge 0$ for all $v \in V$ without loss.[11]

[11] For any buyer type with $u(v) < 0$ (if any), it is a strictly dominant strategy to exercise the outside option immediately. Thus, the presence of such a buyer type does not make any difference in the equilibrium strategies of other players. In our experiment, we consider a treatment (Out60) in which some buyer types indeed have a negative net-value. Theory predicts that those buyer types exercise the outside option in period 0, while all other players behave as discussed in this section.

Strategies and Equilibrium The equilibrium concept is *perfect Bayesian equilibrium (PBE)*. Here, we introduce some notations useful to describe equilibrium strategies and beliefs. Suppose that the buyer has rejected the seller's offers $p_0, p_1, \ldots, p_{n-1}$ and continues the negotiation in period n. We generically denote such a history of the game by $h_n = (p_0, p_1, \ldots, p_{n-1})$, while h_0 refers to the null history. For any $n \geq 1$, let $H_n = [0, \infty)^n$ denote the set of possible histories up to the beginning of period n, and let $H = \bigcup_{n \in \mathbb{N}} H_n \cup \{h_0\}$ denote the set of all possible histories.

For any $p \geq 0$, $h_n \in H$, and $v \in V$, $\sigma^B(p|h_n, v) : \{T, O, D\} \to [0,1]$ generically denotes the behavioral strategy of a buyer type v in response to the seller's offer p at h_n, where $\sigma^B(p|h_n, v)[T]$, $\sigma^B(p|h_n, \theta)[O]$, and $\sigma^B(p|h_n, v)[D] \in [0,1]$ refer to probabilities that the buyer of type v chooses T (accept to trade at p), O (exercise the outside option), and D (delay), respectively. $\sigma^S(h_n) \in \triangle(\mathbb{R}_+)$ denotes the seller's behavioral strategy at h_n. $f^S(v|h_n) \in [0,1]$ denotes the seller's posterior belief that the buyer's type is v, and $F^S(v|h_n) = \sum_{v' \leq v} f^S(v'|h_n)$ denotes the corresponding distribution function.

A PBE assessment is generically denoted by $\sigma = (\sigma^B, \sigma^S, f^S)$. Let $\mathcal{E}_w(\Delta)$ denote the set of all PBEs.[12] We also follow the convention that $\text{supp}(\sigma^S(h_n))$ and $\text{supp}(f^S(h_n))$ denote the supports of $\sigma^S(h_n)$ and $f^S(h_n)$. Finally, $\overline{v}^\sigma(h_n) := \max\{v \in V : f^S(v|h_n) > 0\}$ and $\underline{v}^\sigma(h_n) := \min\{v \in V : f^S(v|h_n) > 0\}$ denote the highest and the lowest buyer types in $\text{supp}(f^S(h_n))$.

Full Commitment Benchmark For a benchmark, suppose that the seller can fully commit to any price path before the negotiation commences. Following the standard lines, we can show that it is optimal for the seller to commit to the following single price indefinitely:

$$p_w^* = \arg\max_{p \geq 0} \underbrace{\left\{ p \sum_{v: u(v) \geq p} f(v) \right\}}_{:= \Pi(p)}. \tag{2.2}$$

Faced with this price path, the buyer accepts p_w^* in period zero if and only if $u(v) \geq p_w^*$. The buyer types with $u(v) < p_w^*$ are excluded by the seller despite positive gains from trade. If $w > 0$, these excluded buyer types exercise the outside option in period zero; if $w = -\infty$, they continue the negotiation indefinitely by rejecting $p_{-\infty}^*$ in all periods. In either case, the seller exercises no inter-temporal price discrimination. We call this outcome *the full-commitment benchmark*. In this section, we will assume that (2.2) admits a unique solution, and thus, p_w^* is well-defined. Let

$$\Pi_w^* := \max_{p \geq 0} \Pi(p) = p_w^* \sum_{v: u(v) \geq p_w^*} f(v)$$

[12] Our notation suppresses the dependence of $\mathcal{E}_w(\Delta)$ on other parameters (e.g., F and r) which we keep fixed throughout the analysis.

denote the seller's profit in the full-commitment benchmark outcome. Clearly, the seller cannot earn strictly more than Π_w^* in any PBE.

Proposition 1. *The following holds in the full-commitment benchmark outcome:*

(i) No Inter-temporal Pricing: The seller offers the same price p_w^* in all periods.

(ii) No Delay: All trade (if any) must occur in period zero.[13]

(iii) Exclusion: The seller does not trade with the buyer types whose $u(v)$ is smaller than p_w^*.

To exclude trivial cases, we further assume $u(\underline{v}) < p_w^*$, so the exclusion occurs with non-zero probability.[14] Without this assumption, (i) the full benchmark outcome coincides with the first best, and (ii) the seller can easily achieve this outcome by offering $p_0 = u(\underline{v})$ in period zero even without the full commitment power; hence, the bargaining problem becomes trivial with or without the seller's commitment power.[15]

2.2 Bargaining with No Outside Option

Now we turn to the case in which the seller cannot commit to future prices. First, consider the case that the buyer has no outside option (i.e., $w = -\infty$). It is intuitively appealing to conjecture that the seller *moves down the demand curve* $1 - F$ in equilibrium. Higher buyer types enjoy a higher payoff from consuming the good, hence a delay of purchase is more costly for them. As a result, high buyer types are more eager to purchase and end the negotiation earlier, inducing *negative selection* in the demand pool.

This intuition is indeed correct. Fix an arbitrary PBE σ, and suppose that the seller offers p_n in period n after a history h. Also, suppose that a buyer type v is willing (at least indifferent) to accept this offer, given his expectation about future offers. One can prove that any buyer type $v' > v$ also *strictly* prefers to accept p_n; hence, all buyer types higher than v must trade. This observation, known as the skimming property in the literature, implies that (i) $\overline{v}^\sigma(h)$ monotonically decreases over time, and (ii) the seller's posterior f^S is always a right-truncation of the prior belief f. Its proof can be found in Fudenberg et al. (1985).

Proposition 2 (Negative Selection)**.** *Suppose $w = -\infty$. The following holds for $h \in H$, $p \geq 0$, and $\sigma = (\sigma^B, \sigma^S, f^S) \in \mathcal{E}_w(\Delta)$.*

[13] For Proposition 1-(i) and (ii) to hold, it is crucial to assume that (a) all buyer types have the same type-independent value for the outside option and (b) the seller and the buyer have identical discount rates. If either assumption fails, the seller may earn a higher profit by committing to a declining price path, thereby inducing a delay in trade. For example, see Landsberger and Meilijson (1985) and Chang (2021) for cases with asymmetric discount rates and type-dependent outside options, respectively.

[14] This assumption holds if, for example, f is the uniform distribution such that $u(\underline{v}) < 2u(\overline{v})$.

[15] All buyer types will accept $p_n = u(\underline{v})$ after any history in any PBE; see, for example Fudenberg et al. (1985, Lemma 2) and Board and Pycia (2014, Lemma 1). Hence, the seller's expected profit is necessarily no less than $u(\underline{v})$ after any history in any PBE.

(i) Suppose that the seller offers p after h, and a buyer type $v \in V$ finds it weakly optimal to accept p. Then, all higher buyer types $\{v' \in V : v' > v\}$ find it strictly optimal to accept p.

(ii) $f^S(h,p)$ is always a right-truncation of $f^S(h)$: There is $M > 0$ such that $f^S(v|h,p) = f^S(v|h)/M$ for any $v < \overline{v}^\sigma(h)$.[16]

Negative selection in the demand pool keeps pushing $\overline{v}^\sigma(h)$ downward. The seller's posterior belief becomes more pessimistic (about the gains from trade) and hence lowers her offer period after period. In other words, the seller exercises inter-temporal price discrimination in equilibrium. Under the assumption that $u(\underline{v}) = \underline{v} < p^*_{-\infty}$, the negotiation takes multiple periods with positive probability. The proof of the following proposition can be found in, for example, Gul et al. (1986).

Proposition 3. *Suppose $w = -\infty$ and $u(\underline{v}) = \underline{v} < p^*_{-\infty}$.*

(i) Delay: In any PBE, the negotiation takes over multiple periods with positive probability.[17]

(ii) Inter-temporal Pricing and Negative-selection: The seller's offer strictly declines toward $\underline{v}$ over time on the path of any PBE. All buyer types trade with the seller eventually, where higher buyer types trade earlier at a higher price than low buyer types.

(iii) Coase Conjecture: There is an integer $L \in \mathbb{N}_0$ such that, in any $\sigma \in \mathcal{E}_{-\infty}(\Delta)$, all buyer types trade with the seller in period L or before, and the seller's offer p_n on the equilibrium path is bounded by

$$\underline{v} \le p_n \le (1 - e^{-rL\Delta})\overline{v} + e^{-rL\Delta}\underline{v} \quad \forall n \le L.$$

Moreover, L is independent of Δ.

The Coase conjecture implies that the bargaining process ends in the "twinkling of an eye" as the duration between consecutive periods, Δ, approaches zero. Consequently, the seller's benefit from inter-temporal price discrimination is completely undermined. Indeed, the upper bound for p_n in Proposition 3-(iii) converges to $\underline{v}$ as $\Delta \to 0$, hence, the seller's equilibrium profit as well as equilibrium offers converge to $\underline{v}$ in the limit. Furthermore, the duration of the bargaining also converges to zero (though the number of periods may be greater than 1), and thus the first-best outcome is achieved in the limit.

[16] More precisely, the statement holds with M being equal to the probability that the buyer rejects p at h. That is, $M = \sum_{v' < \overline{v}^\sigma(h,p)} f^S(v'|h) + f^S(\overline{v}^\sigma(h,p)|h)\sigma^B(p|h, \overline{v}^\sigma(h,p))[D]$.

[17] Even when $p^*_{-\infty}$ is only marginally larger than $\underline{v}$, the seller can secure a profit strictly higher than $\underline{v}$ by offering $\tilde{p}_0 = (1 - e^{-r\Delta})p^*_{-\infty} + e^{-r\Delta}\underline{v} > \underline{v}$ in period 0. Following standard arguments, one can show that the seller never offers a price strictly lower than $u(\underline{v}) = \underline{v}$ in any equilibrium; see, for example, Fudenberg et al. (1985, Lemma 2) and Board and Pycia (2014, Lemma 1). Thus, once the seller offers $\tilde{p}_0$, whether on or off the equilibrium path, the buyer accepts this price if $v - \tilde{p}_0 \ge e^{-r\Delta}(v - \underline{v})$, or equivalently, $v \ge p^*_{-\infty}$. Some buyer types (e.g., the buyer type $\underline{v}$) would reject $\tilde{p}_0$ and delay the negotiation. However, in the next period, the seller can trade with all remaining buyer types by offering $p_1 = \underline{v}$. Hence, by offering $\tilde{p}_0$, the seller can guarantee a profit of $[(1-e^{-r\Delta})p^*_{-\infty} + e^{-r\Delta}\underline{v}]\sum_{v \ge p^*_{-\infty}} f(v) + \delta\, \underline{v} \sum_{v < p^*_{-\infty}} f(v)$. This profit exceeds $\underline{v}$, and thus a delay must occur in equilibrium whenever $p^*_{-\infty} > \underline{v}$.

The bargaining game generically admits a unique PBE in the following sense. Fix an arbitrary probability mass function f over V. For any $-1 < \epsilon < 1$, let f_ϵ denote the perturbation of f such that

$$f_\epsilon(\overline{v}) = (1-\epsilon)f(\overline{v}) + \epsilon \quad \text{and} \quad f_\epsilon(v) = (1-\epsilon)f(v) \quad \forall v \neq \overline{v}.$$

Then, there is $\bar{\epsilon} > 0$ such that all PBEs in $\mathcal{E}_{-\infty}(\Delta)$ induce the identical outcome whenever the seller's prior belief belongs to $\{f_\epsilon : -\bar{\epsilon} < \epsilon < \bar{\epsilon}\}\backslash\{f\}$.[18] Moreover, even when the bargaining game admits multiple PBEs, all buyer types play the same strategies on the path, and the seller's expected equilibrium payoff is identical across all PBEs (Fudenberg et al., 1985; Gul et al., 1986).

2.3 Bargaining with Outside Option

Now suppose that $w > 0$. Proposition 2 does not hold in this case; in particular, $\underline{v}^\sigma(h)$ is not necessarily $\underline{v}$. Lower buyer types have smaller gains from trade, and hence they tend to exit the market earlier than higher buyer types by exercising the outside option, especially when the seller is expected to insist on high prices during the negotiation. The possibility of such *positive selection* in the demand pool changes the PBE outcome dramatically. BP prove that the seller can implement the full-commitment benchmark outcome in the essentially unique PBE.[19]

Proposition 4 (Board and Pycia 2014). *Suppose $w > 0$. The full-commitment benchmark outcome is induced in the essentially unique equilibrium.*

We can understand this result with the language of positive and negative selection. In the essentially unique PBE, the possibility of positive selection allows the seller to resist any temptation of a price cut, holding a sufficiently optimistic belief about the remaining gains from trade. More formally, note that the lowest possible gain from trade is $\underline{v} - w$ among the buyer types in V. As a result, the seller will never offer strictly lower than $\underline{v} - w$ after any history in any PBE (Board and Pycia, 2014, Lemma 1).

This observation implies that the lowest buyer type's payoff from delaying the negotiation is at most

$$e^{-r\Delta}\max\{\underline{v} - (\underline{v} - w), w\} = e^{-r\Delta}w,$$

where $\underline{v} - (\underline{v} - w)$ captures the case that the seller offers $\underline{v} - w$ in the next period, and w captures the case that the buyer exercises the outside option in the next period. Since $e^{-r\Delta}w < w$, it is always better for the lowest type to exercise the outside option in period zero than any delay. Consequently, positive selection would occur immediately, at least for the low type, whenever the

[18]Compared to the original prior f, f_ϵ with a negative ϵ assigns less probability to $\overline{v}$, while f_ϵ with a positive ϵ assigns more probability to $\overline{v}$. $f_\epsilon(v) > 0$ for all $v \in V$ and $\epsilon \in (-\bar{\epsilon}, \bar{\epsilon})$, provided that $\bar{\epsilon} > 0$ is sufficiently small.

[19]A PBE is said to be essentially unique if all PBEs lead to the same payoffs. Board and Pycia (2014) consider a more general case with type-dependent outside options and show that the full benchmark outcome is implemented in the essentially unique PBE.

seller attempts any price cut. This is not the end of the story. Suppose the buyer types $[\underline{v}, \underline{v} + \xi]$ have exited in the first period by the argument above, where $\xi > 0$ is a small positive number. In the next period, $v = \underline{v} + \xi$ becomes the lowest buyer type in the seller's posterior belief. The same argument shows that the seller never offers lower than $p_n = (\underline{v} + \xi) - w$ in the continuation game (i.e., in the next period and beyond), which means that the buyer type $v = \underline{v} + \xi$ can never hope for a payoff larger than w from continuing the negotiation; this buyer type will also find it optimal to exercise the outside option immediately. We can repeat the same argument to obtain the following conclusion: For any $\tilde{v} \in V$, if the seller trades with the buyer types in $[\tilde{v}, \overline{v}]$ in period zero, all buyer types in $[\underline{v}, \tilde{v})$ will find it optimal to exercise the outside option immediately. Hence, neither price cut nor delay occurs in period zero. Given this, it is optimal for the seller to offer $p_0 = p_w^*$ to induce the full-commitment benchmark outcome.

Proposition 4, especially its prediction of no delay, is robust in several ways. Note that the above argument—buyers either accept the offer or exercise the outside option immediately—applies to the equilibrium play in any period (not only period zero) after any negotiation history (both on and off the equilibrium paths). Therefore, all buyer types never choose to delay in response to *any price offer* by the seller *after any history*; see the proof of Proposition 1 in BP. An immediate implication is that the negotiation always ends in period zero with probability 1, even in the case that the seller trembles hands and mistakenly chooses a non-equilibrium offer.

Proposition 4 also remains to hold under a more flexible solution concept. Catonini (2022) shows that the full-commitment benchmark outcome is induced in every *strongly rationalizable* strategy profile. In other words, the assumptions of rationality and common strong belief in rationality suffice to predict that neither price cut nor delay ever occurs in negotiation with the buyer's outside option $w > 0$. It implies that we can dispense with the assumption that all the players (i.e., both the buyer and the seller) hold rational expectations about the opponents' strategies.

We close this section with a discussion on how the bargaining outcome changes as we gradually change w. First, note that a small change in $w > 0$ does not result in a qualitative change in the essentially unique equilibrium outcome as long as w remains strictly positive (however small it is); the only change is that the seller's offer p_w^* continuously varies with w. Next, for any negative but finite $w \in (-\infty, 0)$,[20] the buyer will never exercise the outside option as he can obtain a strictly higher payoff from delaying the negotiation forever. Hence, the case with a negative outside option is effectively equivalent to the case with no outside option (however small its magnitude is); the generically unique equilibrium outcome does not respond to any change in $w \in (-\infty, 0)$. If w is exactly zero, there are multiple PBEs including the two PBEs described in Propositions 3 and 4. In this sense, there is a discontinuity regarding the uniqueness of equilibrium but a continuity regarding the existence of each equilibrium at $w = 0$ from both directions.

[20]The payoff from exercising the outside option may be negative if the significant search and/or transportation costs to identify and pursue it outweigh its value.

The results summarized in this section show that the absence/presence of an outside option changes the theoretical predictions dramatically. The results also have an important policy implication, especially for the market design and regulatory policy in various markets. If the market designer's goal is to protect consumer surplus, then BP's result, together with the Coase conjecture, suggests that it is sufficient for the designer to prevent buyers from accessing any outside options. However, this policy implication seems contrary to the conventional wisdom that access to outside options usually enhances consumer surplus. Such a gap between our conventional wisdom and economic theory motivates our approach of using controlled laboratory experiments.

3 A Uniform Example

In this section, we illustrate the theoretical predictions with an example. In Section 3.1, we describe the equilibrium outcome for the case in which the buyer's value is drawn from a uniform distribution. This example coincides with the experimental design discussed later; hence, it also provides concrete hypotheses for our experiment. In Section 3.2, we demonstrate with the same example how forward induction can restrict the seller's off-the-path beliefs.

3.1 Equilibrium

Suppose that the buyer's type v is randomly drawn from a discrete uniform distribution over $V = \{v_1, v_2, \ldots, v_{N+1}\}$, where the lowest and the highest values are $\underline{v} = v_1 = 50$ and $\overline{v} = v_{N+1} = 400$. In the numerical exercise below, we adopt $N = 1400$ as the number of equally spaced grid points.[21] We compare two cases: one in which the buyer has an outside option worth 50 ($w = 50$) and one in which the buyer has no outside option ($w = -\infty$). The full-commitment profit levels are given by $\Pi^*_{50} \approx 87.50$ and $\Pi^*_{-\infty} \approx 114.28$ for the respective cases. The gain from trade, as measured by $u(v)$, is larger when the buyer has no outside option, and thus $\Pi^*_{-\infty} > \Pi^*_{50}$.

For the case $w = 50$, the seller could achieve the full-commitment benchmark profit level Π^*_{50} even without full commitment power (Proposition 4). In both cases, with and without full commitment power, the seller insists on the same price $p^*_{50} = 175$ in all periods. All buyer types weakly higher than $p^*_{50} + w = 225$ (i.e., $u(v) \geq p^*_{50}$) accept the seller's offer p^*_{50} in period $n = 0$, while all remaining buyer types immediately exercise the outside option. The seller's equilibrium offer and profit are independent of the discount factor, and, notably, no delay occurs in equilibrium.

For the case $w = -\infty$, the seller offers $p^*_{-\infty} = 200$ in the full commitment benchmark case, and the buyer immediately accepts this offer if and only if $v \geq p^*_{-\infty}$. We can characterize the equilibrium without the seller's full commitment power using the dynamic programming method (Fudenberg et al., 1985; Gul et al., 1986; Ausubel et al., 2002). A complete description of the equilibrium is

[21] All theoretical and numerical results remain qualitatively similar across a wide range of grid points, from smaller (e.g., $N = 350$) to larger (e.g., $N = 4200$) values.

convoluted, so we focus on the seller's equilibrium offers and expected profit. Table 1 illustrates the equilibrium price path for four different discount factors $\delta \equiv e^{-r\Delta} \in \{0.65, 0.80, 0.95\}$,[22] where $p_n^{\text{OutNo}}(\delta)$ represents the seller's equilibrium offer in period n. The price path is qualitatively similar across all discount factors considered: the seller initially offers a relatively high price in period $n = 0$, then gradually lowers it whenever the offer is rejected. Finally, the seller offers a price equal to $\underline{v}$, which all buyer types accept for sure.

Table 1: Theoretical Predictions When $v \sim U[50, 400]$ ($w = -\infty$)

	FC	$\delta =$		
		0.65	0.80	0.95
$p_0^{\text{OutNo}}(\delta)$	200.00	169.42	121.56	80.83
$p_1^{\text{OutNo}}(\delta)$	–	110.96	86.38	68.99
$p_2^{\text{OutNo}}(\delta)$	–	70.42	62.05	59.53
$p_3^{\text{OutNo}}(\delta)$	–	51.13	50.75	53.11
$p_4^{\text{OutNo}}(\delta)$	–	50.00	50.00	50.21
$p_5^{\text{OutNo}}(\delta)$	–	–	–	50.00
$\Pi_S^{\text{OutNo}}(\delta)$	114.28	86.97	76.21	59.07

* FC: Full Commitment Benchmark
* $\delta \equiv e^{-r\Delta}$: the common discounting factor
* $p_n^{\text{OutNo}}(\delta)$: the equilibrium offer in period n
* $\Pi_S^{\text{OutNo}}(\delta)$: the seller's equilibrium profit

Three features of $p_n^{\text{OutNo}}(\delta)$ in Table 1 are worth mentioning. First, for all n and δ considered here, $p_n^{\text{OutNo}}(\delta)$ is lower than $p_{50}^* = 175$, the offer that the seller would make if the buyer had an outside option of $w = 50$. Second, the seller's initial offer $p_0^{\text{OutNo}}(\delta)$ is rejected by some buyers, so delay occurs with positive probability when the buyer has no outside option. With $\delta = 0.80$, for example, $p_0^{\text{OutNo}}(\delta)$ is rejected by all buyer types with $v \leq 262$, which implies delay occurs in period $n = 0$ with a probability of approximately $0.60 \approx (262 - \underline{v})/(\overline{v} - \underline{v})$.[23] Third, consistent with the Coase conjecture, all equilibrium prices approach $\underline{v}$ as δ increases.

The seller's equilibrium profit levels without commitment power are indicated by $\Pi_S^{\text{OutNo}}(\delta)$ in Table 1. Recall that, with full commitment power, the seller could achieve a higher expected profit when the buyer has no outside option. However, a reversal occurs without commitment power: when the buyer has no outside option, the seller's equilibrium profit decreases as δ approaches 1.

[22]Among these discounting factors, we adopt $\delta = 0.80$ for our experimental design; see Section 4. The case with zero discounting factor is mathematically equivalent to the full benchmark case.

[23]The buyer prefers to accept $p_1^{\text{OutNo}}(\delta)$ in period $n = 1$ rather than accept $p_0^{\text{OutNo}}(\delta)$ in period $n = 0$ iff $\delta(v - p_1^{\text{OutNo}}(\delta)) \geq v - p_0^{\text{OutNo}}(\delta)$, or equivalently, $v \leq [p_0^{\text{OutNo}}(\delta) - \delta p_1^{\text{OutNo}}(\delta)]/[1 - \delta]$, where this cutoff buyer type is approximately 262 when $\delta = 0.8$.

3.2 Off-the-path Belief

When the buyer has an outside option, the negotiation ends immediately without any delay in equilibrium: for any $p_0 \geq 0$ posted by the seller, the buyer either buys or exercises their outside option in period 0 in any PBE. Thus, the seller's posterior beliefs in periods $n \geq 1$ cannot be theoretically pinned down. Indeed, there are multiple PBEs that lead to the same outcome on the equilibrium path but differ in the seller's posterior belief upon a delay. The lack of a theoretical benchmark on the seller's posterior belief imposes a challenge in experimentally testing positive selection in lab settings.

However, we can still apply the idea of forward induction to further restrict the seller's posterior beliefs. We will first discuss the general idea and then illustrate it with the uniform example. We focus on the seller's belief in period 1, though a similar idea is applicable to other periods as well. Fix any PBE $\sigma = (\sigma^B, \sigma^S, f^S)$, and suppose that the seller offers $p_0 \geq 0$ in period 0 in this PBE. The buyer's highest payoff from delaying the negotiation is $\max\{e^{-r\Delta}v, e^{-r\Delta}w\}$, where $e^{-r\Delta}v$ is the buyer's payoff from trading at $p_1 = 0$, and $e^{-r\Delta}w$ is the payoff from exercising the outside option in the next period. In other words, $\max\{e^{-r\Delta}v, e^{-r\Delta}w\}$ is the buyer's continuation payoff under the most optimistic conjecture regarding the seller's offer in the next period. On the other hand, the buyer can obtain $\max\{v - p_0, w\}$ by accepting p_0 or exercising the outside option in the current period. Define

$$D(p_0) := \left\{v \in V : \max\{e^{-r\Delta}v, e^{-r\Delta}w\} < \max\{v - p_0, w\}\right\}. \tag{3.1}$$

Any buyer type in $D(p_0)$ would find it strictly dominated to delay the negotiation to the next period in response to the seller's offer p_0. Thus, the standard dominance argument imposes that the seller's posterior belief $f^S(p_0)$ in the next period assigns zero probability to any buyer type in $D(p_0)$ whenever $D(p_0)$ is a non-empty proper subset of V.[24]

The following example, pertinent to our experimental design, illustrates how the dominance argument works: the seller's posterior belief after round 1 should assign zero probability to buyer types below a certain threshold, provided that the seller's initial offer, p_0, is greater than a certain level. Consider again the example such that v is drawn from the uniform distribution over $V = \{v_1, \ldots, v_{N+1}\}$ where the lowest and the highest values are $\underline{v} = v_1 = 50$ and $\overline{v} = v_{N+1} = 400$, and the buyer has an outside option worth $w = 50$. We will focus on the case pertinent to our experimental design, where the common discounting factor is $\delta \equiv e^{-r\Delta} = 0.8$, and $\underline{v} < w/\delta$. Suppose the seller offers $p_0 \geq 0$ in period zero. Since $\underline{v} < w/\delta$, there exists $v_0 \in V$ such that $u(v_0) = v_0 - w = p_0$. We can simplify the inequality condition in (3.1) to $\delta v < \max\{v - p_0, w\}$.[25] Thus, $D(p_0)$ is a proper subset

[24] A buyer type v belongs to $D(p_0)$ only if his payoff from delaying is less than $\max\{v - p_0, w\}$ under *all possible conjectures* that he may have regarding how the continuation game would be played. We can strengthen the effectiveness of the dominance argument, so that $f^S(p_0)$ assigns zero probability to a larger set of buyer types, by further limiting the set of conjectures that the buyer may have.

[25] Note that $\delta v \geq \delta(p_0 + w) \geq \delta w$ for buyer types $v \geq v_0$, and thus, the inequality in (3.1) is equivalent to $e^{-r\Delta}v < \max\{v - p_0, w\}$ for these buyer types. For buyer types $v < v_0$, $\delta w < w = \max\{v - p_0, w\}$, and hence, the inequality in

of V if and only if

$$\delta v_0 \geq \max\{v_0 - p_0, w\} = w \quad \Longleftrightarrow \quad p_0 \geq \frac{1-\delta}{\delta} w = 12.5 \tag{3.2}$$

When (3.2) holds, $D(p_0)$ is the union of two sets (intervals) as follows

$$D(p_0) = \{v \in V : v < w/\delta\} \cup \{v \in V : v > p_0/(1-\delta)\} \subsetneq V.$$

Thus, the dominance argument imposes the seller's posterior $f^S(v|p_0)$ to be zero for all buyer types less than $w/\delta = 62.5$, as well as for those higher than $p_0/(1-\delta) = 5 \cdot p_0$.

4 Experimental Design and Hypotheses

4.1 Experimental Design

We consider three infinite horizon bargaining games and one single-round bargaining game between a seller and a buyer with one-sided private information in our experiment. The seller has an indivisible good for sale. It is common knowledge that the seller has zero intrinsic value to the good. The buyer's value of the good v is drawn uniformly from the support $[50, 400]$ prior to the negotiation, and it is private information of the buyer.

In each round[26] $n = 1, 2, \ldots$ of the infinite horizon games, the seller offers a price $p_n \in (0, 400)$, and the buyer decides whether to accept, reject the price offer or to take the outside option. The buyer's value of the outside option $w \in \{\varnothing, 50, 60\}$ is commonly known. To implement an infinitely repeated game in the lab, we introduce the random termination of a supergame (Roth and Murnighan, 1978) with a fixed continuation probability of $\delta = 0.8$. The expected length of each supergame (called a match) is $\frac{1}{1-\delta} = 5$ rounds. When the buyer accepts p_n in round n, the negotiation ends, and the seller and the buyer receive the respective payoffs of p_n and $v - p_n$. When the buyer rejects p_n, the negotiation proceeds to the next period with probability δ. In case of termination, both the seller and the buyer receive a payoff of zero. When the buyer decides to pursue the outside option, the negotiation ends, and the seller and the buyer receive the respective payoffs of 0 and w. The experiment consists of seven matches, and participants are reshuffled to form new pairs after each match so that there are no strategic dynamics between matches. The participants' roles were fixed across matches.

Three treatment conditions differ in the size, if exists, of the outside option. Out50 and Out60 are treatments where the buyers have an outside option whose value is 50 and 60, respectively. Out0 is the control condition where the buyers do not have an outside option. It is worth noting

(3.1) holds if and only if $e^{-r\Delta} v < \max\{v - p_0, w\} = w$.

[26] Round n in the experimental setting refers to period $n-1$ in Section 2.

that we regard the comparison of Out50 and Out60 as a placebo check to examine whether there are some behavioral factors that are not considered in any of the proposed theories. Although the two treatments are not qualitatively different in terms of their theoretical predictions, equilibrium reasoning or cognitive loads for the positive selection might be easier and faster when the outside option is 60, as the probability that the buyer draws a value less than or equal to 60 is distinctively larger than the case of the outside option 50. Thus, confirming similarities between Out50 and Out60 could help us ensure that our observations are not driven by other factors to which we have not paid attention. If Out50 and Out60 are not qualitatively different, we collectively call Out50 and Out60 as OutYes, and accordingly Out0 as OutNo.

After each rejection of the price offer, the sellers were asked to report their beliefs about the paired buyer's types before making another price offer. More precisely, we asked them to report the range of the possible values within which the buyer's value falls. This report enables us to measure the seller's belief about how much positive selection and negative selection have occurred in the current round. Connecting this measurement with the seller's offer in the next round, we can study how the seller responds to positive and negative selections. To make it as incentive compatible as possible for seller participants to report their beliefs truthfully, we presented the minimum and maximum values for the buyer types reported in Round n in the decision screen of the same player in Round $n+1$ so that the player can potentially utilize the information to make a better decision.

The single-round game is essentially Out50 without the second rounds and beyond. In the sense that the seller makes a take-it-or-leave-it (TIOLI) offer knowing that the buy has an outside option whose value is 50, we call this treatment TIOLI50. It is identical to the ultimatum bargaining experiment with private information: The seller makes an ultimatum offer, and the buyer accepts it or rejects it. In the case of rejection, the buyer takes the outside option. Since the buyer only accepts the offer when it is better than taking the outside option worth 50, the seller's optimal strategy is to offer the full commitment price that maximizes the expected payoff. Thus, the equilibrium predictions in TIOLI50 are identical to those of Out50. By reporting the findings from this additional experiment at the end of Section 5, we aim to address other concerns unaccounted for in our initial design, such as possible lack of understanding optimal pricing strategies and the buyers' potential concerns other than their monetary payoffs. However, for expositional consistency, testable hypotheses stated in Section 4.3 will only address infinite horizon bargaining games. Our final experimental design is summarized in Table 2.

Our experiment was conducted by oTree (Chen et al., 2016) at HKUST. A total of 268 (196 for the initial three treatments) subjects were recruited from the graduate and undergraduate population of the university. We had 4 sessions each for treatments Out0, Out50, and TIOLI50, and 5 sessions for treatment Out60. Each session consisted of 12 to 18 participants and we had 44, 58, 66, 72 participants in treatments Out0, Out50, Out60, and TIOLI50, respectively. In all sessions, subjects participated in seven matches of the bargaining game described above under one

Table 2: Experimental Design

		Outside Option Value		
		∅	50	60
#Rounds	infinite horizon	*Out0* (*OutNo*)	*Out50* (*OutYes*)	*Out60* (*OutYes*)
	single round		*TIOLI50*	

* Each participant has seven newly paired supergames (matches).
* Continuation probability to the next round is 0.8.
* Buyer's value v is drawn from $U[50, 400]$.

treatment condition. Sample experimental instructions can be found in Online Appendix A.

All sessions were conducted via the real-time online mode using Zoom and oTree. Upon arrival at the designated Zoom meeting, subjects were instructed to turn on their video. Each received a web link to the experimental instructions. To ensure that the information contained in the instructions is induced as public knowledge, the instructions were presented and read aloud by the experimenter via Zoom. All questions were privately addressed via the chat function in Zoom.

One of the seven matches was randomly selected for each subject's payment. The payoffs a subject earned in the selected match were converted into Hong Kong dollars at a fixed and known exchange rate of HK$1 per token. In addition to these earnings, subjects received HK$40 as a show-up payment. Subjects on average earned HKD 115 (≈ USD 16) by participating in a session that lasted 1.2 hours. They were paid electronically via the autopay system of HKUST into the bank account he or she has registered with the Student Information System (SIS).

4.2 Discussions on the Experimental Design

In this subsection, we mention two important design features that are worth discussion.

First, we asked sellers to report their beliefs without financial incentives. Although this belief reporting is not financially incentivized, we argue no substantial reasons to believe that subjects misreport their beliefs. Incentive compatible mechanisms at the end of every bargaining round could have been considered, but Burdea and Woon (2022) provide evidence that the quality of reported beliefs may depend less on the formal incentive compatibility properties of the elicitation procedure than on the difficulty of comprehending the elicitation task and how well incentives induce cognitive effort. Danz et al. (2022) also report that the information on the incentives for belief elicitation may lead to a systematic bias toward the center. By explicitly stating that the purpose of the belief reporting is to help the subjects make non-erroneous decisions in the subsequent round, we believe that our elicitation is simple yet reasonable enough to induce cognitive effort. Moreover, the belief reporting could function as a suggestion or a guideline for the future self, as similarly done in

Halevy et al. (2021). If we regard this belief reporting as communication between team partners—the current self and the future self—then truthful reporting is more convincing (Burchardi and Penczynski, 2014).

Second, by adding the outside option without adjusting the value distribution, it becomes nontrivial to claim that equilibrium predictions for seller profits and initial prices with outside options are necessarily higher than those without. This is because the net gains from trade are shifted by the value of the outside option. For illustration, consider an extreme case where the buyer's value is uniformly distributed on $[50, 400]$, and the outside option is worth 399. Without the outside option, the seller would surely earn a profit of 50 by setting the price at the lower bound of the distribution, 50. With the outside option, however, any buyer with $v \leq 399$ would strictly prefer the outside option at any nonnegative price, leaving the seller little opportunity to profit. Although our actual experiments with $w = 50$ and 60 do not substantially alter the value distribution, and the equilibrium predictions in Table 1 remain consistent with our hypotheses, the directional effect marginally blurs the qualitative role of outside options in bargaining outcomes. We acknowledge this limitation. Nonetheless, holding the net value distribution constant would have required simultaneously changing both the distribution and the outside option values, which would conflict with our initial goal of isolating the effect of introducing an outside option. Importantly, under our parametric setup, the equilibrium predictions still differ substantially across treatments, so we believe it remains appropriate to compare experimental results with and without outside options.

4.3 Hypotheses

We set the differences between theoretical predictions with and without the outside option as our null hypotheses, so the qualitative support for the null hypotheses is closely related to what we summarized in Section 2.

We list our hypotheses in the order of our thought processes, rather than the order of subjective importance of the theoretical predictions. Although less crucial from the theoretical perspective, our first hypothesis is that the size of the outside option does not affect the delay in bargaining. Whenever with an outside option, no delay occurs in equilibrium.

Hypothesis 1 (Irrelevance of the Size of Outside Option)**.** *No significant differences between Out50 and Out60 exist regarding the bargaining length.*

Although the seller's equilibrium price offer and the buyer's response would depend on the value of the outside option as stated in Proposition 4, the differences between 50 and 60 are meager compared to the support of the type distribution. So, the qualitative difference should be minimal in equilibrium. If Hypothesis 1 is rejected, it might imply that some crucial factors are not accounted for in any of the proposed theories, or that some assumptions we postulate would not hold. Perhaps it could mean that the participants' cognitive loads for unraveling positive selection play an important role, or some other unintended aspects become salient. We present the following hypotheses focusing on the difference between OutYes and OutNo, assuming that rejecting Hypothesis 1 is not the case.

The next hypothesis regards the bargaining efficiency in terms of the delay. When an outside option is unavailable to buyers, theory predicts some delay (although the delay shrinks to zero as the continuation probability approaches 1). On the other hand, when there is an outside option, positive selection takes place such that no delay is expected in equilibrium. We say that bargaining ends either when they reach an agreement or when the buyer takes the outside option.

Hypothesis 2 (Outside Option and Delay)**.** *The average number of bargaining rounds in OutNo is strictly larger than 1, while that in OutYes is 1.*

Hypothesis 2 contains two testable statements. We examine (1) whether the average length of bargaining in OutNo is longer than OutYes and (2) whether bargaining in OutYes ends in round 1.

The next three hypotheses are about the seller's profits, actions, and beliefs.

Hypothesis 3 (Seller Profits)**.** *The seller's profits are higher in OutYes than in OutNo.*

Theory predicts that without an outside option, negative selection arises so that the seller's profit gradually decreases to the lowest level, while the seller achieves the monopoly profit with the presence of an outside option. This prediction leads us to hypothesize that the seller's profit is higher in OutYes than in OutNo. Specifically, the commitment price offer of the seller is 175 in Out50 and 170 in Out60, so the expected profit is 87.5 in Out50 and 82.5 in Out60. In OutNo, not every bargaining agreement is made in round 1; with the parameters we considered for the lab experiment, the upper bound of the initial offer is 122, and it decreases down to 86, 62, 51, and 50 in subsequent rounds. In this path, the seller's expected payoff is 76.21.

Hypothesis 4 (Offering Prices)**.** *The average offer price is smaller in OutNo than in OutYes.*

The following elaborates more on Hypothesis 4. Without the outside option, the seller believes that the higher buyer types trade earlier, and thus, the price offer for the remaining buyers decreases over time until it reaches $p = \underline{v}$. Meanwhile, with the outside option, it is hard to tell the price dynamics because the equilibrium predicts that the seller offers p_w^* in round 1, and bargaining ends immediately. One testable implication from these two predictions is that the average offered price would be smaller in OutNo than in OutYes.

After observing a buyer's rejection, the way the seller updates her belief differs depending on the presence/absence of the outside option. Without the outside option, the buyer whose type is higher would be more likely to accept the previous price offer, so the minimum of the belief distribution (the lowest possible type remaining in the market) would be unaffected. With the outside option, however, positive selection occurs to the seller's belief in any history of any equilibrium. That is, the seller would believe that all buyer types lower than a history-dependent cutoff would end bargaining by immediately exercising the outside option. This reasoning leads to the following hypothesis.

Hypothesis 5 (Belief about the Type Distribution)**.** *When the seller's offer is rejected, her belief about the minimum type of buyer is an increasing function of the size of the outside option in OutYes, but it is constant at* $\underline{v}$ *in OutNo.*

Regarding Hypothesis 5, it is worth mentioning that examining the players' actions after the second round is challenging because when the outside option exists, the standard theory does not provide a concrete testable implication. The equilibrium predicts that bargaining ends instantly, and a delay arises only if the buyer deviates from the equilibrium play. As the standard definition of PBE does not restrict the seller's posterior belief after such a deviation, the equilibrium analysis does not make any testable prediction for the seller's belief after any non-null history. However, as discussed in Section 3.2, we can argue based on the dominance reasoning that the seller's posterior belief after round 1 should assign zero probability to buyer types below w/δ (62.5 in Out50 and 75 in Out60), provided that the seller's initial offer was no less than $(1-\delta)w/\delta$ (12.5 in Out50 and 15 in Out60). From the perspective of such a buyer, even when the price in round 2 is 0, taking the outside option to earn w is better than the expected gains from trade in round 2, $\delta(v-0)$. Thus, we mainly examine whether the minimum of the stated belief after being rejected in round 1 is greater than w/δ.

The next hypothesis regards the buyer's behavior.

Hypothesis 6 (Buyer Actions)**.** *In OutYes, low-type buyers take the outside option, and high-type buyers accept the seller's offer in round 1. In OutNo, some buyers reject the seller's price offer and delay in round 1.*

The first part of Hypothesis 6 restates the theoretical prediction by BP (Propositions 1 and 4 in Section 2): For the case that the buyer has an outside option, there is a cutoff value such that the buyer types with a valuation higher than this cutoff (i.e., high-value buyers) trade with the seller immediately, while other buyer types (i.e., low-value buyers) exercise the outside option immediately. Then, a testable implication is that the buyer never rejects the seller's equilibrium offer or delays the negotiation in round 1. On the other hand, when the buyer has no outside option, we would observe some rejections and delays in round 1 in equilibrium (Proposition 3).[27]

[27] Although Hypothesis 6 focuses on the buyer's action on the path, particularly in round 1, we can also extend the same hypothesis to the play off the path. For example, by generalizing the proof of Proposition 1 in Board and Pycia (2014, p.660), one can prove that a delay never occurs after any history when the buyer has an outside option. On the other hand, when the buyer has no outside option, trading at $p > \underline{v}$ is not acceptable to the lowest buyer type $v = \underline{v}$. Hence, a delay should occur with positive probability after any history unless the standing offer is lower than $\underline{v}$.

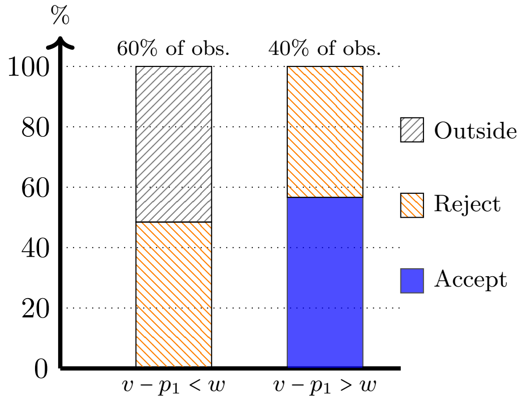

Figure 1: Round-1 States

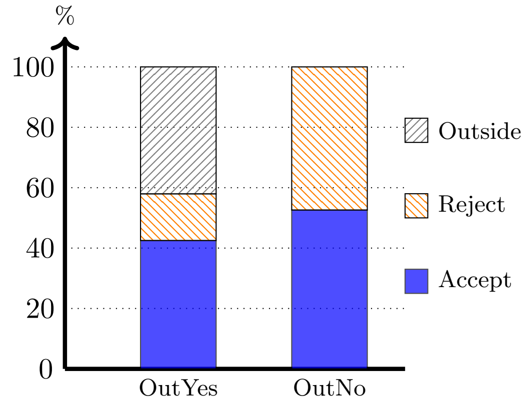

Figure 2: End-of-Bargaining States

5 Results

In this section, we report experimental findings in the corresponding order of our hypotheses.

First of all, we compare the observations from Out50 with those from Out60. As described in Hypothesis 1, theory predicts no differences between the two treatments in terms of the length of bargaining and negligible differences in terms of buyer's and seller's earnings and seller's initial offers. As predicted, the average length of bargaining (Out60=1.53, Out50=1.59) is almost identical (t-test,[28] p=0.558), and the bargaining length distributions cannot reject the null hypothesis that those come from the same population distribution (Kolmogorov-Smirnov two-sample (KS) test, p=0.354). Moreover, the seller's average earnings (Out60=59.03, Out50=68.42) and the seller's average initial offers (Out60=194.45, Out50=211.48) are different in a statistically insignificant manner (p-values are 0.213 and 0.167, respectively). The buyer's average earnings in Out50 (78.72) and in Out60 (94.82) are significantly different (p=0.025). However, this is mostly due to the fact that the exercise of the outside option in Out60 (46.52% of the entire cases) renders an additional earning of 10 relative to Out50. If the buyer's earnings from the outside option were 50 instead of 60, the difference in the buyer's average earnings should become insignificant (p=0.116). Accordingly, the distributions of the seller's earnings are not statistically different between Out50 and Out60 (KS test, p=0.207). From now on, unless stated otherwise, we pool Out50 and Out60 as OutYes and focus on the differences between OutYes and OutNo.[29]

Result 1. *No significant differences between Out50 and Out60 are found.*

Regarding Hypothesis 2, we compare the length of bargaining in OutNo with that in OutYes.

[28]Unless otherwise stated, p-value reported in the parentheses is the result of a two-tailed t-test comparing the means, and the standard error of the mean difference is clustered at the individual level as each individual played the same role for the entire seven matches.

[29]In Online Appendix B, some figures in the main context reappear with Out50 and Out60 being separated from OutYes.

On average, it takes 2.89 rounds to end the bargaining process in OutNo (either in the form of agreement or termination), while it takes 1.56 rounds in OutYes.[30] This difference is statistically significant (p<0.001,) and it is primarily driven by the exercise of the outside option, not by different degrees of acceptance. Figure 1 shows how the buyers in OutYes behave. About 48% of the buyers who are supposed to immediately take the outside option (i.e., buyers with $v - p_1 < w$) rejected the first-round offer. Even 42% of the buyers whose value is high rejected the first-round offer, clearly negating the positive selection reasoning. Figure 2 shows the proportion of bargaining matches that end with an acceptance of the offer (in blue), a random termination after a rejection (in orange), and the exercise of the outside option (in gray), respectively. In OutNo, 47.40% of the bargaining matches end with random termination after the buyer rejects an offer, but in OutYes, that proportion plummets to 15.46%. It is worth noting that a substantial proportion (42.04%) of the bargaining matches end by the buyer exercising the outside option in OutYes, and the proportion of matches ending with the acceptance of an offer (42.50%) is still comparable to that in OutNo (52.60%). These observations support our claim that the difference in the length of bargaining is driven by the exercise of the outside option, which otherwise may end up with random termination after a rejection.

The observed difference in length of bargaining between OutYes and OutNo does not imply that we should entirely confirm Hypothesis 2. In OutYes, where the delay should not be observed in theory, the average length of bargaining rounds is larger than 1 in a statistically significant manner (p<0.001). Moreover, in 52% of the 260 cases where a buyer's payoff of accepting the first-round offer is strictly smaller than the outside option ($v - p_1 < w$), the buyer rejected the first-round offer to move on to the next bargaining round. This evidence is a clear negation of the positive selection reasoning: If buyers anticipate that the next-round price offer would be non-decreasing as the remaining demand pool consists of higher types, those whose value is below $p_1 + w$ should have exercised the outside option immediately.

Result 2. *The average number of bargaining rounds in OutNo is strictly larger than that in OutYes, and the average number of bargaining rounds in OutYes is strictly larger than 1.*

Figure 3 shows the seller's average earnings in OutYes and OutNo, along with the aggregate-level standard deviations, the full-commitment benchmarks, and the equilibrium payoffs. Hypothesis 3 states that the seller's profits are higher in OutYes than in OutNo. We find the opposite from the lab: The seller's average earning in OutYes (63.43) is smaller than that in OutNo (78.25). The difference is statistically significant at the 5% level of significance (one-sided, p=0.039). We thus reject Hypothesis 3 stating the seller's equilibrium profit in OutYes would be larger than that in OutNo. Note that since we have precise numerical equilibrium predictions for each treatment, if we compare the deviation from the theoretical payoff, then the statistical significance would have been

[30] Figure B.1 in Online Appendix B shows the average number of rounds by match.

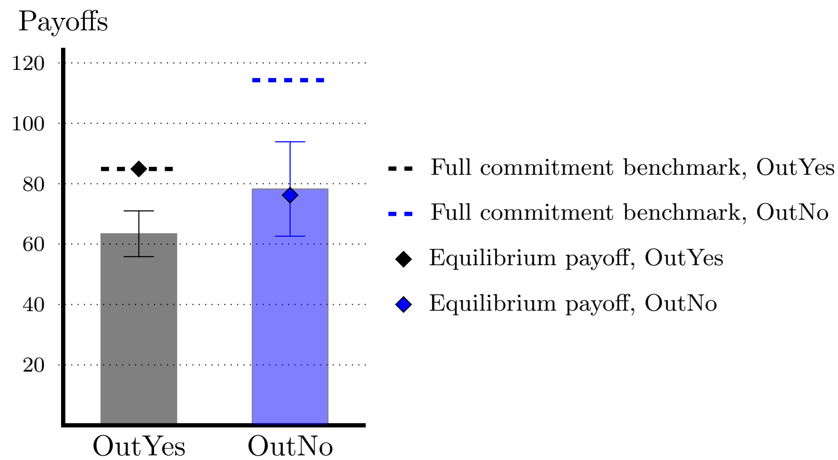

Figure 3: Seller's Earnings

stronger. Another noticeable observation is that the equilibrium payoff level in OutNo falls inside the confidence interval of the seller's average payoff. However, this "coincidence" should not be interpreted as supporting evidence that the negative selection precisely worked in the lab because in fact these payoffs arose from prices that were far from the Coasean equilibrium prediction, as we will next show. We shortly revisit this observation when examining the dynamics of the sellers' price offers.

Result 3. *The seller's average profit in OutYes is smaller than that in OutNo.*

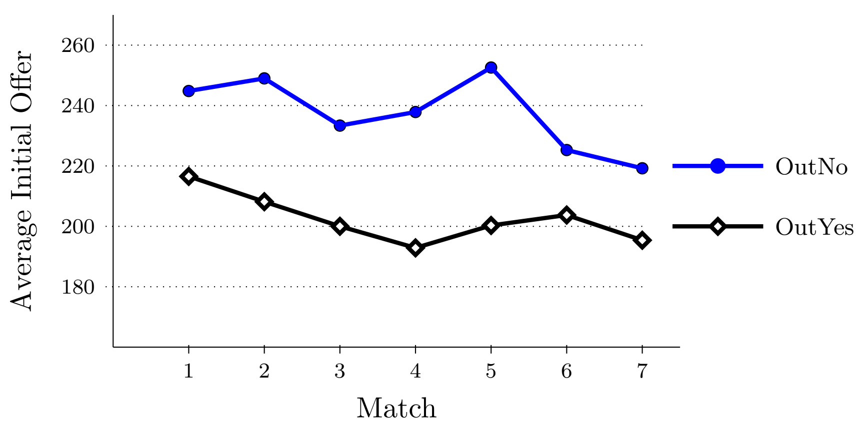

Figure 4: Round 1 Offer across Match

The seller's initial offers are also inconsistent with Hypothesis 4. Figure 4 shows the average initial offers by match. Overall, the average initial offer in OutYes (202.44) is smaller than that in OutNo (237.45), and the difference is statistically significant (p=0.003). The average initial offer in all our treatments is significantly larger than the full commitment price offer (200 in OutNo, 175 in Out50, and 170 in Out60) that can theoretically maximize a monopolist's expected payoff. Equilibrium predictions for initial prices are those full commitment prices in Out50 and Out60, whereas in OutNo the equilibrium prediction is even lower (121.56). Both in OutYes and in OutNo, we can observe a small but significant downward trend of the average initial offer over time.[31] We do not observe important time trends by match for our other variables. The declining trend in initial prices suggests that sellers in both treatments learned to make lower initial offers over time. While those trends represent a move toward the equilibrium in both treatments, they are similar in magnitude so that the (non-hypothesized) difference between the treatments is preserved.[32] These trends suggest that the forces influencing sellers were qualitatively similar in both treatments. In fact, as we will show later, lower initial offers resulted in higher seller profits in both treatments, so the observed trend in initial seller prices improved their profits. Even in the first match, the average initial offers (OutNo=244.82, OutYes=216.55) are different in a statistically significant manner (p=0.023). This observation implies that the sellers, who consider the possibility for the buyers to exercise their outside option, proactively lower the offering price to avoid zero gains from trade. Thus, it strongly rejects Hypothesis 4.

Result 4. *The seller's initial offer is strictly larger in OutNo than in OutYes, and in both treatments it is strictly larger than the respective equilibrium price offer.*

As we acknowledge in Subsection 4.2, having an outside option without changing the range of the values may make the comparison between two treatments harder. Despite this, we believe the fact of slightly higher seller profits (reported in Result 3) and initial prices without outside options (reported in Result 4), as well as other evidence still strongly suggest that bargaining with and without the outside option is qualitatively similar, even though equilibrium reasoning suggests that they should be substantially different.

Figure 5 shows how the average price offers change across rounds among the pairs whose bargaining lasted in round 2 (Figure 5a) and in round 3 (Figure 5b).[33] It is clear that the average price offer tends to decline over rounds in both OutYes and OutNo. The declining trend appears similar to the theoretical prediction of the Coasean equilibrium (presented with a dotted line). Yet the price offers are much larger in level than what the negative selection predicts, implying that negative selection is in force but only partially executed in the lab.

Another noticeable finding is that even with the outside option, the seller's price cut (the difference between the initial offer and the round-2 offer) is much larger than the theoretical bound,

[31] After controlling for the treatment effect, the average initial offer decreases by 2.8853 tokens in the next match (p=0.032).

[32] After controlling for the treatment effect and the learning effect, the interaction effect of the treatment and learning is insignificant (p=0.668).

[33] We did not show the average price offers over rounds using all available price offers because such observations are exposed to the survivorship bias. The data from the second round and beyond are collected only when the first round price offer is rejected and the bargaining process is not terminated, and the number of observations sharply decreases with rounds.

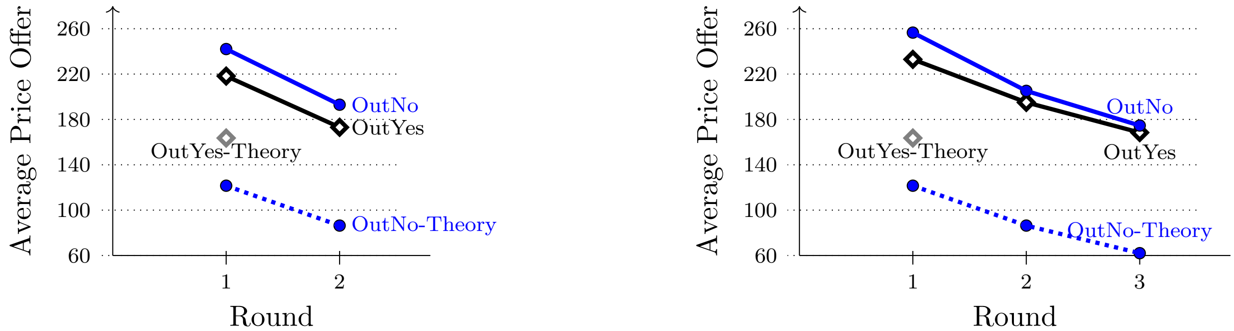

(a) Subsamples: Bargaining Length ≥ 2 (b) Subsamples: Bargaining Length ≥ 3

Figure 5: Offer across Round

$\frac{\delta}{1-\delta}w$. We mean by theoretical bound the largest possible off-the-equilibrium price cut that the seller could make: Let v^* denote the buyer type such that $v^* - p_1 = w$. If the buyer and the seller play a PBE (consistent with BP), v^* finds it suboptimal to delay. Thus, $v^* - p_1 = w \geq \delta\mathbb{E}[v^* - p_2|p_1]$, which leads to $\mathbb{E}[p_1 - p_2|p_1] = \mathbb{E}[v^* - w - p_2|p_1] = \mathbb{E}[v^* - p_2|p_1] - w \leq \frac{w}{\delta} - w = \frac{1-\delta}{\delta}w$. It means that, theoretically speaking, the seller should not decrease round 2 price more than 12.5 ($=\frac{0.2}{0.8}50$) in Out50, and 15 in Out60 ($=\frac{0.2}{0.8}60$).

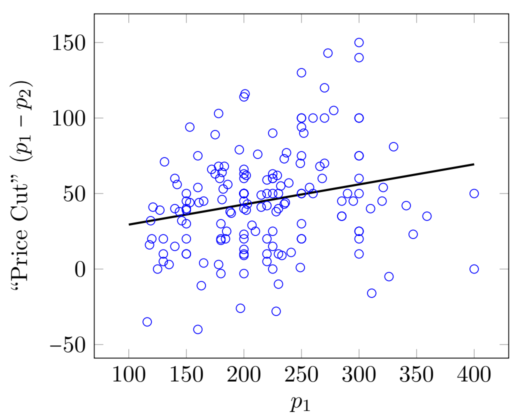

Figure 6: Scatterplot of $p_1 - p_2$ on p_1

Figure 6 shows that the empirical gap between p_1 and p_2 is clearly larger than the theoretical bound. The black solid line is the fitted line for all observations, $p_1 - p_2 = 0.1331p_1 + 16.124$. The expected price cut implied by this trend is strictly greater than $(1-\delta)w/\delta$ (which is at most 15) for all p_1. Thus, any buyer who was close to indifferent between accepting p_1 or taking her outside option, $v - p_1 \approx w$, would have received a strictly higher expected payoff by rejecting and waiting until period 2, given the anticipated reduction of prices. As we will show later, it is precisely those buyers who typically did reject. Such aggressive price cuts are not specific to OutYes.[34]

Hypothesis 5 regards the seller's belief after observing the rejection of the initial offer. At the end of every unfinished bargaining round, we asked the seller to guess the range of the possible valuations that the buyer draws. As explained in Hypothesis 5, the minimum of the guess should have increased, at least, to w/δ (62.5 in Out50 and 75 in Out60). The average minimum of the guess in OutYes (81.58) is larger than that in OutNo (71.36), and the difference is statistically significant only at the 10% level (p=0.059).[35] The average minimum of the guess in OutYes is also significantly larger than w/δ (p=0.016), so we cannot reject Hypothesis 5 directly. Yet it is too hasty to take the observed difference as suggestive evidence that the sellers in OutYes reflect on the positive selection. Since some sellers move on to Round 2 more frequently than others, the (unweighted) averages tend to overweight such sellers' reports. To correct the misinterpretation due to the heterogeneous individual weights, we also examine the individual's average minimum of the reported range. Almost half of the sellers (29 out of 59 sellers) in OutYes report it below w/δ on average, and we cannot reject the null hypothesis that the sample distribution of such averages in OutYes is from the same population distribution for those in OutNo (KS, p-value=0.328). In other words, the individual-level reports on the minimum of the guess in OutYes and those in OutNo are, by and large, the same.

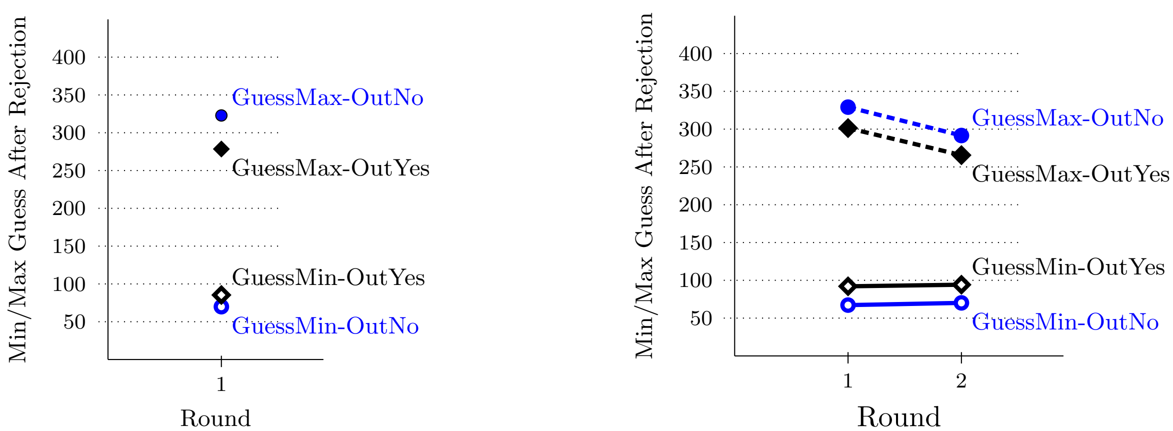

(a) Subsample: Bargaining Length ≥ 2

(b) Subsample: Bargaining Length ≥ 3

Figure 7: Minimum/Maximum of the Guess across Round after Rejection

Figure 7 presents how the average (min and max) guesses across round after rejection among the pairs whose bargaining lasted to round 2 (Figure 7a) and in round 3 (Figure 7b). The decreasing

[34]A related analysis in Online Appendix C shows that a similar pattern of excessive price cutting also arises in OutNo. In Online Appendix C, we construct a lower bound on the probability that rejection of the first-round offer occurs using observed price data, under the assumption that buyers anticipate future price cuts in the next round. The analysis shows that this lower bound exceeds the theoretical benchmark for a wide range of first-round prices, reflecting the aggressive price cuts made by sellers.

[35]When looking at the last four matches only, we still observe that the average minimum of the guess in OutYes is larger than that in OutNo, but the difference is statistically insignificant (p=0.248).

trend of the maximum guess in OutNo provides supplementary evidence that the negative selection takes place, albeit partially, as the Coasean equilibrium predicts. The minimum of the guess is overall stable over time in OutNo, while the maximum of it decreases steadily, which implies that the sellers understand that the buyers with high valuations gradually leave the market when outside options are unavailable. It is noticeable that a similar pattern of negative selection is also observed in OutYes, implying that even when the outside option exists, the sellers believe that buyers with high valuations gradually leave the market.

Result 5. *The seller's reported belief about the lowest type of buyer who rejects the initial offer in OutYes is weakly higher than that in OutNo on average. However, the difference is not substantially large in its magnitude.*

We should, however, acknowledge that Result 5 should be appreciated with caveats for two reasons. First, the sellers' guesses were elicited without monetary incentives. Second, underlying reasoning for examining the minimum of the guesses is that the low type buyers should have exercised the outside option as soon as possible, but this reasoning relies an assumption of no noises and mistakes, e.g., no low value buyers ever "accidentally" reject.

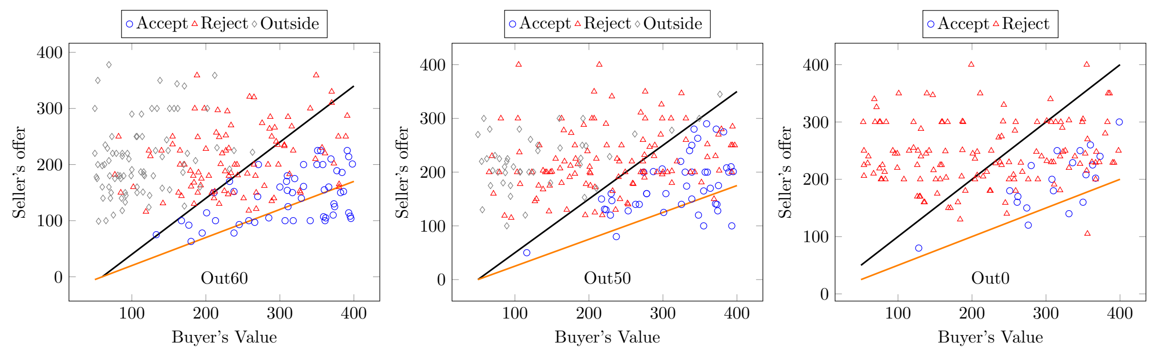

Figure 8: Buyer's Action in Round 1, by Treatment

This figure shows the buyers' round-1 decisions on the plane of the buyer's value (x-axis) and the seller's price offer (y-axis). The scatterplot point (x, y) reads that the buyer's value is x and the seller's offer is y. Blue circles, red triangles, and gray diamonds respectively indicate that the buyers accept the offer, reject the offer, and take the outside option.

Our last set of results regards the buyers' decisions. Figure 8 shows the buyers' actions in Round 1 on the plane of the buyer's value and the seller's offer. Each scatterplot point (x, y) reads that the buyer's value is x and the seller's offer is y. Blue circles, red triangles, and gray diamonds respectively indicate that the buyers accept the offer, reject the offer, and take the outside option. We draw two auxiliary lines to grasp the overall patterns. The black 45-degree lines show the buyer's value minus the outside option, which means that when the value–offer coordinate is on the right-hand side of the line, accepting the offer is strictly better than taking the outside option.

The orange lines show half of the buyer's value minus the outside option, which means that the value–offer coordinate on the orange line equally splits the gains from trade between the buyer and the seller.

Some observations are consistent with positive selection. For example, in both Out60 and Out50, the vast majority of buyers with low values take their outside option in round 1, unlike buyers with high values. This makes the seller more likely to face high-value buyers in later rounds.[36] To provide corroborative evidence, we ran a logit regression of the first-round rejection on the buyer's value, OutYes dummy, and the interaction of those two variables. The fitted line is $\ln(\text{Odds of Reject}) = 5.1287 - 5.8367 OutYes - 0.01202 Value + 0.01491 OutYes \times Value$, where all estimated coefficients are statistically significant at the 1% significance level. It reads that one unit increase in $OutYes \times Value$ will increase the odds of rejection by 1.5%, which implies that the remaining demand pool has, on average, higher value buyers.

However, rejections are prevalent in OutYes, contrary to the prediction that buyers either accept the offer or take the outside option in round 1. Furthermore, across all treatments and holding the seller's offer fixed, we observe that buyers with higher values are more likely to accept the offer; as shown in Figure 8, accepted observations (blue circles) are concentrated in the lower-right region. These observations suggest that Hypothesis 6 is not supported. The observed rejections in OutYes can be partly explained by sellers cutting price offers too aggressively in the next round, as illustrated in Figure 6. Indeed, buyers appear to have anticipated the incentives created by excessive price cutting: buyers with $v - p_1 \approx w$ (depicted around the black 45-degree lines) largely reject the offer in OutYes rather than exercising the outside option. In the absence of the outside option, a majority (87.01%) of buyers also reject the round-1 price offer, exceeding the theoretical prediction (about 60%; see the discussion after Table 1). Two factors contribute to this discrepancy between theory and data. First, the average first-round price (237.45) is higher than the theoretical prediction (121.56) and therefore less attractive to buyers. Second, as in OutYes, sellers cut their price offers too aggressively in the second period.[37]

Another possible story could go along with the orange lines, the equal-split lines of the gains from trade. Below the orange lines, we observe a dominant fraction of acceptance while we observe a substantial fraction of rejection on the opposite area. Although one may interpret that the buyers' inequity aversion would have played a role to make their accept/reject decisions, it is challenging to explain the overall observations using the inequity aversion only. Even if we explicitly consider an inequity-averse player (where a formal model and its results are presented in Online Appendix F), the buyer's decision with the presence of the outside option ought to be either accepting the offer or exercising the outside option, so it does not help us explain the frequent rejections. Moreover,

[36] The average buyer values in each round are monotonically decreasing from 227.2 in Round 1 to 178.4 in Round 4 in Out0 while the decreasing trend is much slower in Out50 (from 228.7 in Round 1 to 219 in Round 4) and it is not monotonically decreasing in Out60 (from 212.9 in Round 1 to 260.9 in Round 4).

[37] A detailed analysis for OutNo appears in Online Appendix C.

the inequity aversion does not explain why the seller's average initial offer in OutYes is smaller than that in OutNo.

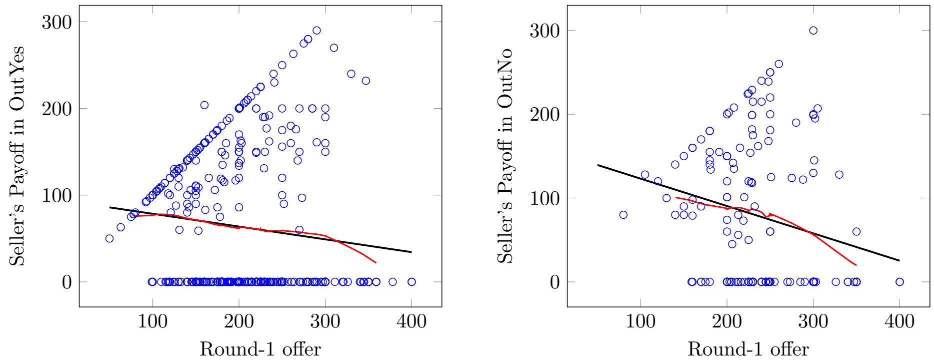

Figure 9: Seller's round-1 offer and payoff

These scatterplot figures show the seller's round-1 price offer (x-axis) and the final payoff (y-axis). Each black solid line is the fitted line, and each red line is a locally-weighted scatterplot smoothed line.

Finally, we briefly explore the drivers of the seller's pricing behavior. While not part of our formal hypotheses stated in Section 4.3, Figure 9 illustrates the relationship between the seller's Round 1 offer and their realized payoff. The scatter plots reveal the binary nature of the outcomes: Offers are either accepted (lying on the 45-degree line) or result in the buyer's exercise of the outside option (lying on the horizontal axis). The fitted line captures the average realized payoff, essentially the seller's expected profit, conditional on the offer price. In both treatments, we observe a negative relationship: lower initial offers are associated with higher average payoffs within the observed range. This empirical pattern suggests that the declining trend in initial offers shown in Figure 4 represents a learning process where sellers adjusted their prices toward levels that yielded higher expected returns. It also indicates that the "optimal" price given actual buyer behavior was lower than the theoretical commitment price, consistent with the presence of Coasean forces in both treatments.

In sum, our experimental evidence shows that no qualitative differences exist between the two environments with and without the outside option. Except that the length of bargaining is shorter with the outside option, our findings go against or do not support the theoretical predictions. Most importantly, positive selection, the key driver for the main theoretical prediction of BP, does not arise in the laboratory.

These discrepancies give us a rationale to explore some factors that may hinder unraveling. Our additional treatment, TIOLI50, serves as another benchmark comparison to the Out50 treatment because the equilibrium payoffs, the seller's equilibrium price offer, and the buyer's acceptance decision should be identical in both treatments in theory. As we reported in Result 4, the seller's initial offer price in Out50 is significantly higher than the equilibrium price, but the average seller price in

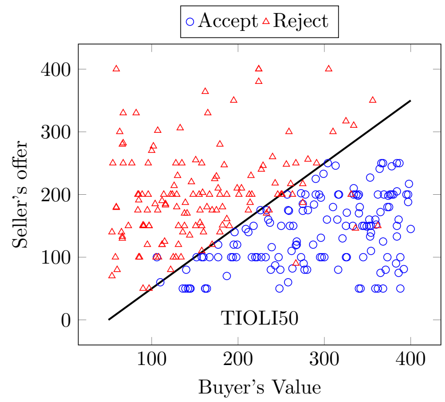

Figure 10: Buyer's Action in TIOLI50

This figure shows the buyer's decision on the plane of the buyer's value (x-axis) and the seller's price offer (y-axis). The scatterplot point (x, y) reads that the buyer's value is x and the seller's offer is y. Blue circles and red triangles respectively indicate that the buyers accept the offer and rejected the offer. The black solid line shows the buyer's indifference line between accepting and rejecting the offer.

TIOLI50 (166.92) is not statistically different from the commitment price of 175 (p-value=0.391). This supplementary finding suggests that the sellers on average know the profit-maximizing offer in a TIOLI setting, so the reason why our experimental findings are inconsistent with BP's predictions is more likely due to the dynamics of the bargaining setup, not the sellers' limited understanding of the pricing strategy.

Also, unlike the typical findings from the ultimatum bargaining experiments that unfair offers are often rejected despite monetary benefits, the buyers in the TIOLI50 accept offers if and only if the price is less than their net value. The scatterplot in Figure 10 clearly shows that buyers in TIOLI50 react almost exactly as theory predicts. The x-axis and y-axis of the scatterplot represent the buyer's value and the seller's offer, respectively. Red triangles indicate rejected offers, and blue circles indicate accepted ones. The black solid line shows the buyer's indifference line between accepting and rejecting the offer. It is evident that most of the red triangles are on the left of the black line, and most of the blue circles are on the right of it, indicating that most of buyers react as theory predicts. This finding suggests that fairness concern does not affect the buyers' decisions, which serves as another good justification for using the equilibrium with money-maximizing subjects as a baseline for other hypotheses.

6 Discussions

We claim that a key driver making the positive selection fragile is the belief that some low-type buyers may remain in the market. In this section, we explore various scenarios in which this belief naturally emerges, including buyer optimism, stubbornness in bargaining, delays in the arrival of outside options, insufficient skepticism in information disclosure, and a lack of first-order rationality.

6.1 Optimism

The discrepancies between our experimental findings and the prediction of BP give us a rationale to explore some factors that may hinder unraveling. Motivated by our experimental observation that some buyers reject the current round price offer instead of taking the outside option optimistically believing that the next round offer will be likely more favorable, we present and analyze a model of bargaining with optimistic buyer types[38] in Online Appendix D. Specifically, in our model, the buyers incorrectly believe that the seller will tremble and offer a low price in the next period with some probability. We show that the theoretical difference created by the absence/presence of an outside option collapses once we introduce a small fraction of optimistic buyers (Proposition D.2).

The intuition is straightforward. As they believe that the seller will concede to a low price with some probability in the future, the optimistic buyer types will refuse to end the negotiation by exercising the outside option; these optimistic buyer types are the ones who are tough to deal with from the seller's perspective. Then, rational buyer types can leverage the seller's fear that the buyer is of an optimistic type by mimicking those optimistic types and refusing to exercise the outside option. As a result, positive selection among buyer types does not occur sufficiently, and hence the seller still has an incentive to practice inter-temporal price discrimination for the remaining buyer types.

6.2 Obstinacy in reputational bargaining

The reputational bargaining literature has demonstrated that an inefficient delay may arise when a small fraction of bargainers are known to be (irrationally) obstinate. Rational bargainers pretend to be obstinate to derive a more favorable deal, hence an inefficient delay may arise in equilibrium (Myerson, 1991; Abreu and Gul, 2000). Embrey et al. (2017) report experimental evidence consistent with the theoretical predictions about the inefficient delay in reputational bargaining. Compte and Jehiel (2002) show that the existence of an outside option may cancel such effect of obstinacy. The bargaining counterpart's outside option limits the rational bargainer's benefit from mimicking the obstinate type because the counterpart would exercise the outside option immediately rather than haggling with the obstinate type.

The standard assumption in the reputational bargaining literature is that the gains from trade

[38] Considering optimistic buyers in the negotiation process is not new in the literature. The deadline effect and the learning about bargaining power would make optimistic players costly delay the agreement to a later period (Yildiz, 2011). The second-order optimism—belief about the other party's optimism toward her prospects—could also play a role in causing a delay (Friedenberg, 2019). Mutual optimism in the Rubinstein bargaining environment can be a source of delayed agreement and substantial efficiency loss (Li and Wong, 2009). Empirical evidence also supports the role of optimism in the negotiation process. For pretrial settlement negotiations, the involving parties' optimism about the judge's decision may lead to a substantial delay (Vasserman and Yildiz, 2019).

(net-values) are common knowledge, hence whether a bargainer is an obstinate type is the only source of information asymmetry. Fanning (2023) considers a bargaining problem such that the buyer's obstinacy as well as his value of the seller's good and outside option are the buyer's private information. Fanning shows that, in contrast to the BP's result, the Coase conjecture is partially restored when the buyer is possibly of an obstinate type. A rational buyer type may benefit from mimicking an obstinate type rather than exercising the outside option immediately, and such a possibility undermines the effect of positive selection.

Although the precise logical steps for reaching the above-mentioned theoretical conclusions differ, one shared driver behind the results seems to be that the buyers' belief that the sellers might make more favorable offers could lead to inefficient delay. In the sense that our model with optimistic buyers directly postulates such a belief, we might say that the optimism de facto captures the obstinate type behaviors.

6.3 Delay in exercising outside options

Unraveling in BP takes off when the buyer types whose net-value is lowest find it *strictly* more profitable to exercise the outside option *immediately* than to continue the negotiation. Thus, it is not difficult to imagine that positive selection may break down if (i) the buyer's outside option is not always available at the very beginning of the negotiation, or (ii) the lowest buyer types' outside option fails to yield a strictly better payoff than the payoff from the continuation of the negotiation in the first place.

Hwang and Li (2017) show that the Coase conjecture is restored if there is friction in the arrival of the buyer's outside option, and hence it becomes available only after a random time. The friction in the arrival of the outside option "physically" delays low buyer types' exit from the market in early stages of bargaining, and hence erodes the positive selection. Lomys (2020) considers a bargaining problem in which both parties are initially uninformed of the existence of the buyer's outside option as well as its arrival time. He discusses how the equilibrium outcome (e.g., whether the efficient outcome occurs) varies with the nature of the outside option.

On the other hand, BP (Section I.B) and Chang and Lee (2022) discuss that the seller may fail to achieve the full commitment profit if some low buyer types only hold zero outside options. Such buyer types are indifferent between different timings of exercising their outside options, and multiple equilibria arise depending on how they break the tie. If they break the tie by delaying the exercise of the outside option, negative selection dominates positive selection, and thus, the Coase conjecture is restored. On the other hand, if they break the tie by exercising the outside option immediately, positive selection dominates and the BP's equilibrium arises.[39]

[39] The main focus of Chang and Lee (2022) is on when and how the seller can achieve the full commitment profit level, as identified by Chang (2021), as an outcome of a perfect Bayesian equilibrium without any commitment device.

6.4 Information disclosure and its inference

In the context of information disclosure where a privately-informed party discloses some information and an uninformed party infers the type of the informed party, an inductive process of information unraveling leads to the full information disclosure (Grossman, 1981; Milgrom, 1981). The key driver of this result is skepticism, i.e., the uninformed party attributes any incomplete disclosure to the informed party concealing unfavorable information. The full information disclosure based on skepticism is hardly supported by empirical and experimental data. For example, some film studios withhold movies from critics before their release, and some moviegoers do not infer low quality from such cold opening (Brown et al., 2012). In the experimental setting, information receivers are insufficiently skeptical about non-disclosed information (Jin et al., 2021). This implies that the informed party may be able to exploit the insufficient inference of the uninformed party.

Unlike the information disclosure environment, there is no explicit device to disclose private information in our environment. Instead, observing the buyer's decision of accepting/rejecting an offer or taking an outside option allows the seller to infer and update her belief, allowing for a similar unraveling to happen in our environment. More precisely, the sellers should have inferred that those who take the outside option in the first round must be the lowest type buyers such that the remaining demand pool in the second round consists of higher types. However, some sellers might not have fully inferred it perhaps due to similar reasons for why the uninformed party in the information disclosure environment does not fully update the belief about non-disclosed information. Optimistic buyers may be one way to parsimoniously capture those who try to exploit the insufficient inference of the seller.

6.5 Reciprocity

Theories of reciprocity (Fehr and Gächter, 2000) also provide an alternative explanation about why some low-type buyers do not exercise the outside option and delay the negotiation in the first round. If a buyer refrains from exercising the outside option, it may create an opportunity to yield a positive payoff for the seller at a risk of yielding herself zero payoff. The buyers may intentionally make this seemingly self-destructive decision, with hoping that the seller will charge a low price reciprocally, which eventually results in a higher payoff for the buyer (Rabin, 1993; Dufwenberg and Kirchsteiger, 2004; Falk and Fischbacher, 2006). However, empirically verifying this narrative is out of the scope of this study as the buyer's action of not exercising the outside option may not be necessarily interpreted as the buyer's intention of appealing reciprocity.

6.6 Failure of rationality

Perhaps a more direct way of explaining the departure of experimental observations from the theoretical prediction is to consider the failure of rationality. In our experimental setup, if the seller worries that low-type buyers, who are supposed to immediately take the outside option, suboptimally continue the negotiation, the key mechanism of Board and Pycia (2014) would not work. The positive selection cannot take off as strongly as theory predicts at the start of unraveling, and hence the stark contrast in the presence/absence of an outside option may disappear.

We explore this idea formally in Online Appendix E. We consider the ϵ-PBEs such that the buyer may choose to delay sub-optimally, provided that the payoff loss relative to the best response is less than ϵ. It turns out that, even in the absence of any optimistic buyer types, such ϵ-irrationality results in the outcome qualitatively identical to the quasi-Coasean equilibrium studied in Section D. The basic intuition is essentially identical to the one for Proposition D.2. The seller's belief that some buyers may sub-optimally remain in the negotiation undermines positive selection among buyer types.[40]

6.7 Failure of higher-order rationality

So far, we have focused on discussing why positive selection may fail in the presence of an outside option. However, our data from the OutNo treatment also significantly deviated from the Coasean equilibrium. Thus, we require an explanation not only for the failure of positive selection when an outside option is present but also for the failure of negative selection in its absence. Given the qualitative similarities in the behaviors of our subjects observed in both treatments, one might conjecture that a common underlying force could explain both phenomena.

We argue that a lack of higher-order rationality is at the core of both phenomena. First, let us consider the situation without an outside option. In general, higher buyer types have a stronger incentive to accept the seller's offer than lower buyer types, resulting in weaker demand over time. Consequently, sellers tend to offer lower prices in later periods. This dynamic may lead even higher buyer types to reject the seller's initial offer, anticipating lower prices in the future (first-order demand withholding). To encourage any trade with higher buyer types in the initial period, the seller must therefore lower initial offer. This suggests that the concepts of demand withholding and the Coase conjecture require players to engage in higher-order reasoning. In our companion paper (Chang et al., 2024), we tested a similar (though not identical) form of higher-order reasoning in a bargaining environment with an outside option and found that the uninformed party failed to conduct such reasoning correctly. While this finding is only suggestive, we believe that a similar failure of higher-order reasoning also accounts for the deviation of our experimental results from the equilibrium prediction in the case of no outside option.

[40] Chang et al. (2024) provide experimental evidence that the majority of participants fail to reason according to positive selection.

7 Conclusion

In this paper, we experimentally consider positive and negative selection in an infinite-horizon bargaining with one-sided, incomplete information. Consistent with theory, we find that bargaining lasts longer when an outside option is unavailable than otherwise. Inconsistent with theory, however, a substantial proportion of price offers are rejected even when the outside option is available. Inconsistent with theory, the seller's initial price offer is strictly higher in the absence of an outside option than in the presence. Similarly, the seller's average profit is larger when the outside option is unavailable. The reported beliefs from the sellers after a price offer is rejected further confirm that positive selection is not fully executed in the lab when the outside option is available, while we find partial evidence of negative selection when the outside option is unavailable.

Our findings provide the first evidence in the literature suggesting that the distributional effects of an outside option available to buyers may not be as pronounced in real-world scenarios as theoretical models predict. Nonetheless, our results imply that the existence of an outside option in a market can be beneficial, as it facilitates quicker agreements between bargainers. This raises the question: will this conclusion hold in real-world markets? We strongly believe it is crucial to validate our main findings obtained in a controlled laboratory setting within a field context to fully understand the impact of outside options on market participants. However, given that our primary focus involves private information on one side of the market, obtaining a suitable data set in the field to address this question poses significant challenges. We leave this for future research.

Finally, we address the external validity of our experiment. A potential criticism is that Coasian bargaining may be strategically complex due to the unraveling process it entails. One might argue that only experts can effectively adhere to equilibrium behavior, while university students may not accurately represent the appropriate population for this context. However, we question whether market participants, in the real-world scenarios of bargaining, durable-good monopolies, and dynamic screening problems, can be deemed "experts" who, on average, possess greater strategic sophistication. Although theoretical models based on full rationality provide a clear depiction of behavior in markets characterized by one-sided private information, we contend that this portrayal does not fully capture the complexities of real-world market dynamics. Our experimental analysis offers an alternative perspective that contrasts with traditional theoretical frameworks, suggesting that the reality likely exists somewhere between these two extremes.

References

Abreu, Dilip and Faruk Gul, "Bargaining and Reputation," *Econometrica*, 2000, *68* (1), 85–117.

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

_ , 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*, Elsevier, 2002, pp. 1897–1945.

Board, Simon and Marek Pycia, "Outside options and the failure of the Coase conjecture," *American Economic Review*, 2014, *104* (2), 656–71.

Brown, Alexander L., Colin F. Camerer, and Dan Lovallo, "To Review or Not to Review? Limited Strategic Thinking at the Movie Box Office," *American Economic Journal: Microeconomics*, May 2012, *4* (2), 1–26.

Burchardi, Konrad B. and Stefan P. Penczynski, "Out of your mind: Eliciting individual reasoning in one shot games," *Games and Economic Behavior*, 2014, *84*, 39–57.

Burdea, Valeria and Jonathan Woon, "Online belief elicitation methods," *Journal of Economic Psychology*, 2022, *90*, 102496.

Cason, Timothy N. and Stanley S. Reynolds, "Bounded rationality in laboratory bargaining with asymmetric information," *Economic Theory*, 2005, *25* (3), 553–574.

_ and Tridib Sharma, "Durable goods, Coasian dynamics, and uncertainty: Theory and experiments," *Journal of Political Economy*, 2001, *109* (6), 1311–1354.

Catonini, Emiliano, "The failure of Coase conjecture with outside options: a rationalizability approach," 2022. Working paper.

Chang, Dongkyu, "Optimal Sales Mechanism with Outside Options," *Journal of Economic Theory*, 2021, *195*, 105279.

_ and Jong Jae Lee, "Price Skimming: Commitment and Delay in Bargaining with Outside Option," *Journal of Economic Theory*, 2022, *205*, 105528.

_ , Duk Gyoo Kim, and Wooyoung Lim, "Unveiling the Failure of Positive Selection," 2024. Working paper.

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*, 88–97.

Cho, In-Koo, "Stationarity, Rationalizability and Bargaining," *The Review of Economic Studies*, 04 1994, *61* (2), 357–374.

Coase, Ronald H., "Durability and Monopoly," *The Journal of Law and Economics*, 1972, *15* (1), 143–149.

Compte, Olivier and Philippe Jehiel, "On the Role of Outside Options in Bargaining with Obstinate Parties," *Econometrica*, 2002, *70* (4), 1477–1517.

Danz, David, Lise Vesterlund, and Alistair J. Wilson, "Belief Elicitation and Behavioral Incentive Compatibility," *American Economic Review*, September 2022, *112* (9), 2851–2883.

Deneckere, Raymond J. and Meng-Yu Liang, "Bargaining with Interdependent Values," *Econometrica*, 2006, *74* (5), 1309–1364.

Dufwenberg, Martin and Georg Kirchsteiger, "A Theory of Sequential Reciprocity," *Games and Economic Behavior*, 2004, *47* (2), 268–298.

Embrey, Matthew, Friederike Mengel, and Ronald Peeters, "Eliciting strategies in indefinitely repeated games of strategic substitutes and complements," Working Paper Series 0317, Department of Economics, University of Sussex Business School January 2017.

Evans, Robert, "Sequential Bargaining with Correlated Values," *Review of Economic Studies*, 1989, *56*, 499–510.

Falk, Armin and Urs Fischbacher, "A Theory of Reciprocity," *Games and Economic Behavior*, 2006, *54* (2), 293–315.

Fanning, Jack, "Outside options, reputations, and the partial success of the Coase conjecture," 2023. Working paper.

_ **and Andrew Kloosterman**, "A simple experimental test of the Coase conjecture: fairness in dynamic bargaining," *The RAND Journal of Economics*, 2022, *53* (1), 138–165.

Fehr, Ernst and Klaus M. Schmidt, "A theory of fairness, competition, and cooperation," *The Quarterly Journal of Economics*, 1999, *114* (3), 817–868.

_ **and Simon Gächter**, "Fairness and Retaliation: The Economics of Reciprocity," *The Journal of Economic Perspectives*, 2000, *14* (3), 159–181.

Friedenberg, Amanda, "Bargaining Under Strategic Uncertainty: The Role of Second-Order Optimism," *Econometrica*, 2019, *87* (6), 1835–1865.

Fudenberg, Drew, David Levine, and Jean Tirole, "Infinite Horizon Models of Bargaining with One-sided Incomplete Information," in A. E. Roth, ed., *Bargaining with Incomplete Information*, London/New York: Cambridge University Press, 1985, pp. 73–98.

Grossman, Sanford J., "The Informational Role of Warranties and Private Disclosure about Product Quality," *The Journal of Law & Economics*, 1981, *24* (3), 461–483.

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

Güth, Werner, Peter Ockenfels, and Klaus Ritzberger, "On durable goods monopolies an experimental study of intrapersonal price competition and price discrimination over time," *Journal of Economic Psychology*, 1995, *16* (2), 247–274.

_ , Sabine Kröger, and Hans-Theo Normann, "Durable-Goods Monopoly with Privately Known Impatience: A Theoretical and Experimental Study," *Economic Inquiry*, 2004, *42* (3), 413–424.

Halevy, Yoram, Johannes Hoelzemann, and Terri Kneeland, "Magic Mirror on the Wall, Who Is the Smartest One of All?," 2021. Working paper.

Hwang, Ilwoo and Fei Li, "Transparency of outside options in bargaining," *Journal of Economic Theory*, 2017, *167*, 116–147.

Jin, Ginger Zhe, Michael Luca, and Daniel Martin, "Is No News (Perceived As) Bad News? An Experimental Investigation of Information Disclosure," *American Economic Journal: Microeconomics*, May 2021, *13* (2), 141–173.

Landsberger, Michael and Isaac Meilijson, "Intertemporal Price Discrimination and Sales Strategy under Incomplete Information," *The RAND Journal of Economics*, 1985, *16* (3), 424–430.

Li, Duozhe and Yat Fung Wong, "Optimism and Bargaining Inefficiency," *The B.E. Journal of Theoretical Economics*, 2009, *9* (1), Article 12.

Liu, Qingmin, Konrad Mierendorff, Xianwen Shi, and Weijie Zhong, "Auctions with Limited Commitment," *American Economic Review*, 2019, *109* (3), 876–910.

Lomys, Niccolò, "Learning while Bargaining: Experimentation and Coasean Dynamics," 2020. Working paper.

McAfee, R. Preston and Daniel Vincent, "Sequentially Optimal Auctions," *Games and Economic Behavior*, 1997, *18* (2), 246–276.

Milgrom, Paul R., "Good News and Bad News: Representation Theorems and Applications," *The Bell Journal of Economics*, 1981, *12* (2), 380–391.

Myerson, Roger, *Game Theory: Analysis of Conflict*, Cambridge, M.A.: Harvard University Press, 1991.

Rabin, Matthew, "Incorporating Fairness into Game Theory and Economics," *The American Economic Review*, 1993, *83* (5), 1281–1302.

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.

Reynolds, Stanley S., "Durable-goods monopoly: laboratory market and bargaining experiments," *The RAND Journal of Economics*, 2000, pp. 375–394.

Rochet, Jean-Charles and Lars Stole, "The Economics of Multidimensional Screening," in Mathias Dewatripont, Lars Peter Hansen, and Stephen J. Turnovsky, eds., *Advances in Economics and Econometrics: Theory and Applications, Eight World Congress, Volume I*, Cambridge University Press, 2003, pp. 150–197.

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

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

Tirole, Jean, "From Bottom of the Barrel to Cream of the Crop : Sequential Screening with Positive Selection," *Econometrica*, 2016, *84* (4), 1291–1343.

Vasserman, Shoshana and Muhamet Yildiz, "Pretrial negotiations under optimism," *The RAND Journal of Economics*, 2019, *50* (2), 359–390.

Vincent, Daniel R., "Bargaining with Common Values," *Journal of Economic Theory*, 1989, *48* (1), 47–62.

Yildiz, Muhamet, "Bargaining with Optimism," *Annual Review of Economics*, 2011, *3* (1), 451–478.

Online Appendix for
Positive and Negative Selection in Bargaining

A Sample Experimental Instructions

Welcome to the experiment. Please read these instructions carefully. There will be a quiz around the end of the instructions, to make sure you understand this experiment. The payment you will receive from this experiment depends on your decisions.

Your Role and Match

At the beginning of the experiment, one-half of the participants will be randomly assigned to the role of a **seller** and the other half the role of a **buyer**. Your role will remain fixed throughout the experiment.

The experiment consists of 7 **matches**. At the beginning of each match, one seller participant and one buyer participant are randomly paired. The pair is fixed **within the match**. After each match, participants will be reshuffled to form new pairs. You will not learn the identity of the participant you are paired with, nor will that participant learn your identity—even after the end of the experiment.

In each match, the seller holds an asset. The value of the asset is 0 for the seller. However, the buyer has a strictly positive value of the asset. Let B denote the buyer's value of the asset. At the beginning of a match, a computer randomly and independently draws B between 50 and 400. Every integer in $[50, 400]$ has an equal chance to be drawn. The value B is fixed within each match, and a new B is independently drawn for a new match. **The buyer knows the value B, but the seller does not.**

Your Decisions in Each Match

Each match consists of at least one round of bargaining. In a round, a seller offers a price to sell the asset, and the buyer responds. If the offer is rejected, the match may move on to the next round of bargaining. The details follow.

Your Task as a Buyer: Suppose your role is a buyer. At the beginning of Round 1, you will see the following figure. The horizontal position of the dark blue line represents B, your value of the asset. (B is 330

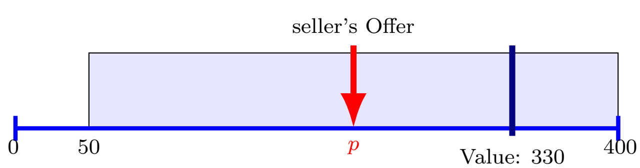

Figure A.1: Buyer's Screen in Each Round

in this example, but your value will vary.) Once the seller in your pair makes a price offer, p, a red vertical arrow will appear on the figure. The position of the red arrow represents p. After that, decide whether to

- **accept** the offer and earn $(B - p)$,

- **reject** it and move on to the next round with an 80% chance, or
- **take an outside option** to earn 50 tokens.

Your task is the same for every round. It is important to understand that when you reject the offer, **there is a 20% chance that the match is terminated, and both of you and the seller in your pair earn 0 tokens.**

Beware that you cannot accept the offered price p if it is strictly greater than your value B, otherwise your payoff becomes negative.

<u>Your Task as a Seller</u>: Suppose your role is a seller. At the beginning of Round 1, you will see the following figure. The blue shaded area between 50 and 400 represents the range of all possible buyer's values.

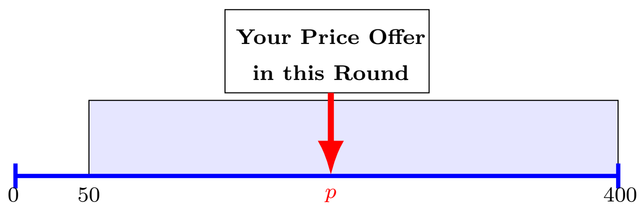

Figure A.2: Seller's Screen: Round 1

Choose your price offer by clicking on the line. A red vertical arrow, whose position represents your price offer, p, will move to the point you click. You are free to choose any point in the range **[0, 400]** for your price offer, and you can adjust it as much as you wish. After that, click the submit button, and wait for the buyer's decision. You expect one of three possible outcomes.

- If the buyer **accepts** the offer, you earn p tokens.
- If the buyer **takes an outside option**, you earn 0 tokens.
- If the buyer **rejects** the offer, then the match moves to Round 2 with an 80% chance. Note that if the match is terminated with a 20% chance, both you and the buyer earn 0 tokens.

If the match moves to Round 2 and beyond, then you will see the following figure below. The red vertical arrow represents your (rejected) previous offer.

Before submitting a new offer price, adjust an orange slider and a purple slider to indicate the updated range of possible values B in your mind. **The reported range will appear in your decision screen but will not be shared with the buyer.** Its sole objective is to help you think about an appropriate price offer. There is nothing to gain by indicating a range that differs from what you actually believe, so please report your belief as accurately as possible. To indicate the range, move the orange and purple sliders. The horizontal positions of the sliders respectively represent the minimum and the maximum of the range in your mind. The minimum can't exceed the maximum.

Note that the buyer's value of the asset (B) and the value of the outside option (50 tokens) will remain the same across rounds within a match.

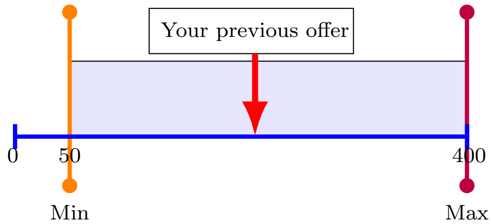

(a) Before

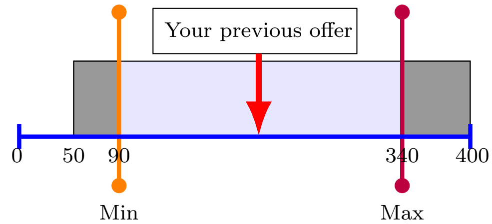

(b) After

Figure A.3: Reporting Beliefs in Round n

Probability of Match Termination

As described, after a buyer rejected a price offer, the match continues to the next round with an 80% chance. Your screen presents a spinning wheel that consists of red area (20%) and green area (80%) as illustrated below. Once you click the "Spin" button, the wheel starts spinning. If the spinning wheel stops at the green area, the match continues. Otherwise, the match terminates. Note that the seller and buyer in the same match always see the same outcome from the wheel for each round.

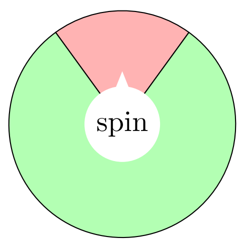

Figure A.4: Spinning Wheel

Information Feedback

- At the end of each **round**, you will know the seller's price offer and the buyer's decision. If the buyer rejects the offer, you will know whether the match is continued to the next round or terminated.
- At the end of each **match**, you will know how many tokens you receive from the match.

Your Monetary Payments

At the end of the experiment, a computer will randomly select one match out of 7 for your payment. Every match has an equal chance to be selected for your payment, so it is in your best interest to take each match equally seriously. Participants will receive the amounts of tokens according to the outcome from the selected match with the exchange rate of 1 token = 1 HKD. Also, every participant will receive a **show-up fee of HKD 40**.

Completion of the Experiment

After the 7th match, the experiment will be over. You will be instructed to fill in the receipt for your payment. The amount you earn will be paid **electronically via the HKUST Autopay System to the bank account you provide to the Student Information System (SIS)**. The Finance Office of HKUST will arrange the auto-payment. An email notification will be sent to your HKUST email address on the pay date under the name of the sender "FOPSAP" (Finance Office Payment System Auto Payment).

Comprehension Check

To ensure your comprehension of the instructions, you will answer four multiple-choice questions. You can proceed only with all correct answers. Afterwards, you will participate in a practice match.

Q1 Suppose you are a seller. Which of the followings is NOT TRUE? (a) I do not know how much the buyer values the asset. (b) If the buyer takes an outside option, I earn 50 tokens. (c) If I offer 200 tokens, and the buyer accepts it, then I earn 200 tokens. (d) If the buyer rejects my offer, then I can make a new offer with an 80% chance.

Q2 Suppose you are a buyer, and the value of the asset is 300. Which of the followings is TRUE? (a) If I accept a price offer of 200 tokens, I earn 200 tokens. (b) If I take an outside option, I earn 250 tokens. (c) If I accept a price offer of 200 tokens, I earn 100 tokens. (d) In Round 2 of this match, the value of the asset will be different from 300.

Q3 Suppose the price offer in Round 1 is rejected. Which of the followings CAN HAPPEN? (a) The match is terminated, and both the seller and the buyer earn 0 tokens. (b) The match is continued forever, even after continuous rejections. (c) The match is terminated, and each participant in the pair earns a half of value B. (d) The match initiates an open chat to negotiate.

Q4 Suppose the first match is done. Which of the followings is TRUE? (a) It is almost sure that I will be paired with the same participant in the first match. (b) I may play another role different from what I did in the first match. (c) The buyer's value of the asset in the second match will be the same as the one in the first match. (d) My previous actions do not affect the value of the asset in the new match.

1 Practice Match and 7 Actual Matches

[After passing the quiz] Thank you for paying attention to the instructions. Before you will play the 7 actual matches, you will have one practice match (Match #0) which is not relevant to your payment. Its objective is to get you familiar with the computer interface and the flow of the decisions in each round of a match. Once the practice match is over, it moves to the actual matches.

B Supplementary Figures

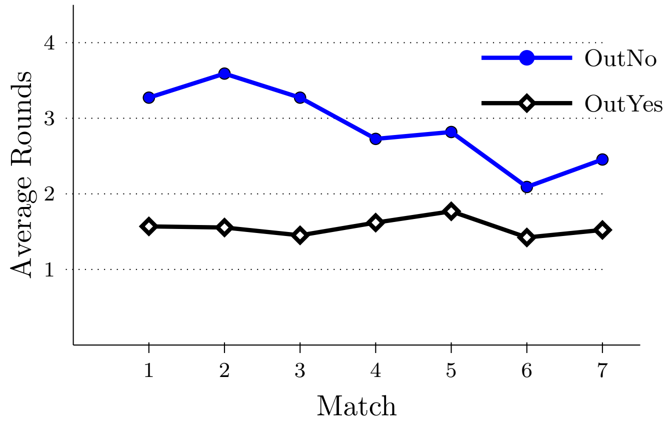

Figure B.1: Average Length of Bargaining across Match

This figure shows the average length of bargaining rounds by match. The bargaining process was ended either in the form or agreement or termination. On average, it takes 2.89 rounds to end the bargaining process in OutNo, while it takes 1.56 rounds in OutYes.

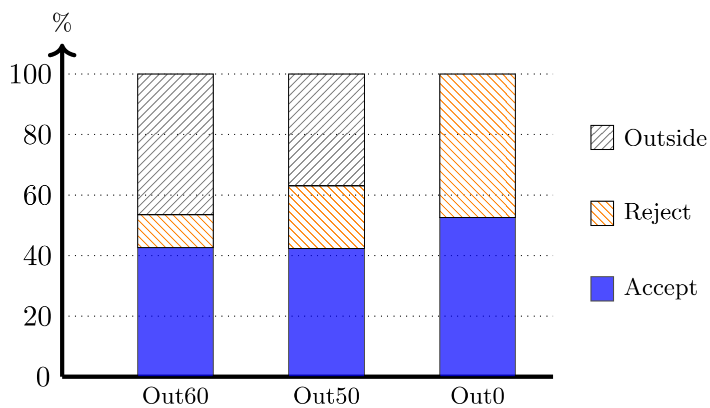

Figure B.2: End-of-Bargaining States by Treatment

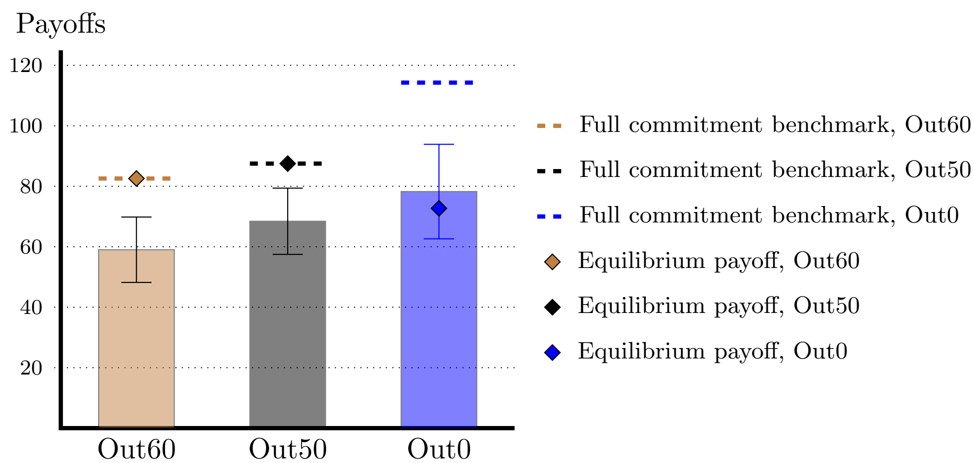

Figure B.3: Seller's Earnings by Treatment

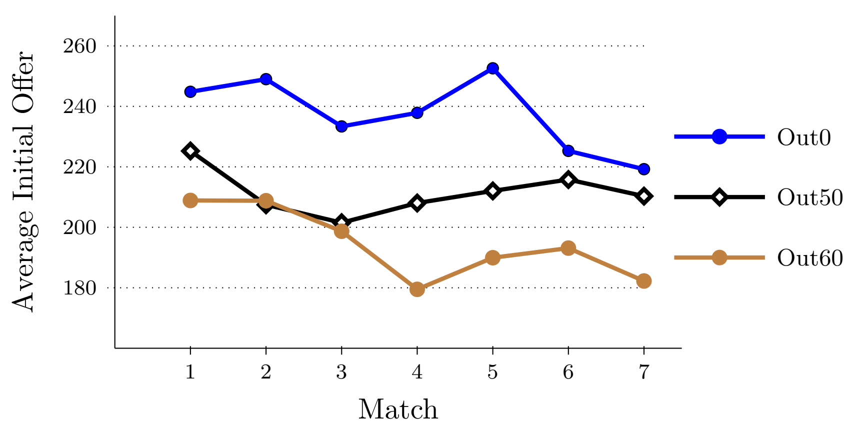

Figure B.4: Round 1 Offer across Match

C Additional Analysis

Consider the case in which the buyer has no outside option (OutNo). Suppose the seller offers $p_1 \geq 0$ in the first round and, following rejection, plans to offer p_2 in the next round. If the buyer correctly anticipates the seller's next-round offer (as in equilibrium, though not necessarily in the lab experiment), the buyer rejects p_1 whenever $v - p_1 < \delta(v - p_2)$, or equivalently,

$$v < p_1 + \frac{\delta}{1-\delta}(p_1 - p_2). \tag{C.1}$$

Here, v denotes the buyer's realized value and δ denotes the continuation probability to the next round. For any given p_1, rejection is therefore more likely when the anticipated price cut $p_1 - p_2$ is larger.

Inequality (C.1) provides only a sufficient condition for rejecting p_1, as it does not account for the possibility that the buyer optimally waits beyond the second round. Under the assumption that buyers correctly anticipate p_2, the ex ante probability that the buyer rejects the seller's initial offer p_1 is bounded below by

$$\Phi(p_1, p_2) := \mathbb{P}\left\{v < p_1 + \frac{\delta}{1-\delta}(p_1 - p_2)\right\} = \min\left\{1, \max\left\{\frac{p_1 + \frac{\delta}{1-\delta}(p_1 - p_2) - 50}{350}, 0\right\}\right\},$$

where the last equality follows because v is uniformly distributed on $[\underline{v}, \overline{v}] = [50, 400]$.

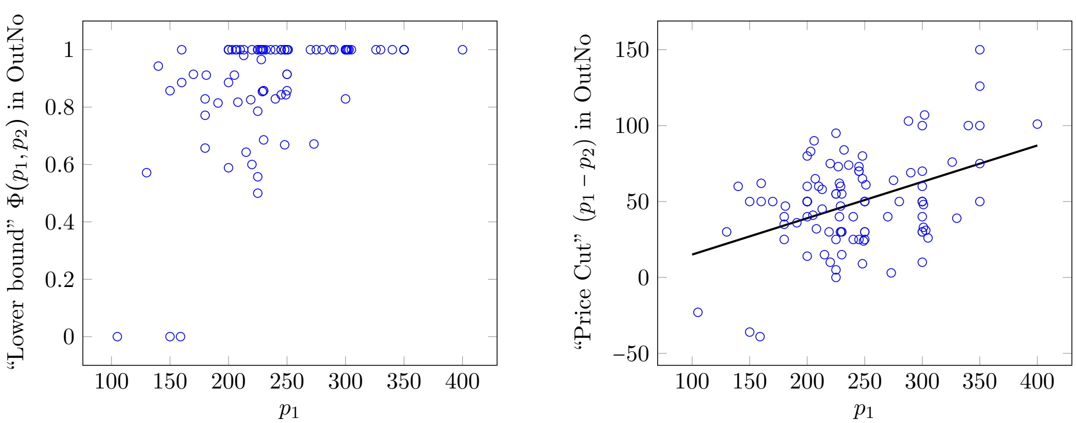

(a) Scatterplot of $\Phi(p_1, p_2)$ on p_1

(b) Scatterplot of $p_1 - p_2$ on p_1

Figure C.1: Lower bound on rejection probability and observed price cuts in OutNo.

Panel (a) plots the lower bound $\Phi(p_1, p_2)$ on the rejection probability against the first-round offer p_1 for all matches that reach the second round. Panel (b) plots the corresponding price cut $p_1 - p_2$ against p_1.

Figure C.1a plots the observed first-round offer p_1 (x-axis) against the lower bound $\Phi(p_1, p_2)$ on the rejection probability (y-axis) for all matches that reach the second round. The bunching at $\Phi(p_1, p_2) = 1$ arises because the right-hand side of (C.1) exceeds $\overline{v}$ in a substantial number of matches. For a large majority of observed (p_1, p_2) pairs, this lower bound exceeds the theoretical prediction (about 0.60), which is consistent with the pervasiveness of rejections observed in Figure 8.[41]

This discrepancy between theory and data can be explained by two factors. First, in many matches the seller's initial offer exceeds the theoretical prediction: the average initial price in OutNo is 237.45, whereas the theory predicts an initial offer of 121.56 in the first round. The observed average initial price of 237.45 exceeds the valuations of more than half of buyer types between $\underline{v} = 50$ and $\overline{v} = 400$, for whom this offer is unacceptable regardless of their anticipation of the seller's next offer. Second, sellers cut their price offers in the second round too aggressively, so that most buyer types, if they correctly anticipate this behavior, would find rejection to be profitable.

To see the second factor more clearly, note that the 6th decile of the uniform distribution over $[50, 400]$ is $v = 260$. Thus, with $p_1 = 237.45$ and $\delta = 0.80$, the lower bound $\Phi(p_1, p_2)$ for the rejection probability falls below the theoretical prediction of 0.60 only if

$$p_1 + \frac{\delta}{1-\delta}(p_1 - p_2) \leq v_{(60)} \quad \Longleftrightarrow \quad 237.45 + 4(p_1 - p_2) \leq 260 \quad \Longleftrightarrow \quad p_1 - p_2 \leq 5.6375,$$

where $v_{(60)} = 260$ is the cutoff type such that $\mathbb{P}\{v \leq v_{(60)}\} = 0.60$. Thus, conditional on an initial offer of $p_1 = 237.45$, the rejection probability equals the theoretical prediction only if the buyer expects a price cut of at most 5.6375. Figure C.1b shows that, in a substantial number of matches with initial prices in a neighborhood of 237.45, the observed price cut substantially exceeds this threshold. This pattern indicates that sellers cut prices too aggressively, thereby leading buyers to reject the initial offer with probability exceeding the theoretical prediction.

[41] For the theoretical prediction for OutNo, see Table 1 and the subsequent discussion.

D Optimism in Bargaining with Outside Option

D.1 Bargaining Model with Optimistic Buyer Types

In this section, we show that the stark contrast between the two cases with and without the buyer's outside option disappears when we introduce the buyer's optimism about the bargaining process. To illustrate this idea, we modify the negotiation game as follows. First, extend the buyer's type space to

$$\Theta := V \times \{\tau_r, \tau_o\},$$

where a finite set $V \subset (0, \infty)$ represents the support of the buyer's valuation of the seller's good as in Section 2.1. We respectively denote the highest and the lowest valuation in V by $\overline{v}$ and $\underline{v}$ and also denote a generic element in Θ by $\theta = (v, \tau)$.

Each buyer type with $\tau = \tau_r$ holds *correctly specified (rational) model* about the seller's behavior in the negotiation: The seller will play as predicted by a standard equilibrium notion (which will be defined shortly). On the other hand, each buyer type with $\tau = \tau_o$ has the following *misspecified model*: In every period,

- the seller offers $p^\dagger \in [0, \infty)$ with probability $\eta \in (0, 1)$, regardless of which equilibrium strategy the seller prepares, and
- the seller plays her equilibrium strategy with probability $1 - \eta \in (0, 1)$.[42]

We focus on the case that $p^\dagger$ is sufficiently small (see Assumption D.2 for the exact condition), so that this misspecified model makes more optimistic predictions about the future bargaining process than the correctly specified one. For simplicity, we assume that both η and $p^\dagger$ are common across all buyer types with $\tau = \tau_o$.

The buyer types with $\tau = \tau_o$ are called *optimistic types* while the buyer types with $\tau = \tau_r$ are called *rational types*. The set of all rational buyer types and the set of all optimistic buyer types are denoted by $\Theta_r = \{(v, \tau) \in \Theta : \tau = \tau_r\}$ and $\Theta_o = \{(v, \tau) \in \Theta : \tau = \tau_o\}$, respectively. The realizations of v and τ are assumed to be stochastically independent. For any finite support V being fixed, let $f_V(v)$ denote the seller's prior probability that the buyer's valuation is v, and let $\phi \in [0, 1]$ denote the prior probability that the buyer is an optimistic type. We write the joint probability mass function as $f_\Theta(\theta) = f_\Theta(v, \tau)$, hence $f_\Theta(v, \tau_o) = f_V(v)\phi$ and $f_\Theta(v, \tau_r) = f_V(v)(1 - \phi)$.

The negotiation procedure is identical to that presented in Section 2.1. The buyer still has a type-independent outside option. In this section, we focus on the case $\underline{v} \geq w > 0$ without loss. Each optimistic buyer type $\theta = (v, \tau_o)$ (mistakenly) believes that he can always guarantee a continuation payoff

$$\sum_{k=1}^{\infty} e^{-rk\Delta}(1 - \eta)^{k-1}\eta(v - p^\dagger) = \frac{e^{-r\Delta}\eta}{1 - e^{-r\Delta}(1 - \eta)}(v - p^\dagger)$$

[42]In other words, all buyer types with $\tau = \tau_o$ incorrectly believe the following. At the beginning of each period $n \geq 0$, nature draws the number u from a uniform random variable $U[0, 1]$. If $u \leq \eta$ the seller is *forced* to offer $p_n = p^\dagger$. If $u > \eta$, the seller may choose any price p_n as she wishes.

by waiting for the seller to offer $p^\dagger$. Hence, we may define the net-value of each optimistic buyer type as

$$u(\theta) := v - \max\left\{\frac{e^{-r\Delta}\eta}{1-e^{-r\Delta}(1-\eta)}(v-p^\dagger), w\right\} \quad \forall\theta = (v,\tau_o) \in \Theta_o.$$

We continue to define the net-value of each rational buyer type as (2.1) in Section 2.1.

Provided that $\phi \geq 0$ is sufficiently close to zero, the presence of the optimistic buyer type does not alter the full-commitment benchmark qualitatively. With full commitment power, it is optimal for the seller to ignore the negligible portion of the optimistic type and to insist on a single price that maximizes the profit from the rational buyer type, which is given as follows:

$$\Pi_r(p) := p \sum_{\theta\in\Theta_r : u(\theta)\geq p} f_\Theta(\theta).$$

We continue to denote the benchmark price that maximizes $\Pi_r(p)$ as p_w^* and the seller's full-commitment profit level (i.e., the payoff from insisting on p_w^*) as Π_w^*, respectively. We also maintain the following two assumptions.

Assumption D.1. *There is $u^* \in \{u(v,\tau_r) : v \in V\}$ such that the following two conditions hold:*

(i) Π_r *is single-peaked: For any u', u'' with $u' < u'' \leq u^*$, $\Pi_r(u') < \Pi_r(u'')$. For any u', u'' with $u^* \leq u' < u''$, $\Pi_r(u') > \Pi_r(u'')$.*

(ii) $u(\underline{v},\tau_r) < u^*$.

Assumption D.2. $0 \leq p^\dagger < p_w^*$.

Assumption D.1-(i) requires Π_r to be single-peaked, and hence p_w^* is well-defined whenever ϕ is sufficiently small. Assumption D.1-(ii) allows us to avoid the case that the full-commitment benchmark outcome induces trading with all rational buyer types. Assumption D.2 requires that the optimistic buyer type finds $p^\dagger$ strictly more favorable than the full-commitment benchmark price p_w^*. Note that, once other parameters such as $p^\dagger$, w, η, Δ, and r being fixed, p_w^* and Π_r are determined by the seller's prior (V, f_V, ϕ). In this sense, Assumptions D.1 and D.2 are the conditions imposed on the seller's prior belief.

D.2 Equilibrium Concept

We employ *a version of ϵ-PBE* as the equilibrium concept. A formal definition of ϵ-PBE requires additional notations. First of all, we inherit all the notations for equilibrium strategies and beliefs (i.e., σ^B, σ^S, and f^S) from Section 2.1. For any assessment $\sigma = (\sigma^B, \sigma^S, f^S)$, history $h_n \in H$, and offer $p \geq 0$, define $V_S^\sigma(p; h_n)$ as the seller's continuation payoff from a one-shot deviation to p at h_n. Also, define

$$P^D(h_n;\sigma) := \{p \geq 0 : \sigma^B(p|h_n,\theta)[D] > 0 \text{ for some } \theta \in \text{supp}(f^S(\theta|h_n))\}$$

as the collection of offers in response to which some remaining buyer types will choose to delay. We call the offers in $P^D(h_n;\sigma)$ *delay-inducing* (at history h_n).

Next, a seller's strategy σ^S is called *ϵ-best at $h_n \in H$*, where $\epsilon \geq 0$, if and only if the following inequality holds:

$$V_S^\sigma(p'; h_n) \leq V_S^\sigma(p; h_n) + \epsilon \mathbb{1}\{p \text{ or } p' \in P^D(h_n;\sigma)\} \quad \forall p \in \text{supp}(\sigma^S(h_n)) \text{ and } p' \geq 0. \tag{B}$$

A strategy σ^S is called ϵ-best if it is ϵ-best at any histories. Note that the condition (B) allows a margin of error *only at a subclass of histories*, and in this sense, our notion of ϵ-best strategy is closer to a fully optimal strategy than alternatives that allow a small error *at any history.* Precisely, it allows a small margin of improving the seller's continuation payoff for the case that either p or p' is delay-inducing; the condition coincides with the standard best-response condition for all other cases. This asymmetric treatment to the delay-inducing price offers is motivated by the observation that a precise assessment of $V_S^\sigma(p; h_n)$ requires a more demanding cognitive capacity (e.g., a prediction of the buyer's future behaviors) when the seller expects further delay and haggling.

Now we are ready to define our notion of ϵ-PBE. The definition of ϵ-PBE is identical to the definition of the standard PBE except that the seller is only required to play an ϵ-best strategy. In any ϵ-PBE, as in the standard PBE, all buyer types exactly play the optimal strategies after all histories.[43] Finally, the seller's posterior belief is also rationally obtained by Bayes rule whenever possible. For any $\epsilon > 0$, $\Delta > 0$, and the seller's prior belief (V, f_V, ϕ), let $\mathcal{E}(\epsilon|\Delta, V, f_V, \phi)$ denote the set of all ϵ-PBEs. Note that our notion of ϵ-PBE includes the standard PBE as a special case.[44]

D.3 Equilibrium Analysis

In this subsection, we investigate the ϵ-PBEs under Assumptions D.1 and D.2. Fix $p^\dagger \geq 0$, $\eta \in [0,1]$, $w > 0$, and $r > 0$ throughout the section, and let $\mathcal{F}$ denote the collection of the seller's prior beliefs, generically denoted by (V, f_V, ϕ), such that (i) $|V| < \infty$, (ii) $\phi \in [0,1]$, and (iii) Assumptions D.1 and D.2 hold. We also maintain the assumption $\underline{v} = \min V \geq w$ without any loss (see footnote 11).

We first consider the case with $\phi = 0$. Note that the model coincides with the one discussed in Section 2 in this case (i.e., the model with all the buyer types being rational). The following proposition shows that all the key theoretical predictions continue to hold. Recall that h_0 generically denotes the null history of the game.

Proposition D.1. *Fix $\epsilon > 0$, $\Delta > 0$ and $(V, f_V, \phi) \in \mathcal{F}$ such that $\phi = 0$. Then, for any f, there is a unique ϵ-PBE in $\mathcal{E}(\epsilon|\Delta, V, f_V, \phi)$. Furthermore, in this unique ϵ-PBE:*

(i) The seller earns the full-commitment benchmark profit Π_w^*.

(ii) There is a cutoff $v^* \in V$ such that all buyer types with $v \geq v^*$ trade, and all other buyer types exercise the outside option in period 0.

[43] More precisely, the optimistic types choose the optimal strategies (i.e., the best response to the seller's strategy) under the misspecified model as discussed above. All rational types choose the optimal strategies under the correct model.

[44] ϵ-PBE is employed to resolve a technical issue due to the multi-dimensional type space. If all buyer types are rational (hence, the type space is one-dimensional) as in Section 2.1, the order of purchase is monotonic in buyer's type regardless of the seller's strategy in the sense that a high-value buyer type *always* accept the seller's offer earlier than lower buyer types. In contrast, with the two-dimensional type space $\Theta = V \times \{\tau_r, \tau_o\}$, the order of purchase among buyer types is endogenous to the seller's strategy. The full equilibrium analysis thus requires to inspect the effect of each deviation by the seller on the order of purchase, which imposes substantial technical challenges. The use of ϵ-PBE as the equilibrium concept eases this issue by allowing us to disregard the optimistic type's purchase as its impact on the seller's profit is limited when ϕ is small. Also, note that a similar technical issue arises in the multi-dimensional mechanism design literature where the set of binding incentive-compatibility constraints depends on the mechanism's rule. See, e.g., Rochet and Stole (2003).

(iii) No delay occurs after any offer $p_0 \in [0, \infty)$ by the seller in period 0 (i.e., $P^D(h_0; \sigma) = \varnothing$ in the unique ϵ-PBE σ).

Proof. The proof is essentially identical to the proof of Proposition 1 in BP, which we sketch here.

Step I: Fix any ϵ-PBE $\sigma = (\sigma^B, \sigma^S, f^S)$. The seller never offers $p_n < \underline{u}^\sigma(h_n) := \min\{u(\theta) : \theta \in \text{supp} f^S(h_n)\}$ at any h_n. Suppose for contradiction $G := \sup_{h_n \in H} \underline{u}^\sigma(h_n) - \underline{p}^\sigma(h_n) > 0$, where $\underline{p}^\sigma(h_n) := \inf(\text{supp} f^S(h_n))$. Choose h_n and $0 < \beta < (1 + e^{-r\Delta})/2$ such that $\underline{u}^\sigma(h_n) - \underline{p}^\sigma(h_n) > (1 - \beta)G$. By the argument identical to Lemma 1 in BP, any deviation to $p \in [\underline{p}^\sigma(h_n), \underline{p}^\sigma(h_n) + \beta G)$ would induce an immediate trade with the buyer, yielding $V_S^\sigma(p; h_n) = p$ as the seller's final payoff. Hence, $\sigma^S(h_n)$ never chooses any price in $[\underline{p}^\sigma(h_n), \underline{p}^\sigma(h_n) + \beta G)$, contradicting the supposition $G > 0$.

Step II: By the argument identical to Proposition 1 in BP, the observation in Step I implies the following: in any ϵ-PBE, any offer p_0 at h_0 never induces a delay. Hence, the seller will charge $p_0 = p_w^*$ to achieve the full-commitment benchmark profit in any ϵ-PBE. ☐

Now we turn our attention to the case with $\phi > 0$. Our main result shows that, in the presence of optimistic buyer types, there is an ϵ-PBE such that the seller practices the inter-temporal price discrimination as in the case of no outside option (Section 2.2). We refer to such equilibrium as *quasi-Coasean equilibrium.*

Definition D.1. *An ϵ-PBE is called a quasi-Coasean equilibrium if the following outcome is induced on its path:*

- *Delay and Inter-temporal Pricing: The negotiation takes multiple periods with positive probability and p_n declines over time on the equilibrium path.*
- *There is $v^* \in V$ such that the following holds on the equilibrium path:*
  - *(i) Positive selection: any rational buyer type $\theta = (v, \tau_r)$ such that $v < v^*$ exercises the outside option immediately, and*
  - *(ii) Negative selection: any rational buyer type $\theta = (v, \tau_r)$ such that $v \geq v^*$ trades with the seller (possibly after a delay). Among these buyer types, one with a higher valuation v trades earlier than others.*

The observed negotiation process in a quasi-Coasean equilibrium is a mixture of the two equilibrium plays described in Propositions 3 and 4, featuring both positive and negative selection as well as inter-temporal price discrimination and exclusion. A chunk of buyer types in $\{(v, \tau_r) : v < v^*\}$ exercise the outside option in period zero. In addition, the seller practices inter-temporal price discrimination for the remaining buyer types in $\{(v, \tau_r) : v \geq v^*\}$, where high buyer types trade earlier among these buyer types. Note that the PBE discussed in Section 2.2 (for the case $w = -\infty$) satisfies all the conditions in the definition of quasi-Coasean equilibrium. On the other hand, quasi-Coasean equilibria differ from the equilibrium discussed in Section 2.3 (for the case $w > 0$) in which no price discrimination occurs between the non-excluded buyer types. Furthermore, any quasi-Coasean equilibrium fails to achieve the full-commitment benchmark outcome.

The next proposition identifies the condition under which a quasi-Coasean equilibrium exists for the case with $w > 0$. For any $(V, f_V, \phi) \in \mathcal{F}$, let $\overline{f}_V := \max\{f_V(v) : v \in V\} > 0$ denote the maximum value that the probability mass function $f_V(v)$ may take.

Proposition D.2. *Fix $\epsilon > 0$. Then, there is $\bar{\phi}(\epsilon) \in (0,1)$ such that the following holds for any (V, f_V, ϕ) with $0 \le \overline{f}_V p^\dagger < \epsilon/2$ and $0 < \phi < \bar{\phi}(\epsilon)$: $\mathcal{E}(\epsilon|\Delta, V, f_V, \phi)$ includes a quasi-Coasean equilibrium whenever Δ is sufficiently small.*

See Section 6.1 for the intuition for Proposition D.2.

D.4 Proof of Proposition D.2

Throughout the proof, denote the common discounting factor by $\delta \equiv e^{-r\Delta} \in (0,1)$. Define

$$v_c := \min\{v \in V : \delta\eta(v - p^\dagger - w) > (1-\delta)w\}, \quad u_c := v_c - w, \quad \text{and} \quad p_c := \frac{(1-\delta)v_c + \delta\eta p^\dagger}{1 - \delta + \delta\eta}.$$

Assumptions D.1 and D.2 guarantee that v_c is well-defined. In particular, v_c is well-defined as V is finite. Also,

$$v^\dagger := \min\{v \in V : v - p^\dagger - w > 0\}.$$

D.4.1 Step I: Preliminary Observations

First, we make several preliminary observations on v_c and p_c. Note that v_c converges to $v^\dagger$ as $\delta \to 1$. Furthermore, given that $|V| < \infty$, there is a cutoff $\bar{\delta} \in (0,1)$ such that $v_c = v^\dagger$ and $v_c - p^\dagger - w = v^\dagger - p^\dagger - w > \frac{1-\delta}{\delta\eta} w > 0$ whenever $\delta \in (\bar{\delta}, 1)$. Finally, $p_c > p^\dagger$ for all $\delta \in (0,1)$ and $p_c \downarrow p^\dagger$ as $\delta \to 1$.

Lemma 1. *There is $\delta^* \in (0,1)$ such that the following properties hold whenever $\delta \in (\delta^*, 1)$.*

(i) $v - w > p_c$ *for* $v \ge v_c$.

(ii) $v_c - p_c = \delta(1-\eta)(v_c - p_c) + \delta\eta(v_c - p^\dagger) = \frac{\delta\eta}{1-\delta+\delta\eta}(v_c - p^\dagger) > w$.

(iii) $v - p_c > \delta(1-\eta)(v - p_c) + \delta\eta(v - p^\dagger) > \frac{\delta\eta}{1-\delta(1-\eta)}(v - p^\dagger) > w$ *for any* $v > v_c$.

(iv) $v - p_c < \delta(1-\eta)(v - p_c) + \delta\eta(v - p^\dagger) < \frac{\delta\eta}{1-\delta(1-\eta)}(v - p^\dagger) \le w$ *for any* $v < v_c$.

Proof. Throughout the proof, we assume without loss $v_c = v^\dagger$ and $v^\dagger - p^\dagger - w > \frac{1-\delta}{\delta\eta} w > 0$. The part (i) directly follows the following observation:

$$v - w - p_c = \frac{\delta\eta(v^\dagger - p^\dagger - w) - (1-\delta)w}{1 - \delta(1-\eta)} + v - v_c > v - v_c \ge 0 \quad \forall v \ge v_c.$$

The part (ii) follows the definition of v_c and p_c. To show (iii), recall from (ii) that the following equation exactly holds at $v = v_c$:

$$v - p_c = \delta(1-\eta)(v - p_c) + \delta\eta(v - p^\dagger) = \frac{\delta\eta}{1 - \delta(1-\eta)}(v - p^\dagger). \tag{D.1}$$

The comparison between the coefficients of v on each side of (D.1) reveals that the first two inequalities in (iii) hold strictly whenever $v > v_c$. The last inequality $\frac{\delta\eta}{1-\delta(1-\eta)}(v - p^\dagger) > w$ in (iii) follows the definition of v_c. The proof of (iv) is similar to the proof of (iii). □

Next, we make observation on the seller's static profit:

$$\Pi(p) := \sum_{\theta\in\Theta:u(\theta)\geq p} pf_\Theta(\theta) = (1-\phi)\Pi_r(p) + \phi \sum_{\theta\in\Theta_o:u(\theta)\geq p} pf_\Theta(\theta).$$

$\Pi(p)$ is the seller's payoff from insisting on $p_n = p$ in all periods. In the next lemma, u^* denotes the argument maximizing Π_r (see Assumption D.1). Note that $\Pi(p)$ is also maximized at $p = u^*$ (hence, $p_w^* = u^*$), provided that ϕ is sufficiently small; we choose $\bar{\phi}(\epsilon)$ as a cutoff such that u^* indeed coincides with p_w^* for all $0 < \phi < \bar{\phi}(\epsilon)$.

Lemma 2. *Suppose $0 \leq \overline{f}_V p^\dagger < \epsilon/2$. There are $\bar{\phi}(\epsilon) \in (0,1)$ and $\bar{\delta}(\epsilon) \in (0,1)$ such that the following holds true whenever $0 < \phi < \bar{\phi}(\epsilon)$ and $\bar{\delta}(\epsilon) < \delta < 1$: For any $\underline{v} - w \leq p' \leq p'' \leq p_c$, $\Pi(p') \leq \Pi(p'') + \epsilon/2$.*

Proof. $\{u(\theta) : \theta \in \Theta_r\}$ is a finite set, and hence we may enumerate its elements by $u_1 < u_2 < \ldots < u_{|V|}$. Note that $\Pi_r(p)$ is linearly increasing in p over each interval (u_k, u_{k+1}) and has a downward jump at each u_k. Also, let M be the integer such that $p^\dagger \in [u_M, u_{M+1})$. Recall that p_c decreases in δ and converges to $p^\dagger$ as $\delta \to 1$. Hence, we may choose $\bar{\delta}(\epsilon)$ such that $p_c \in (u_M, u_{M+1})$ whenever $\bar{\delta}(\epsilon) < \delta < 1$, which we will indeed suppose throughout the proof.

Fix any p' and p'' such that $\underline{v} - w \leq p' \leq p'' \leq p_c$. The inequality $\Pi(p') \leq \Pi(p'') + \frac{\epsilon}{2}$ trivially holds if both p' and $p'' \in (u_k, u_{k+1}]$ for some k. Hence, we may assume without loss that there is u_k such that $p' \leq u_k < p''$; if there are multiple such u_k's, pick the largest one among them. Then, $p' \leq u_k < p'' \leq p_c < u_{M+1}$, and therefore,

$$\Pi_r(p'') - \Pi_r(p') \geq \Pi_r(u_k+) - \Pi_r(u_k) = -u_k f_V(v_k) \geq -p_c \overline{f}_V,$$

where $\Pi_r(u_k+)$ stands for the right-limit of Π_r at $p = u_k$ and $v_k = u_k + w$. Hence,

$$\Pi_r(p'') - \Pi_r(p') > -p_c\overline{f}_V = -\left[p^\dagger - \frac{1-\delta}{1-\delta+\delta\eta}(v_c - p^\dagger)\right]\overline{f}_V > -\frac{\epsilon}{2} + \overline{f}_V \frac{1-\delta}{1-\delta+\delta\eta}(v_c - p^\dagger) > -\frac{\epsilon}{2}.$$

Finally,

$$\Pi(p'') - \Pi(p') \geq (1-\phi)\left[\Pi_r(p'') - \Pi_r(p')\right] - \phi\overline{v} > (1-\phi)\left[-\frac{\epsilon}{2} + \overline{f}_V \frac{1-\delta}{1-\delta+\delta\eta}(v_c - p^\dagger)\right] - \phi\overline{v}.$$

This bound for $\Pi(p'') - \Pi(p')$ is independent of p' and p''. Hence, there is $\overline{\phi}(\epsilon)$ such that, whenever $\overline{\phi}(\epsilon) < \phi < 1$, $\Pi(p'') - \Pi(p') \geq -\epsilon/2$ for all $p' \leq p'' \leq p_c$. □

D.4.2 Step II: Quasi-Coasean Equilibrium Assessment

Fix $\epsilon > 0$, and choose ϕ and Δ sufficiently small (i.e., $0 < \phi < \bar{\phi}(\epsilon)$ and $\bar{\delta}(\epsilon) < \delta = e^{-r\Delta} < 1$) so that Lemmas 1 and 2 hold. In particular, we may assume

$$(1-\delta)\overline{v} < \epsilon/2 \quad \text{and} \quad \delta(v - p_c)\begin{cases} > w & \text{if } v \geq v_c, \\ < w & \text{if } v < v_c. \end{cases} \tag{D.2}$$

Also, from Lemma 1-(ii), (iii) and (iv),

$$u(\theta) = v - \max\left\{\frac{\delta\eta}{1-\delta(1-\eta)}(v - p^\dagger), w\right\} = \begin{cases} \frac{(1-\delta)v+\delta\eta p^\dagger}{1-\delta(1-\eta)} & \text{for } \theta \in \Theta_o \text{ such that } v \geq v_c, \\ v - w & \text{for } \theta \in \Theta_o \text{ such that } v < v_c. \end{cases} \tag{D.3}$$

From Lemma 1-(i) and (iv),

$$u(\theta) = v - w \begin{cases} > p_c & \text{for } \theta \in \Theta_r \text{ such that } v \geq v_c, \\ < p_c & \text{for } \theta \in \Theta_r \text{ such that } v < v_c. \end{cases} \tag{D.4}$$

We construct an assessment $\sigma_c = (\sigma_c^B, \sigma_c^S, f_c^S)$ as follows. Later, we will show that this assessment constitutes a quasi-Coasean equilibrium.

The Buyer's Strategy: Suppose that the seller offers p_n at $h_n \in H$. Any buyer type $\theta = (v, \tau)$ with $v \geq v_c$ accepts p_n if $p_n \leq P(\theta)$ and delays if $p_n > P(\theta)$, where

$$P(\theta) = P(v, \tau) := \begin{cases} (1-\delta)v + \delta p_c & \text{if } v \geq v_c \text{ and } \tau = \tau_r, \\ (1-\delta)v + \delta p_c - \delta\eta(p_c - p^\dagger) & \text{if } v \geq v_c \text{ and } \tau = \tau_o. \end{cases}$$

We call $P(\theta)$ the reservation price of a buyer type θ. Any buyer type $\theta = (v, \tau)$ with $v < v_c$ accepts p_n if $p_n \leq u(\theta) = v - w$ and exercises the outside option if $p_n > u(\theta) = v - w$.

In what follows, let $\Theta_c := \{(v, \tau) : v \geq v_c\}$ denote the buyer types with valuation weakly larger than v_c. Note that the reservation price $P(\theta)$ is defined only for buyer types in Θ_c. Note that $P(v_c, \tau_o) = p_c$ and $P(v, \tau) > p_c$ for all $\theta \in \Theta_c \backslash \{(v_c, \tau_o)\}$ by Lemma 1-(ii) and (iii).

The Seller's Strategy and Posterior Beliefs: Let p_c^+ denote the second lowest element in $\{P(v, \tau) : (v, \tau) \in \Theta_c\}$ (recall the p_c is the lowest element in this set). At any history h_n, the seller offers

$$p_n = p(h_n) := \begin{cases} p_c^+ & \text{at the null history } h_n = h_0, \\ p_c & \text{at all other histories.} \end{cases} \tag{D.5}$$

The seller never randomizes at any history. At any non-null history $h_n = (p_0, \ldots, p_{n-1})$ such that $p_k > P(v_c, \tau_o) = p_c$ for all k, the buyer type (v_c, τ_o) still remains in the negotiation, and thus, the seller's posterior belief $f_c^S(h_n)$ is well-defined via the Bayes' rule. Any other non-null histories are off the path, and we assume that the seller's posterior belief assigns probability 1 to (v_c, τ_o) at all such histories. Note that $f_c^S(h_n)$ assigns zero probability to the buyer types with $v < v_c$ at any non-null history both on and off the path.

Equilibrium Outcome: σ_c induces the following play on the path. In period 0, the seller offers $p_0 = p(h_0) > p_c$ and all buyer types in $\Theta_c \backslash \{(v_c, \tau_o)\}$ accept it, while the buyer type (v_c, τ_o) chooses to delay. All other buyer types exercise the outside option in period 0. In period 1, the seller offers $p_1 = p_c < p(h_0)$ and the remaining buyer type (v_c, τ_o) accepts it.

The outcome of σ_c is consistent of the definition of quasi-Coasean equilibrium. In the remaining part of the proof, we show that σ_c is indeed an ϵ-PBE. We first prove the optimality of the buyer's strategy σ_c^B.

Suppose that the seller offers p_n at $h_n \in H$ (in period n). If the buyer rejects p_n, the seller will offer $p_k = p_c$ in all subsequent periods. We analyze the responses of the rational and optimistic buyer types separately. First, suppose that the buyer is of a rational type. The buyer's payoff from rejecting p_n is $\max\{\delta(v-p_c), w\}$, where $\delta(v-p_c)$ is the payoff from accepting $p_{n+1} = p_c$ in the next period, and w is the payoff from exercising the outside option in the current period. Recall from (D.2) that $\delta(v-p_c) \geq w$ for any $\theta \in \Theta_c$ with $v \geq v_c$, hence these buyer types find it optimal to accept p_n iff $v - p_n \geq \delta(v-p_c)$, or equivalently, $p_n \leq P(v, \tau_r)$. On the other hand, $\delta(v-p_c) \leq w$ for all rational buyer types with $v \leq v_c$, and hence these buyer types find it optimal to accept p_n iff $u(v, \tau_r) = v - w \geq p_n$.

Next, suppose that the buyer is of an optimistic type. The buyer's payoff from rejecting p_n is

$$W_o(v) := \max\left\{\delta(1-\eta)(v-p_c) + \delta\eta(v-p^\dagger), \frac{\delta\eta}{1-\delta(1-\eta)}(v-p^\dagger), w\right\},$$

where $\delta(1-\eta)(v-p_c)+\delta\eta(v-p^\dagger)$ is the buyer's expected payoff from trading in the next period (an optimistic buyer believes that $p_{n+1} = p_c$ with probability $1-\eta$ and $p_{n+1} = p^\dagger$ with probability η), and $\frac{\delta\eta}{1-\delta(1-\eta)}(v-p^\dagger)$ is the payoff from waiting indefinitely until the seller trembles hands and offers $p^\dagger$. By Lemma 1-(iii) and (iv),

$$W_o(v) = \begin{cases} \delta(1-\eta)(v-p_c) + \delta\eta(v-p^\dagger) & \text{if } v \geq v_c \\ w & \text{if } v < v_c \end{cases}$$

Hence, each optimistic buyer type with $v \geq v_c$ indeed finds it optimal to accept p_n iff $v - p_n \geq \delta(1-\eta)(v-p_c) + \delta\eta(v-p^\dagger)$, or equivalently, $p_n \leq P(v, \tau_o)$. On the other hand, any optimistic buyer type with $v < v_c$ finds it optimal to accept p_n iff $u(v, \tau_o) = v - w \geq p_n$.

Next, we prove that $\sigma_c^S(h_n)$ is ϵ-best at any history $h_n \in H$, which will be a direct consequence of the observations in Claims D.1–D.3.

Claim D.1. *(i) $P(\theta) \leq u(\theta)$ for any $\theta = (v, \tau)$ such that $v \geq v_c$. (ii) $u(\theta) < p_c$ for any $\theta = (v, \tau)$ such that $v < v_c$.*

Proof. The part (i) follows the following observation: For any $\theta = (v, \tau)$ such that $v \geq v_c$,

$$P(\theta) = (1-\delta)v + \delta p_c = v - \delta(v-p_c) < v - w = u(\theta) \qquad \text{if } \tau = \tau_r,$$
$$P(\theta) = v - \delta(1-\eta)(v-p_c) - \delta\eta(v-p^\dagger) < v - \max\left\{\frac{\delta\eta}{1-\delta(1-\eta)}(v-p^\dagger), w\right\} = u(\theta) \qquad \text{if } \tau = \tau_o,$$

where the two inequalities follow (D.2) and Lemma 1, respectively. The part (ii) follows (D.3) and (D.4). □

Claim D.2. *(i) $P^D(h_n;\sigma_c) = (p_c, \infty)$. (ii) $V_S^{\sigma_c}(p_c; h_n) + \epsilon/2 \geq V_S^{\sigma_c}(p'; h_n)$ for any $h_n \in H$ and $p' \leq p_c$.*

Proof. The buyer type (v_c, τ_o) always belongs to $\mathrm{supp}(f_c^S(h_n))$ and chooses to delay in response to any $p_n > p_c = P(v_c, \tau_o)$; hence, $(p_c, \infty) \subset P^D(h_n; \sigma_c)$ at any h_n. On the other hand, by construction of σ_c, all buyer types $\theta \in \Theta$ never choose to delay in response to any $p_n \leq p_c$. Thus, $P^D(h_n;\sigma_c) = (p_c, \infty)$ at any h_n.

To show the part (ii) of the claim for the case $h_n = h_0$, suppose that the seller offers $p' \leq p_c$ at h_0. Because $p' \notin P^D(h_0;\sigma_c)$, all buyer types accept p' iff $u(\theta) \geq p'$, and thus, $V_S^{\sigma_c}(p'; h_0) - V_S^{\sigma_c}(p_c; h_0) = \Pi(p') - \Pi(p_c) \leq \epsilon/2$, where the last inequality holds due to Lemma 2. For the case $h_n \neq h_0$, note that the seller believes that the buyer's valuation v is weakly larger than v_c for any non-null history. Hence, any price $p' \leq p_c$ is accepted at $h_n \neq h_0$, yielding $V_S^{\sigma_c}(p'; h_n) - V_S^{\sigma_c}(p_c; h_n) = p' - p_c \leq 0 \leq \epsilon/2$. ☐

Claim D.3. *$|V_S^{\sigma_c}(p; h_n) - V_S^{\sigma_c}(p'; h_n)| \leq \epsilon/2$ for any $p, p' \geq p_c$ and $h_n \in H$.*

Proof. Fix a history $h_n \in H$. Suppose that the seller offers $p \geq p_c$ in period n at h_n. By construction, the seller will offer $p_{n+1} = p_c$ in the next period. All the buyer types with $v \geq v_c$ will accept either p_n or $p_{n+1} = p_c$ for sure, while all buyer types with $v < v_c$ will exercise the outside option in period n. Hence,

$$V_S^{\sigma_c}(p; h_n) \geq \delta p_c \sum_{\theta \in \Theta: v \geq v_c} f_\Theta(v, \tau) \quad \forall p \geq p_c.$$

On the other hand, because $P(v, \tau) \leq (1-\delta)\overline{v} + \delta p_c$ for all buyer types with $v \geq v_c$, hence,

$$V_S^{\sigma_c}(p'; h_n) \leq ((1-\delta)\overline{v} + \delta p_c) \sum_{\theta \in \Theta: v \geq v_c} f_\Theta(v, \tau) \quad \forall p' \geq p_c.$$

Combining these two inequalities,

$$|V_S^{\sigma_c}(p; h_n) - V_S^{\sigma_c}(p'; h_n)| \leq |\delta p_c - (1-\delta)\overline{v} - \delta p_c| \leq (1-\delta)\overline{v} < \epsilon/2 \quad \forall p, p' \geq p_c,$$

where the last inequality holds due to (D.2). ☐

Finally, we are ready to prove that σ_c^S is ϵ-best. First, consider the seller's strategy at the null history. By construction of σ_c^S, the seller will offer $p(h_0) \in (p_c, \infty) = P^D(h_0; \sigma_c)$. Claims D.2 and D.3 jointly imply that any deviation cannot increase the seller's payoff by more than ϵ, and hence σ_c^S is indeed ϵ-best at h_0.

Next, consider any non-null history $h_n = (p_0, p_1, \ldots, p_{n-1})$. There are two subcases.

- First, suppose that $p_k \leq p_c^+$ for some k. In this case, $f_c^S(h_n)$ assigns probability 1 to the buyer type (v_c, τ_o), and hence, the equilibrium offer $p(h_n) = p_c = P(v_c, \tau_o)$ is exactly optimal for the seller.
- Second, suppose that $p_k > p_c^+$ for any k. In this case, the seller offers $p(h_n) = p_c \notin P^D(h_n; \sigma_c)$. By construction, all buyer types (v, τ) with $v < v_c$ already have exercised the outside option, while any remaining buyer type will accept any price weakly lower than p_c; hence, $V_S^{\sigma_c}(p'; h_n) = p' \leq V_S^{\sigma_c}(p(h_n); h_n) = p_c$ for any $p' \notin P^D(h_n; \sigma_c)$. Finally, Claim D.3 guarantees that the seller cannot increase her continuation payoff by more than ϵ by deviating to $p' \in (p_c, \infty) = P^D(h_n; \sigma_c)$.

Hence, the seller's strategy σ_c^S is indeed ϵ-best.

E Bargaining with First-Order ϵ-Irrationality

E.1 Model

In this section, we consider the case that a small proportion of buyer types are possibly not fully rational. Suppose that the buyer's type space is given as follows:

$$\Theta = [\underline{v}, \overline{v}] \times \{\tau_r, \tau_{nr}\}.$$

A generic element of Θ is denoted by $\theta = (v, \tau)$, where v represents the buyer's valuation of the good, and τ indicates whether the buyer is rational ($\tau = \tau_r$) or non-rational ($\tau = \tau_{nr}$). The realizations of v and τ are stochastically independent and revealed only to the buyer before the negotiation begins. We denote the marginal probability of $\tau = \tau_{nr}$ by $\rho \in (0,1)$. The marginal density and distribution functions of v are denoted by $g : [\underline{v}, \overline{v}] \to \mathbb{R}$ and $G : [\underline{v}, \overline{v}] \to [0,1]$, respectively.[45] We assume that the hazard rate $g(v)/(1 - G(v))$ increases in v and $\underline{v}g(\underline{v}) < w$.[46]

The negotiation procedure is identical to the bargaining game in Section 2.1, including information structure, timing of moves, and payoff functions. In particular, the buyer has the identical payoff function and the set of feasible moves as specified in Section 2.1, regardless of whether he is rational or non-rational.[47] We also maintain the assumptions that (i) all buyer types have the same outside option w, with the focus on the case that $w \in (0, \infty)$, and (ii) all players have the common discounting factor $e^{-r\Delta} \in (0,1)$, where $r > 0$ and $\Delta > 0$ denote the discounting rate and the time duration between two consecutive periods, respectively.

The non-rational buyer types differ from the rational types only in that they may make ϵ-optimal decisions in equilibrium, while the rational types always make the *exactly* optimal decision. Formally, for any assessment $\sigma = (\sigma_B, \sigma^S, f^S)$, $v \in [\underline{v}, \overline{v}]$, $h_n \in H$, and $p_n \geq 0$, define $V_B^\sigma(v; h_n, p_n)$ as the buyer type v's expected payoff (discounted to period n) for the case that all players play according to σ after the seller offers p_n at h_n. Define $\widetilde{V}_B^\sigma(v; h_n, p_n)$ as the highest payoff (discounted to period n) that the buyer type v can obtain by making an one-shot deviation in response to the seller's offer p_n (and then playing according to σ in all subsequent periods). The buyer's strategy is called ϵ-optimal if $V_B^\sigma(v; h_n, p_n) + \epsilon \geq \widetilde{V}_B^\sigma(v; h_n, p_n)$ for any $v \in V$, $h_n \in H$, and $p_n \geq 0$.

We employ *ϵ-equilibrium* as our equilibrium concept in this section. The definition of ϵ-equilibrium is identical to the standard definition PBE except that the behavioral strategies of non-rational buyer types are required to be only ϵ-optimal as defined above. On the other hand, as in the standard definition of PBE, the seller and all rational buyer types are required to play *exactly* optimal behavioral strategies after any history, and the seller's posterior belief is obtained by Bayes rule whenever possible. For any $\epsilon > 0$, $\Delta > 0$, and $\rho \in (0,1)$, let $\mathcal{E}^*(\epsilon|\Delta, \rho)$ denote the set of all ϵ-equilibrium.[48]

Remark E.1. *The above model aims to capture the potential irrationality of some buyer types in a parsimo-*

[45] Departing from Section 2.1, we assume here that v is drawn from continuous probability distribution. This assumption simplifies the proof of the results in this section. All results remain valid if we assume that v is drawn from a discrete probability distribution.

[46] This assumption guarantees that the full commitment benchmark (2.2) in Section 2 admits a unique interior solution.

[47] Hence, the non-rational buyer types should *not* be equated with commitment types (who blindly stick to certain strategies) in reputational bargaining games.

[48] Note that ϵ-equilibrium differs from ϵ-PBE in Section D which allows the seller plays an ϵ-optimal strategy.

nious way. Note that we keep the model's deviation from the standard model (such that all players are fully rational) as minimally as possible. Formally, we will focus on the case that (i) the fraction of non-rational buyer types is small (i.e., $\rho \approx 0$), and moreover, (ii) the degree of those non-rational types' irrationality is also small (i.e., $\epsilon > 0$ is small).

Note that the model degenerates to the bargaining model in Section 2.1 if $\rho = 0$ or $\epsilon = 0$. In this case, the result in BP applies, and hence, the full-commitment benchmark outcome is achieved in the essentially unique equilibrium; in particular, the Coase conjecture fails and neither inter-temporal price discrimination nor delay occurs when the seller does not have full commitment power. The BP's result fails if both ρ and ϵ are positive.

Definition E.1. *For any $\epsilon > 0$, an ϵ-equilibrium is called a* quasi-Coasean ϵ-equilibrium *if the following outcome is induced on its path.*

- *Delay and Inter-temporal Pricing: The negotiation takes multiple periods with positive probability, and p_n declines over time on the equilibrium path.*
- *There is v_r^* and $v_{nr}^* \in [\underline{v}, \overline{v}]$ such that the following holds on the equilibrium path:*
  - *(i) Positive selection: Any rational buyer type with $v < v_r^*$ exercises the outside option immediately. Similarly, any non-rational buyer type with $v < v_{nr}^*$ exercises the outside option immediately.*
  - *(ii) Negative selection: Any rational buyer type with $v > v_r^*$ trades with the seller (possibly after a delay); among these buyer types, one with a higher valuation v trades earlier than others. Similarly, any non-rational buyer type with $v > v_{nr}^*$ trades with the seller (possibly after a delay); among these buyer types, one with a higher valuation v trades earlier than others.*

Proposition E.1. *Suppose $w > 0$. For any $\epsilon > 0$ and $\rho \in (0, 1)$, there is $d_{\rho,\epsilon} > 0$ such that $\mathcal{E}^*(\epsilon|\Delta, \rho)$ admits a quasi-Coasean ϵ-equilibrium whenever $\Delta \in (0, d_{\rho,\epsilon})$.*

The proof of Proposition E.1 is constructive. In this proof, we adopt the notations from Section 2. Fix $\epsilon > 0$ and $\rho > 0$. Choose $v^* \in (\underline{v}, \overline{v})$ and $\Delta > 0$ such that

$$\frac{1 - e^{-r\Delta}}{e^{-r\Delta}} w < \epsilon \quad \text{and} \quad v_\Delta^* := v^* + \frac{1 - e^{-r\Delta}}{e^{-r\Delta}} w < \overline{v} \tag{E.1}$$

$$v_\Delta^* - \frac{1 - G(v_\Delta^*)}{g(v_\Delta^*)} = w \quad \iff \quad (v - w)(1 - G(v)) \text{ is maximized at } v = v_\Delta^*. \tag{E.2}$$

v_Δ^* and v^* will play the role of v_r^* and v_{nr}^* in Definition E.1. In what follows, we will construct an ϵ-equilibrium such that the following play is observed on its path:

(i) The seller offers

$$p_0^C := (1 - e^{-r\Delta}) v_\Delta^* + e^{-r\Delta}(v^* - w) = v_\Delta^* - w \quad \text{and} \quad p_1^C := v^* - w < p_0^C,$$

in periods 0 and 1, respectively.

(ii) The rational buyer types with $v \geq v_\Delta^*$ trades with the seller in period 0.

(iii) The non-rational buyer types with $v \in [v_\Delta^*, \overline{v}]$ trade with the seller in period 0.

(iv) The non-rational buyer types with $v \in (v^*, v^*_\Delta)$ trades with the seller in period 1.

(v) All other buyer types exercise the outside option in period 0.

Equilibrium Assessment Consider the following strategy profile. Define

$$p_0^C := (1 - e^{-r\Delta})v^*_\Delta + e^{-r\Delta}(v^* - w) = v^*_\Delta - w \quad \text{and} \quad p_1^C := v^* - w,$$

so that

$$\max\{v - p_0, w\} \geq e^{-r\Delta}(v - p_1^C) \quad \text{if} \quad v \in [\underline{v}, \overline{v}] \quad \text{and} \quad p_0 \in [0, p_0^C], \quad \text{and} \tag{E.3}$$

$$v - p_0 = w = e^{-r\Delta}(v - p_1^C) \quad \text{if} \quad v = v^*_\Delta \quad \text{and} \quad p_0 = p_0^C. \tag{E.4}$$

We divide all non-null histories into the following two groups:

$$H_A := \{(p_0, p_1, \ldots, p_{n-1}) \in H \backslash \{h_0\} : p_k \in (p_1^C, p_0^C] \text{ for all } k \leq n-1\},$$
$$H_B := \{(p_0, p_1, \ldots, p_{n-1}) \in H \backslash \{h_0\} : p_k \notin (p_1^C, p_0^C] \text{ for some } k \leq n-1\}.$$

(C1) The seller offers

$$\begin{aligned} p_0 &= p_0^C && \text{at the null history,} \\ p_n &= p_1^C && \text{in any period } n \text{ after } h_n \in H_A, \\ p_n &= \overline{v} - w && \text{in any period } n \text{ after } h_n \in H_B. \end{aligned}$$

(C2) Suppose that the seller offers $p_n \in (p_1^C, p_0^C]$ after any history $h_n \in H_A \cup \{h_0\}$.

(C2-a) All the rational buyer types with valuation $v \in [\underline{v}, \overline{v}]$ and the non-rational buyer types with valuation $v \in [\underline{v}, v^*] \cup [v^*_\Delta, \overline{v}]$ employ the following behavioral strategy:

Accept p_n	(i.e., $\sigma^B(p_n; h_n, \theta)[T] = 1$)	if $v - p_n \geq w$,
Exercise the outside option	(i.e., $\sigma^B(p_n; h_n, \theta)[O] = 1$)	if $w > v - p_n$.

(C2-b) All non-rational buyer types with $v \in (v^*, v^*_\Delta)$ delay the negotiation to the next period.

(C3) Suppose that the seller offers $p_n \notin (p_1^C, p_0^C]$ after any history $h_n \in H_A \cup \{h_0\}$. Then, all buyer types accept p_n if $v - p_n \geq w$ and exercise the outside option in all other cases.

(C4) Suppose that the seller offers $p_n \geq 0$ after any history in $h_n \in H_B$. Then, all buyer types accept p_n if $v - p_n \geq w$ and exercise the outside option in all other cases.

Next, we specify the seller's posterior beliefs. Note that, on the equilibrium path, the seller offers $p_0 = p_0^C$ and then p_1^C in the first two periods, and all buyer types either exercise the outside option or trade with the seller by the end of period 1. In particular, any history in H_B lies off the equilibrium path, and hence we may assign any posterior belief for those cases.

(C5) According to (C1)–(C4), non-rational buyer types with $v \in (v^*, v^*_\Delta)$ choose to delay any offer in $(p_1^C, p_0^C]$ after any history in $H_A \cup \{h_0\}$. Hence, the Bayes rule is applicable to pin down the seller's posterior belief after any history in $H_A \cup \{h_0\}$.

(C6) After any history in H_B, we assume that the seller believes $(v,\tau) = (\overline{v}, \tau_r)$ with probability 1.

(Approximate) Optimality of the Buyer's Behavioral Strategy We first show (C2-a) and (C2-b) are exactly optimal and ϵ-optimal, respectively. By (C1), the seller will offer $p_{n+1} = p_1^C$ in the next period because (h_n, p_n) still belongs to H_A. If the buyer rejects $p_{n+1} = p_1^C$ again, the seller will offer $p_{n+k} = \overline{v} - w$ in all future periods. Hence, the buyer's highest payoff from delaying in period n is

$$e^{-r\Delta} \max\{w, v - p_1^C, e^{-r\Delta}(v - \overline{v} + w)\},$$

which is less than $\max\{v - p_n, w\}$ by (E.3). Hence, it is exactly optimal (i.e., ϵ-optimal with $\epsilon = 0$) for the buyer to accept p_n if $v - p_n \geq w$ and exercises the outside option if $v - p_n < w$. This shows that (C2-a) is exactly optimal for the buyer.

To show ϵ-optimality of (C2-b) for the buyer types with $v \in (v^*, v_\Delta^*)$, it suffices to show

$$v - p_n \leq e^{-r\Delta} \max\{v - p_1^C, w\} + \epsilon \quad \text{and} \quad w \leq e^{-r\Delta} \max\{v - p_1^C, w\} + \epsilon, \tag{E.5}$$

where $e^{-r\Delta} \max\{v - p_1^C, w\}$ represents the payoff from delaying the negotiation in period n. To show (E.5), first note that

$$v - p_n - e^{-r\Delta}(v - p_1^C) < (1 - e^{-r\Delta})(v_\Delta^* - p_1^C) = \frac{1 - e^{-r\Delta}}{e^{-r\Delta}} w < \epsilon, \tag{E.6}$$

where the first inequality in (E.6) holds because $v < v_\Delta^*$, and the last inequality in (E.6) follows (E.1). We also have $w \leq e^{-r\Delta} w + \epsilon$ by (E.1). Summing up all the observations so far, (E.5) holds and therefore (C2-b) is indeed ϵ-optimal for the buyer with $v \in (v^*, v_\Delta^*)$.

Next, to show the optimality of (C3), suppose that the buyer rejects $p_n \notin (p_1^C, p_0^C]$ after $h_n \in H_A \cup \{h_0\}$. Then, $(h_n, p_n) \in H_B$, and hence, the seller will offer $p_k = \overline{v} - w$ in all subsequent periods $k \geq n$. Note that the offer $p_k = \overline{v} - w$ is not acceptable to almost all buyer types, and therefore, the buyer's continuation payoff from delaying is dominated by the payoff from exercising the outside option in period n. This shows (C3) is exactly optimal for all buyer types. Similarly, (C4) is also exactly optimal for all buyer types.

Optimality of the Seller's Behavioral Strategy On the equilibrium path, the seller offers $p_0 = p_0^C$ and then $p_1 = p_1^C$. The negotiation ends by the end of period 1 for sure. We first check the seller has no incentive to deviate from the equilibrium after any non-null history in H_A. To see this, note that only the non-rational buyer types with $v \in (v^*, v_\Delta^*)$ still remain active after any $h_n \in H_A$. By (C2)–(C4), in all subsequent periods, those buyer types never accept any offer strictly higher than p_1^C; hence, it is indeed exactly optimal for the seller to offer $p_n = p_1^C$ after h_n.

Next, we show that the seller has no incentive to deviate in period 0 (at the null history). We first characterize the seller's expected payoff from offering p_0 in period 0, assuming that all players will follow the equilibrium strategy thereafter. Let $V_S(p_0)$ denote this payoff. First, consider the case $p_0 \in [\underline{v} - w, p_1^C] \cup (p_0^C, \overline{v} - w]$. By (C3) and (C4), all the buyer types accept p_0 iff $v - p_0 \geq w$ and exercise the outside option iff $v - p_0 < w$. Hence,

$$V_S(p_0) = p_0[1 - G(p_0 + w)] \quad \forall p_0 \in [\underline{v} - w, p_1^C] \cup (p_0^C, \overline{v} - w]. \tag{E.7}$$

Next, consider the case $p_0 \in (p_1^C, p_0^C]$. By (C1), the seller will offer $p_1 = p_1^C$ in the next period if p_0 is rejected in period 0. Let $v_0 \in [v^*, v^*_\Delta]$ denote the buyer's value such that $v_0 - p_0 = w$.

- By (C2-a), each rational buyer type $\theta = (v, \tau_r)$ trades with the seller at p_0 iff

$$v - p_0 \geq w \qquad \Longleftrightarrow \qquad v \geq v_0$$

  and exercises the outside option in period 0 in all other cases.

- By (C2-b), each non-rational buyer type with $v \notin (v^*, v^*_\Delta)$ responses to the seller's offer p_0 exactly identically to the rational buyer types with the same valuation (as described in the last bullet point). In particular, all non-rational buyer types with $v > v^*_\Delta$ accept p_0, while all non-rational buyer types with $v < v^*$ exercise the outside option. On the other hand, each non-rational buyer type with $v \in (v^*, v^*_\Delta)$ rejects p_0 and then trades with the seller at p_1^C in the next period.

Hence, the seller's expected payoff from offering $p_0 \in (p_1^C, p_0^C]$ is bounded as follows:

$$\begin{aligned} V_S(p_0) &= (1-\rho)p_0[1 - G(p_0 + w)] + \rho p_0[1 - G(v^*_\Delta)] + e^{-r\Delta}\rho[G(v^*_\Delta) + G(v^*)]p_1^C \\ &\leq (1-\rho)p_0[1 - G(p_0 + w)] + \rho p_0[1 - G(v_0)] + e^{-r\Delta}\rho[G(v_0) - G(v^*)]p_1^C \\ &= p_0[1 - G(p_0 + w)] + e^{-r\Delta}\rho[G(v_0) - G(v^*)]p_1^C \end{aligned} \tag{E.8}$$

for any $p_0 \in (p_1^C, p_0^C]$, and

$$V_S(p_0^C) = p_0^C[1 - G(p_0^C + w)] + e^{-r\Delta}\rho[G(v^*_\Delta) + G(v^*)]p_1^C. \tag{E.9}$$

Finally, we are ready to prove that the seller has no incentive to deviate in period 0. By (E.7)–(E.9), it suffices to show

$$p_0^C[1 - G(p_0^C + w)] \geq \max_{p_0 \in [\underline{v} - w, \overline{v} - w]} p_0[1 - G(p_0 + w)] = \max_{v \in [\underline{v}, \overline{v}]} (v - w)(1 - G(v)).$$

Because $p_0^C = (1 - e^{-r\Delta})v^*_\Delta + e^{-r\Delta}(v^* - w) = v^*_\Delta - w$ by definition of v^*_Δ and p_0^C, the left-hand side of the last inequality equals $(v^*_\Delta - w)(1 - G(v^*_\Delta))$. Hence, the last inequality directly follows (E.2).

It remains to check the seller's incentive after any non-null history in H_B. By (C1), the seller insists on $p_n = \overline{v} - w$ in all subsequent periods after $h_n \in H_B$. Additionally, by (C6), the seller believes that the seller believes $v = \overline{v}$ with probability 1 after any deviation from the equilibrium play; hence, the seller's insisting on $p_n = \overline{v} - w$ is clearly exactly optimal for the seller.

F Bargaining Between Inequity-Averse Buyer and Seller

In this section, we show that the inequity aversion (with a reasonable choice of parameters) does not make a qualitative change in the bargaining model in Section 2.1. Suppose the same bargaining environment, except that the buyer and the seller have the inequity aversion à la Fehr and Schmidt (1999).

- Suppose that the buyer exercises the outside option in period $n \geq 0$. In this case, the buyer and the seller obtain the final discounted utilities

$$U_B = e^{-rn\Delta}\Big[w - \alpha_B \max\{0 - w, 0\} - \beta_B \max\{w - 0, 0\}\Big] = e^{-rn\Delta}(1 - \beta_B)w \tag{F.1}$$

$$U_S = e^{-rn\Delta}\Big[0 - \alpha_S \max\{w - 0, 0\} - \beta_S \max\{0 - w, 0\}\Big] = e^{-rn\Delta}(-\alpha_S w), \tag{F.2}$$

respectively, where

(i) $e^{-r\Delta}$ represents the common discounting factor;

(ii) the pecuniary payoff of the buyer is w in (F.1), and that of the seller is 0 in (F.2);

(iii) for both $k = B$ and S, α_k and β_k respectively capture each player's distaste for disadvantageous and advantageous inequities.

- Suppose that the buyer accepts the seller's offer p_n in period n. The seller obtains pecuniary outcome p_n, and the buyer does $v - p_n$. Hence, the buyer's and the seller's final discounted utilities are

$$U_B = e^{-rn\Delta}\Big[v - p_n - \alpha_B \max\{p_n - (v - p_n), 0\} - \beta_B \max\{(v - p_n) - p_n, 0\}\Big] \tag{F.3}$$

$$U_S = e^{-rn\Delta}\Big[p_n - \alpha_S \max\{(v - p_n) - p_n, 0\} - \beta_S \max\{p_n - (v - p_n), 0\}\Big] \tag{F.4}$$

respectively.

The negotiation procedure is identical to the bargaining game in Section 2.1, including information structure, timing of moves, and payoff functions. As in Section 2.1, the buyer's value of the good (i.e., the buyer's type) $v \in V = \{v_1, v_2, \ldots, v_N\}$ is his private information, where $0 < \underline{v} := v_1 < v_2 < \ldots < \overline{v} := v_N$. The seller holds the prior belief that $v = v_j$ with probability $f(v_j) > 0$. Let $F(v) := \sum_{v' \le v} f(v')$ denote the cumulative distribution function. We also maintain the assumptions that all buyer types have the same outside option w and $0 < w < \underline{v}$. The solution concept is PBE.

Assumption We impose the following assumption throughout this section.

$$\alpha_k \ge 0 \text{ and } \beta_k < 1/2 \text{ for both } k = B \text{ and } S. \tag{A^*}$$

The condition $\beta_k < 1/2$ in Assumption (A^*) guarantees that both (F.3) and (F.4) are monotone in p_n. If $\beta_B > 1/2$, for example, the buyer's distaste for advantageous inequality becomes so strong that the buyer's overall utility increases as the buyer makes more payment to the seller; the condition $\beta_k < 1/2$ excludes such extreme cases. Note that $\beta_k < 1/2$ still allows β_k to be moderately positive or even negative.

Notations Throughout this section, we adopt the notations from Section 2. Additionally, define

$$\widetilde{U}_B(v,p) := v - p - \alpha_B \max\{p - (v - p), 0\} - \beta_B \max\{(v - p) - p, 0\} = \begin{cases} (1 - \beta_B)v - (1 - 2\beta_B)p & \text{if } p \le v/2 \\ (1 + \alpha_B)v - (1 + 2\alpha_B)p & \text{if } p \ge v/2 \end{cases}$$

as the buyer's utility from trading at price p (in period 0). Similarly,

$$\widetilde{U}_S(v,p) := p - \alpha_S \max\{(v - p) - p, 0\} - \beta_S \max\{p - (v - p), 0\}$$

denotes the seller's utility from trading with the buyer type v at price p (in period 0). Finally, define

$$P^* := \arg\max_{p \geq 0} \sum_{k=1}^{N} \mathbb{1}\{\widetilde{U}_B(v_k, p) \geq (1-\beta_B)w\} \widetilde{U}_S(v_k, p) f(v_k)$$

P^* is the solution of the seller's profit maximization problem if she is endowed with full commitment power.

Now, we are ready to state and prove our main proposition in this section. For any $p^* \in P^*$, there is a PBE, denoted by $\sigma(p^*)$, such that

(i) after any history, the buyer accepts p iff $\widetilde{U}(v,p) \geq (1-\beta_B)w$ or, otherwise, exercises the outside option;

(ii) the seller offers p^* in period zero; in all other cases, the seller offers p such that $\widetilde{U}_B(\overline{v}, p) = (1-\beta_B)w$, with the posterior belief that $v = \overline{v}$ with probability 1.[49]

It is straightforward to check this assessment indeed constitutes a PBE, and hence, we omit its proof. The next proposition shows that all other PBEs are payoff-equivalent to one of those PBEs.

Proposition F.1. *All PBEs are payoff-equivalent to one of the PBEs in* $\{\sigma(p^*) : p^* \in P^*\}$.

This proposition directly follows the following three preliminary lemmas.

Lemma F.1. $\widetilde{U}_B(v,p)$ *increases in* v *and decreases in* p. $\widetilde{U}_S(v,p)$ *increases in* p.

Proof. The monotonicity of $\widetilde{U}_S(v,p)$ with respect to p is straightforward, hence we focus on $\widetilde{U}_B(v,p)$. Assumption (A^*) guarantees that $\widetilde{U}_B(v,p)$ clearly decreases in p for any fixed $v \in V$. To show that $\widetilde{U}_B(v,p)$ increases in v, first note that the case $\alpha + \beta < 0$ is trivial, because in this case

$$\widetilde{U}_B(v,p) = \max\{(1-\beta_B)v - (1-2\beta_B)p, (1+\alpha_B)v - (1+2\alpha_B)p\}$$

and both $(1-\beta_B)v-(1-2\beta_B)p$ and $(1+\alpha_B)v-(1+2\alpha_B)p$ increase in v, and so does their maximum $\widetilde{U}_B(v,p)$. Hence, let us focus on the case

$$1 - 2\beta_B \leq 1 + 2\alpha_B \quad \iff \quad \alpha_B + \beta_B \geq 0,$$

and therefore

$$\widetilde{U}_B(v,p) = \min\{(1-\beta_B)v - (1-2\beta_B)p, (1+\alpha_B)v - (1+2\alpha_B)p\} \quad \forall v, p.$$

Fix $p \geq 0$ and consider two buyer types v_H and v_L such that $v_H \geq v_L$. The only non-trivial case arises when $\frac{v_L}{2} \leq p \leq \frac{v_H}{2}$. In this case,

$$\widetilde{U}(v_H, p) = (1-\beta_B)v_H - (1-2\beta_B)p \quad \text{and} \quad \widetilde{U}(v_L, p) = (1+\alpha_B)v_L - (1+2\alpha_B)p$$

and therefore, $\widetilde{U}(v_H,p) \geq \widetilde{U}(v_L,p)$ iff $(1-\beta_B)v_H - (1+\alpha_B)v_L \geq -2(\alpha_B+\beta_B)p$. Indeed,

$$-2(\alpha_B+\beta_B)p \leq -(\alpha_B+\beta_B)v_L = (1-\beta_B)v_L - (1+\alpha_B)v_L \leq (1-\beta_B)v_H - (1+\alpha_B)v_L \tag{F.5}$$

[49] There is a unique p that satisfies this property. See Lemma F.1 and the discussion that follows this lemma.

where the first and second inequalities hold due to conditions $\frac{v_L}{2} \leq p$ and $1 - \beta_B > 0$, respectively. $\square$

For any $v \geq w$, $\widetilde{U}_B(v,0) = (1-\beta_B)v \geq (1-\beta_B)w$ and $\widetilde{U}_B(v,v) = -\alpha_B v \leq 0 \leq (1-\beta_B)w$. Hence, there is a unique p_v such that

$$\widetilde{U}_B(v,p_v) = (1-\beta_B)w$$

where $(1-\beta_B)w$ is the buyer's utility from exercising the outside option. That is, the buyer type v is indifferent between accepting p_v and taking the outside option. Lemma F.1 implies that p_v increases in v. For any buyer type $v \geq w$, denote such price level $u^*(v)$ (i.e., $u^*(v) \equiv p_v$) and call this *net-value*. The buyer never accepts a price higher than $u^*(v)$.

Fix any PBE $\sigma = (\sigma^B, \sigma^S, f^S)$. For any non-null history $h_n \in H$ and price $p_n \geq 0$, let $\widetilde{\Gamma}^\sigma(h_n, p_n)$ denote the continuation game after (i) the buyer has rejected $p_0, p_1, \ldots, p_{n-1}$ in the first n periods, and then, (ii) the seller charges p_n in period n (but the buyer has not responded in period n yet). $\widetilde{\Gamma}^\sigma(h_n,p_n)$ begins with the buyer's decision node at which the buyer decides whether to accept p_n, exercise the outside option, or delay. Similarly, define $\widetilde{\Gamma}^\sigma(h_0,p_0)$ as the continuation game after the seller charges p_0 in period 0 (but the buyer has not responded yet). For any $(h_n,p_n) \in H \times [0,\infty)$, a history $h_m = (h_n, p_n, p_{n+1}, \ldots, p_{m-1})$, where $m \geq n+2$, is called *reachable in* $\widetilde{\Gamma}^\sigma(h_n,p_n)$ if it lies on the path of σ. Formally:

(i) $p_{n+1} \in \operatorname{supp} \sigma^S(h_n,p_n)$ and $p_k \in \operatorname{supp} \sigma^S(h_n,p_n,p_{n+1},\ldots,p_{k-1})$ for all $k = n+2,\ldots,m-1$.

(ii) $\sum_{v\in V} f^S(v;h_n)\sigma^B(p_n;h_n,v)[D] > 0$ and $\sum_{v\in V} f^S(v;h_n,p_n,\ldots,p_{k-1})\sigma^B(p_k;(h_n,p_n,\ldots,p_{k-1}),v)[D] > 0$ for all $k = n+1,\ldots,m-1$.

Lemma F.2. *In any PBE $\sigma = (\sigma^B,\sigma^S,f^S)$, $\inf \operatorname{supp} \sigma(h_n) \geq u^*(\underline{v}^\sigma(h_n))$ for all history $h_n \in H$.*

Proof. For any history h_n, define $\underline{p}^\sigma(h_n) \coloneqq \inf \operatorname{supp} \sigma^S(h_n)$ and $G \coloneqq \sup_{h_n\in H} u^*(\underline{v}^\sigma(h_n)) - \underline{p}^\sigma(h_n)$. Suppose for contradiction $G > 0$. Pick a positive number ϵ and a history $h_m \in H$ such that

$$u^*(\underline{v}^\sigma(h_m)) - \underline{p}^\sigma(h_m) > G - \epsilon > \epsilon \quad \text{and} \quad \epsilon < \frac{(1-\delta)w}{2}\frac{1-\beta_B}{\max\{1+2\alpha_B, 1-2\beta_B\}}. \tag{F.6}$$

Also, pick $p_m \in \operatorname{supp} \sigma^S(h_m)$ such that $p_m \in [\underline{p}^\sigma(h_m), \underline{p}^\sigma(h_m)+\epsilon)$.

We claim that all buyer types in $\operatorname{supp} f^S(h_m)$ accept p_m with probability 1. It suffices to show that all these buyer types strictly prefer accepting p_m to any other alternatives. First of all, because $p_m < u^*(\underline{v}^\sigma(h_m))$, all buyer types in $\operatorname{supp}(f^S(h_m))$ strictly prefer accepting p_m to exercising the outside option in period m or any future period. Next, we show that all of these buyer types also strictly prefer accepting p_m to any delayed purchase. Suppose for contradiction that there are $v^\diamond \in \operatorname{supp} f^S(h_m)$ and $h_\ell = (h_m, p_m, p_{m+1},\ldots,p_\ell)$, where $\ell > m$, such that (i) h_ℓ is reachable in $\widetilde{\Gamma}^\sigma(h_m,p_m)$, (ii) $\underline{v}^\sigma(h_m) \leq \underline{v}^\sigma(h_\ell) \leq v^\diamond$, and (iii) $\widetilde{U}_B(v^\diamond,p_m) \leq \delta^{\ell-m}\widetilde{U}_B(v^\diamond,\underline{p}^\sigma(h_\ell))$. By the definition of G and (F.6),

$$p_m \leq \underline{p}^\sigma(h_m) + \epsilon \leq u^*(\underline{v}^\sigma(h_m)) - G + 2\epsilon \quad \text{and} \quad \underline{p}^\sigma(h_\ell) \geq u^*(\underline{v}^\sigma(h_\ell)) - G \geq u^*(\underline{v}^\sigma(h_m)) - G.$$

Hence, because $\widetilde{U}_B(v^\diamond,p)$ decreases in p,

$$\widetilde{U}_B(v^\diamond, u^*(\underline{v}^\sigma(h_m)) - G + 2\epsilon) \leq \widetilde{U}_B(v^\diamond,p_m) \leq \delta^{\ell-m}\widetilde{U}_B(v^\diamond,\underline{p}^\sigma(h_\ell)) \leq \delta^{\ell-m}U_B(v^\diamond, u^*(\underline{v}^\sigma(h_m)) - G).$$

and therefore,

$$(1-\delta^{\ell-m})\widetilde{U}_B(v^\diamond, u^*(\underline{v}^\sigma(h_m))-G) \leq \widetilde{U}_B(v^\diamond, u^*(\underline{v}^\sigma(h_m))-G) - \widetilde{U}_B(v^\diamond, u^*(\underline{v}^\sigma(h_m))-G+2\epsilon).$$

The right-hand side on the last inequality is bounded from above by $2\epsilon\max\{1+2\alpha_B, 1-2\beta_B\}$, while the left-hand side is bounded from below by $(1-\delta)w(1-\beta_B)$.[50] Hence, the last inequality implies

$$\epsilon \geq \frac{(1-\delta)w}{2}\frac{1-\beta_B}{\max\{1+2\alpha_B, 1-2\beta_B\}}.$$

This contradicts (F.6).

By the essentially same argument, any deviation to a price in $(p_m, \underline{p}^\sigma(h_m)+\epsilon)$ also leads to an immediate trade with probability 1. This means that p_m is not a best response at h_m, contradicting the supposition that $p_m \in \operatorname{supp} f^S(h_m)$. □

Finally, the next lemma shows that all buyer types never delay the negotiation in response to any seller's offer in period 0.

Lemma F.3. *In any PBE, all buyer types never choose to delay in response to any seller's offer p_0 in period 0.*

Proof. This lemma easily follows from the previous lemma. Fix a PBE σ, and suppose that the seller offers p_0 in period 0. Suppose for contradiction that a positive measure of buyer types choose to delay in response to p_0. Let $\underline{v}^\sigma(p_0)$ denote the lowest buyer type among them. Then, the last lemma shows that the seller never offers lower than $u(\underline{v}^\sigma(p_0))$ in the continuation game. Hence, each buyer's payoff from delaying is bounded from above by

$$e^{-r\Delta}\max\{\widetilde{U}_B(v, u^*(\underline{v}(p_0)), (1-\beta_B)w\}.$$

The buyer type $\underline{v}^\sigma(p_0)$ would choose to delay in period 0 only if

$$\max\{\widetilde{U}_B(\underline{v}^\sigma(p_0), p_0), (1-\beta_B)w\} \leq e^{-r\Delta}\max\{\widetilde{U}_B(\underline{v}^\sigma(p_0), u^*(\underline{v}(p_0)), (1-\beta_B)w\} \tag{F.7}$$

where the left-hand side is the payoff from concluding the negotiation in period 0. By the definition of $u^*(\cdot)$, the right-hand side of the last inequality equals

$$e^{-r\Delta}\max\{\widetilde{U}_B(\underline{v}^\sigma(p_0), u^*(\underline{v}(p_0)), (1-\beta_B)w\} = e^{-r\Delta}\widetilde{U}_B(\underline{v}^\sigma(p_0), u^*(\underline{v}(p_0)) = e^{-r\Delta}(1-\beta_B)w > 0.$$

Hence, (F.7) holds only if

$$(1-\beta_B)w \leq \max\{\widetilde{U}_B(\underline{v}^\sigma(p_0), p_0), (1-\beta_B)w\} \leq e^{-r\Delta}(1-\beta_B)w.$$

This is impossible given that $w > 0$. □

[50]Because $v^\diamond \geq v^\sigma(h_m)$ by construction, $u^*(\underline{v}^\sigma(h_m)) - G \leq u^*(v^\diamond)$. From the earlier observation that $\widetilde{U}_B(v,p)$ decreases in p, we obtain $\widetilde{U}_B(v^\diamond, u^*(\underline{v}^\sigma(h_m)) - G) \geq \widetilde{U}_B(v^\diamond, u^*(v^\diamond)) = (1-\beta_B)w$.

Given that all buyer types never choose to delay, the seller's equilibrium offer in period 0 must be chosen from P^*. This completes the proof of Proposition F.1.