Northwestern | Kellogg

CMS-EMS
Center for Mathematical Studies in Economics and Management Sciences

Discussion Paper #1590

"Reputational Bargaining and Inefficient Technology Adoption"

Harry Pei & Maren Vairo
Northwestern University

May 23, 2023

J.L. Kellogg School of Management at Northwestern University
2211 Campus Drive, Evanston, Illinois 60208 | 847.491.3300 | kellogg.northwestern.edu

Reputational Bargaining and Inefficient Technology Adoption*

Harry Pei† Maren Vairo‡

May 23, 2023

Abstract: A buyer and a seller bargain over the price of an object. Both players can build reputations for being obstinate by offering the same price over time. Before players bargain, the seller decides whether to adopt a new technology that can lower his production cost and the buyer cannot observe this adoption decision. We show that players' reputational incentives can lead to inefficient adoption and significant delays in reaching agreement, and that these inefficiencies arise in equilibrium *if and only if* the social benefit from adoption is large enough. As a result, an increase in the social benefit from adoption may lead to a lower adoption probability and a longer expected delay.

Keywords: hold-up problem, inefficient technology adoption, reputational bargaining, delay.

1 Introduction

Suppose a supplier needs to decide whether to adopt a new technology that can lower his cost of production. Even when the gain from adoption outweighs its cost, the supplier might be reluctant to adopt due to the concern that after his investment becomes sunk cost, his clients will offer low prices and expropriate the gains from adoption. This is the well-known *hold-up problem*, which is a fundamental determinant of people's incentives to make relationship-specific investments, firms' incentives to adopt new technologies, as well as the boundaries of firms and organizations.

The severity of the hold-up problem depends on the bargaining process that determines the terms of trade as well as players' information about others' investment decisions. For example, Grossman and Hart (1986) assume efficient Nash bargaining and that investments are publicly observed. They show that investments are inefficient unless the player who makes the investment decision has all the bargaining power. Gul (2001) shows that investments are approximately efficient even when the investing player cannot make any offer and hence has no bargaining power, as long as his opponents *frequently revise their offers* and *cannot observe* how much he has invested.

*An earlier version was circulated under the title "Reputational Bargaining with Unknown Values". We thank Dilip Abreu, Sandeep Baliga, Francesc Dilmé, Jack Fanning, Lucas Maestri, Juan Ortner, Karthik Sastry, Bruno Strulovici, Chris Udry, and Alex Wolitzky for helpful comments and NSF Grant SES-1947021 for financial support.

†Department of Economics, Northwestern University. Email: harrydp@northwestern.edu

‡Department of Economics, Northwestern University. Email: mvairo@u.northwestern.edu

This paper revisits the hold-up problem by incorporating an important concern in practice, that players may have incentives to build reputations for being obstinate in the bargaining process and therefore, might be *reluctant to revise their offers*. Our main result shows that even when investment decisions cannot be observed by other players, players' reputational incentives can lead to *inefficient investments* and *costly delays in reaching agreements*. We also show that under-investment occurs in equilibrium *if and only if* there are large social gains from investment.

We augment the reputational bargaining model of Abreu and Gul (2000) with *a technology adoption stage* before the *bargaining stage*. A buyer and a seller bargain over the price of an object. The buyer's value is commonly known. The seller's production cost depends on his choice of production technology *before* the bargaining stage, which is his private information. In our baseline model, the seller either uses a *default technology*, or adopts a *new technology* that has a lower production cost compared to the default one but requires a positive adoption cost. Section 5 extends our results to settings where the seller chooses between a finite set of technologies.

In the bargaining stage, the buyer offers a price. The seller either accepts the buyer's offer, or demands a higher price after which players engage in a continuous-time war-of-attrition.[1] With positive probability, each player is one of the commitment types who makes an exogenous offer and never concedes. With complementary probability, they are rational and decide what to offer and whether to concede in order to maximize their payoff. To be consistent with the reputational bargaining literature pioneered by Abreu and Gul (2000), we focus on the limiting case where there is a rich set of commitment types and the probability of commitment types goes to zero.

Theorem 1 characterizes the equilibria of a reputational bargaining game where the distribution of the seller's production cost is *exogenous*. It shows that inefficient delays arise in equilibrium *if and only if* (i) the difference between the two production costs is large enough, and (ii) the seller has a low production cost with probability above some cutoff. Our inefficient bargaining result stands in contrast to the results in Kambe (1999), Abreu and Gul (2000), and Abreu, Pearce and Stacchetti (2015), which show that bargaining is efficient when players have no private information about their preferences, or when the only private information is about their discount rate.

The bargaining inefficiencies in our model stem from the buyer's incentive to *screen* the seller, that is, to induce sellers with different production costs to offer different prices. The buyer *cannot*

[1] The uninformed player making the offer first is also assumed in Abreu, Pearce and Stacchetti (2015) and Fanning (2022). In a working paper version, we consider the case where players make offers simultaneously, in which case the qualitative features of our results remain robust. When players endogenously decide who makes the first offer, there always exists an equilibrium in which the uninformed buyer makes their offer before the informed seller does.

screen the seller when she faces uncertainty about the seller's discount rate, as shown in Abreu, Pearce and Stacchetti (2015), but she can do so when she faces uncertainty about the seller's cost.

Why can the buyer screen the seller? Recall from Abreu and Gul (2000) that when each player chooses their offer, they face a trade-off between demanding more surplus and increasing their speed of building reputation. After the buyer offers a price *between the two production costs*, the high-cost seller's payoff from conceding is negative, so he can credibly commit not to concede. Due to his commitment power, the high type does not need to increase his speed of building reputation. This motivates him to demand a larger share of the surplus. Due to the low type's incentive to speed up reputation building, he has less incentive to pool with the high type when the latter demands more surplus. When the high type's demand approaches the entire surplus, the low type will separate from the high type and demand a low price. This explains why the buyer can induce different types of the seller to demand different prices. In equilibrium, the high type will demand the entire surplus, never concede, and trade with delay due to the low type's incentive constraint.

When is screening profitable? As in Abreu, Pearce and Stacchetti (2015), the buyer can offer a high price (we refer to as a *pooling offer* which, in equilibrium, equals the high-cost seller's Rubinstein bargaining price) and induce both types of the seller to trade immediately. She can also offer a low price that is between the high and the low production costs (we refer to as a *screening offer*), after which she will lose all her surplus when she faces the high-cost type. Hence, screening is profitable for the buyer only if (i) she can pay a lower price to the low type, and (ii) the low type occurs with high enough probability. The former is true *if and only if* the difference between the two costs is large enough. This is because when the cost difference is small, all the screening offers are too low relative to the low type's Rubinstein bargaining price, in which case the low type can demand something greater than the high type's Rubinstein bargaining price while still inducing the buyer to concede immediately. If this is the case, then screening is unprofitable for the buyer under all cost distributions since she ends up paying a higher price to the low type.

Theorem 2 shows that, in the game with *endogenous technology adoption*, the seller's adoption decision is bounded away from efficiency (under an open set of adoption costs) *if and only if* bargaining is inefficient under some exogenous distribution over production costs. Otherwise, his adoption decision is *approximately efficient* regardless of the adoption cost. Our theorems imply that inefficient adoption occurs in equilibrium *if and only if* there are inefficient delays in bargaining, and the latter occur *if and only if* the social benefit from adoption is large enough. Therefore, an increase in the social benefit from adoption may *decrease* the probability of adoption.

In order to understand why, let the seller's *private gain* from adoption be the increase in his payoff from bargaining once he lowers his production cost. We use an observation that the seller's private gain from adoption equals the social benefit from adoption when the buyer makes the pooling offer, but is bounded below the social benefit when the buyer makes the screening offer.

Let us focus on the case where the adoption cost is between the seller's private gain from adoption and the social benefit from adoption. In equilibrium, the buyer cannot strictly prefer the screening offer, since the seller's gain from adoption will be lower than his adoption cost, in which case he will never adopt and screening will be unprofitable. The buyer cannot strictly prefer the pooling offer, since the seller's gain from adoption will exceed his adoption cost, in which case he will adopt for sure, which will provide the buyer a strict incentive to make the screening offer.

Therefore, in every equilibrium, the buyer must be indifferent between the screening offer and the pooling offer, and the seller must be indifferent between adopting and not adopting. When the social benefit from adoption is large enough such that the buyer prefers the screening offer under some distribution over production costs, the seller will adopt with a probability that makes the buyer indifferent between the screening offer and the pooling offer, and the buyer will mix between the pooling offer and the screening offer. There will be costly delay in reaching agreement since the screening offer is made with positive probability. The equilibrium adoption decision will be *inefficient* since it is efficient to adopt yet the seller's adoption probability is bounded below 1.

In the complementary scenario where the social benefit from adoption is small, Theorem 1 suggests that the buyer prefers to make the pooling offer regardless of the adoption probability, in which case the seller's private gain from adoption equals the social benefit, which is assumed to be strictly greater than the cost of adoption. This seems to suggest that there is no adoption probability under which the seller is indifferent between adopting and not adopting.

The above contradiction arises since Theorem 1 *cannot* be applied to settings where the distribution of production cost is *endogenous*: It only applies once we *fix* the distribution of production cost *as the probability of commitment types vanishes*. But in the game with endogenous technology adoption, the distribution of production cost may depend on the probability of commitment types. We show that in equilibrium, the adoption probability approaches 1 as the probability of commitment types vanishes. The buyer makes a screening offer, the low-cost seller accepts with probability close to 1, and the high-cost seller counteroffers a higher price and trades with delay. However, the expected loss from delay is close to 0 since the adoption probability is close to 1.

Our inefficient adoption result suggests an explanation for the under-adoption of cost-saving technologies, which is widely documented in agriculture and in other industries. Wolitzky (2018) provides a complementary explanation based on social learning, which fits applications where producers *do not know* the effectiveness of the new technology. Our theory fits when (i) producers *know* that the new technology is effective from earlier adoption results, (ii) it is hard for potential buyers to observe the producers' adoption decisions (unlike in Grossman and Hart 1986 where investments are observable), but (iii) producers are reluctant to adopt due to the fear of being held-up.

One example that fits our setting is the under-adoption of Bt cotton and other genetically-modified crops. It is well-known among farmers that Bt cotton can significantly reduce insecticide applications which can lower the cost per unit yield (Qaim and de Janvry 2003). It is hard for potential buyers to observe the farmers' adoption decisions: Although Bt cotton is more resistant to pests compared to traditional cotton, it is hard to distinguish the two from their appearances. However, the adoption rates of Bt cotton are very low in developing countries inspite of its effectiveness in lowering production costs. For example, Ali and Abdulai (2010) find that the adoption rate is only 62% in the Punjab province of Pakinstan. Although there are other explanations, such as the lack of access to credit markets, farmers' concerns about the hold-up problem also seem to be relevant since (i) the adopted farmers in Pakistan are more likely to be members in organizations that have more bargaining power over prices, and (ii) the farmers' share of surplus in Pakistan is much lower than that in countries that have higher adoption rates (Falck-Zepeda, Traxler, and Nelson 2000). We conclude this section by discussing our contributions to the related literature.

Hold-Up Problem: Grossman and Hart (1986) assume that investments are *observable* and show that investments are bounded below the socially optimum unless the player who invests has all the bargaining power. In order to protect players against others' opportunism, one needs additional tools such as vertical integration (Williamson 1979), relational incentives (Baker, Gibbons and Murphy 2002), and investing over time (Che and Sákovics 2004). When investments are *unobservable*, Gul (2001) shows that investments are approximately efficient even when the investing player has no bargaining power, as long as their opponent frequently revise their offers.

We incorporate the concern that players might be reluctant to revise their offers due to their reputational incentives. We show that the hold-up problem emerges even when investments are *unobservable*,[2] and that inefficiencies occur if and only if the benefit from investment is large enough.

[2] When the seller's investment can be *observed* by the buyer in a reputational bargaining game, bargaining will be efficient but the seller's adoption decision is *socially inefficient* unless he is arbitrarily more patient than the buyer.

Reputational Bargaining: We contribute to the reputational bargaining literature from two angles. First, compared to the seminal works of Kambe (1999), Abreu and Gul (2000), and Compte and Jehiel (2002), we incorporate *heterogeneity* in players' costs in reputational bargaining.[3] In contrast to the heterogeneity in players' discount rates in Abreu, Pearce and Stacchetti (2015, or APS), heterogeneity in players' costs or values enables the uninformed player to *screen* the informed player, leading to inefficient delays.[4] Fanning (2022) shows that bargaining is efficient when players' values or costs are drawn from a rich set. In contrast, we show that bargaining is inefficient as long as there exist two adjacent types who have sufficiently different production costs. This violates Fanning's richness assumption.[5] His assumption fits when investment is a continuous choice variable (e.g., human capital). Our assumption fits when the cost heterogeneity is driven by the *differences in production technologies*, since adoption decisions are usually *indivisible* and adopting an innovative technology may significantly reduce the cost of production.

Second, compared to APS and Fanning (2022) that focus on the case where the distribution of preference is *exogenous*, we also analyze the case where the distribution of production cost is *endogenous*. To the best of our knowledge, this has not been studied in the existing literature on reputational bargaining. We explain why the existing characterization results on reputational bargaining with *exogenous* type distributions cannot be applied to settings with an *endogenous* type distribution. The arguments in our proof are portable to other settings with endogenous type distributions, such as reputational bargaining with endogenous information acquisition.

Bargaining with Incomplete Information: Our paper is also related to the literature on bargaining with incomplete information such as Gul, Sonnenschein and Wilson (1986), Ausubel and Deneckere (1989), Kim (2009), Strulovici (2017), and Liu, Mierendorff, Shi and Zhong (2019).

Our bargaining inefficiencies stem from players' reputational incentives, which are conceptually different from the ones that are caused by higher-order uncertainty (Feinberg and Skrzypacz 2005),

[3]Ekmekci and Zhang (2022) study reputational bargaining where a player's payoff from their outside option depends on whether the other player is the commitment type. Their model has interdependent values and only one rational type for each player. In contrast, we study a private value model where the seller has multiple rational types.

[4]APS also consider the case in which there are commitment types using *non-stationary strategies*. We assume that all commitment types demand the same price over time. This is because our motivation is to revisit the hold-up problem when players are unwilling to revise their offers (due to their incentives to build reputations for being obstinate). This can be better captured by commitment types who demand the same price over time.

[5]Both our model and Fanning (2022) assume that there is *a rich set of commitment types*, that is to say, a player can be a commitment type with positive probability regardless of the price they offer. The key difference is that Fanning (2022)'s efficiency result requires another richness assumption on *the set of production costs*, namely, every production cost occurs with positive probability and is much greater compared to the probability of commitment types. This assumption is violated in our model in the case whenever there are bargaining inefficiencies.

interdependent values (Deneckere and Liang 2006, Baliga and Sjöström 2023), costly concessions (Dutta 2022), risk aversion (Dilmé and Garrett 2022), and new arrivals (Fuchs and Skrzypacz 2010).

Ortner (2017) shows that when a durable good monopolist's cost may decrease over time, the bargaining outcome is efficient if and only if the consumers' values are drawn from an interval. Although our analysis reaches a similar conclusion that bargaining is efficient if and only if the gap between adjacent types' production costs is small enough, the underlying logic is different from Ortner (2017): The gap in production cost in our model makes it profitable for the uninformed buyer to screen the informed seller when both of them can build reputations for being obstinate.

2 Baseline Model

Time is continuous, indexed by $t \in [0, +\infty]$. A buyer (she) and a seller (he) bargain over the price of an object. The buyer's value is common knowledge, which is normalized to 1.

Time 0 consists of two stages. In the first stage, the seller decides whether to adopt a new technology at an *adoption cost* $c > 0$. This adoption decision determines his cost of producing the object, which cannot be observed by the buyer. If he adopts the new technology, then his production cost is θ_1. If he uses the default technology, then his production cost is θ_2. We assume that $0 < \theta_1 < \theta_2 < 1$. Therefore, trade is efficient regardless of the seller's production cost.

In the second stage, the buyer makes an offer $p_b \in [0, 1]$.[6] The seller either accepts the offer and sells at price p_b, or rejects the offer and counteroffers $p_s \in (p_b, 1]$, after which players engage in a continuous-time war-of-attrition. If a player concedes, then players trade at the price offered by their opponent. If players concede at the same time, then they trade at price $\frac{p_b+p_s}{2}$.

Players have the same discount rate $r > 0$. If trade happens at time $\tau \in [0, +\infty)$ and at price $p \in [0, 1]$, then the buyer's payoff is $e^{-r\tau}(1-p)$ and the seller's payoff is $e^{-r\tau}(p-\theta) - \widetilde{c}$, where θ stands for the seller's production cost and $\widetilde{c} \in \{0, c\}$ stands for his adoption decision. If players never trade, then $\tau = +\infty$, in which case the buyer's payoff is 0 and the seller's payoff is $-\widetilde{c}$.

Each player is rational with probability $1-\varepsilon$ and is one of the commitment types with probability $\varepsilon > 0$. Each commitment type of the buyer is characterized by $p_b \in \mathbf{P}_b \equiv \{p_b^1, p_b^2, ..., p_b^{N_b}\}$, who offers p_b and never accepts any price greater than p_b. Each commitment type of the seller is characterized by $p_s \in \mathbf{P}_s \equiv \{p_s^1, p_s^2, ..., p_s^{N_s}\}$, who offers p_s and never accepts any price lower than p_s. Let $\mu_b \in \Delta(\mathbf{P}_b)$ and $\mu_s \in \Delta(\mathbf{P}_s)$ be the distributions of players' types conditional on being

[6]The uninformed player making the offer first is also assumed in APS and Fanning (2022).

committed. We assume that $0 = p_b^1 < p_b^2 < ... < p_b^{N_b} = 1$ and $0 = p_s^1 < p_s^2 < ... < p_s^{N_s} = 1$. Let

$$\nu \equiv \max_{i \in \{s,b\}} \max_{j \in \{2,...,N_i\}} \left| p_i^j - p_i^{j-1} \right|. \tag{2.1}$$

Intuitively, ν measures the richness of the sets of commitment types. In order to be consistent with Abreu and Gul (2000), our analysis focuses on the case where both ν and ε are close to 0, that is, both sets of commitment types are rich and players are rational with probability close to one.

The public history consists of players' offers and whether any player has conceded. The buyer's private history consists of the public history and whether she is committed. The seller's private history consists of the public history, whether he is committed, and his adoption decision. The buyer's strategy consists of her offer $\sigma_b \in \Delta[0,1]$ and a mapping from players' offers to her concession time $\tau_b : [0,1]^2 \to \Delta(\mathbb{R}_+ \cup \{+\infty\})$. The seller's strategy consists of his adoption decision, or equivalently, the distribution of his production cost $\pi \in \Delta\{\theta_1, \theta_2\}$, a mapping from his production cost and the buyer's offer to his offer $\sigma_s : \{\theta_1, \theta_2\} \times [0,1] \to \Delta[0,1]$, and a mapping from his production cost and players' offers to his concession time $\tau_s : \{\theta_1, \theta_2\} \times [0,1]^2 \to \Delta(\mathbb{R}_+ \cup \{+\infty\})$. The solution concept is Perfect Bayesian equilibrium (or *equilibrium* for short).

Benchmark: Observable Adoption Decision Suppose the buyer *can observe* the seller's adoption decision, or equivalently, the buyer knows the seller's production cost θ. Abreu and Gul (2000) show that as $\varepsilon \to 0$, players will trade with almost no delay at price approximately $p_\theta \equiv \frac{1+\theta}{2}$. We call p_θ type-θ seller's *Rubinstein bargaining price*, since it is the equilibrium price in the bargaining game of Rubinstein (1982) between a buyer with value 1 and a seller with cost θ.

The intuition is that the buyer can secure payoff $1 - p_\theta$ by offering p_θ and the seller can secure payoff $p_\theta - \theta$ by demanding p_θ. Their guaranteed payoffs are their equilibrium payoffs since the sum of these payoffs equals the social surplus from trade $1 - \theta$. The seller's gain from adoption is

$$(p_{\theta_1} - \theta_1) - (p_{\theta_2} - \theta_2) = \frac{\theta_2 - \theta_1}{2}. \tag{2.2}$$

Hence, the seller will adopt when $c \leq \frac{\theta_2 - \theta_1}{2}$. Since it is socially efficient to adopt as long as $c < \theta_2 - \theta_1$, the equilibrium adoption decision is inefficient when $c \in \left(\frac{\theta_2-\theta_1}{2}, \theta_2 - \theta_1\right)$. In summary, when the seller's adoption decision is observable, bargaining will be efficient but the seller's adoption decision will be inefficient. Later on, we compare this to the case with unobservable adoption, in which there is inefficient adoption *if and only if* there are costly delays in bargaining.

Remarks: Our baseline model assumes that the seller chooses between two production technologies. Section 5 extends our results to settings where the seller chooses between three or more technologies. Our baseline model assumes that players have the same discount rate and the same commitment probability. In Section 6, we comment on extensions of our results to environments with heterogeneous discount rates and heterogeneous commitment probabilities. We also discuss how players' bargaining powers affect the severity of the hold-up problem and the expected delay.

We only analyze the benchmark in which the buyer observes the seller's adoption decision and the case in which the buyer cannot observe the seller's adoption decision. If the buyer observes a noisy private signal of the seller's adoption decision, then it introduces further complications such as higher-order uncertainty about the seller's cost, which we abstract away from in our analysis.

We assume that every obstinate type demands the same price over time, which is also assumed in Abreu and Gul (2000), Ekmekci and Zhang (2022), and Fanning (2018, 2022). In contrast, Abreu and Pearce (2007), Wolitzky (2012), and APS allow obstinate types to make time-dependent demands. Our assumption that obstinate types' demands are *time-independent* is motivated by the concern that once a player lowers their demand, it might be hard for them to convince others that they are obstinate. This assumption fits the motivation of our analysis, which is to revisit the hold-up problem when players find it costly to revise their offers due to their reputation concerns.

Our baseline model assumes that the uninformed buyer makes their offer before the informed seller. This assumption is standard in reputational bargaining models with one-sided uncertainty about players' payoffs, such as APS and Fanning (2022). In a working paper version, we assume that players make their initial offers simultaneously and obtain qualitatively similar results. Section 6 revisits this issue by commenting on the case where the timing of offers is endogenous.

While our results are stated in environments where the seller decides whether to lower their production cost, they can be applied to several other settings, such as (i) when the seller decides whether to *divest* and receive benefits for becoming less cost-efficient, (ii) when the buyer chooses whether to increase her value, such as a downstream firm investing to expand its customer base.

3 Main Results

Section 3.1 analyzes a reputational bargaining game in which the distribution of production cost is *exogenous*. Theorem 1 provides conditions under which players will reach an agreement after significant delays and highlights the sources of inefficiencies. Section 3.2 analyzes the reputational bargaining game with endogenous technology adoption. Theorem 2 shows that reputation concerns lead to inefficient adoption *if and only if* there are large social gains from adoption. We explain the subtleties when analyzing reputational bargaining games with endogenous type distributions.

3.1 Reputational Bargaining with Exogenous Production Cost

This section analyzes a reputational bargaining game when θ is drawn from an *exogenous* full support distribution $\pi \in \Delta\{\theta_1, \theta_2\}$. We refer to the rational seller with production cost θ as *type* θ. Our result in this section characterizes the common properties of all equilibria as the probability of commitment types vanishes and the set of commitment types $\mathbf{P}_b$ and $\mathbf{P}_s$ are rich enough.

We describe players' limiting equilibrium strategies and later explain why they arise in equilibrium. Let $\underline{\sigma}_b$ denote the buyer's strategy of offering $\min\{p_{\theta_1}, \theta_2\}$. Let $\overline{\sigma}_b$ denote the buyer's strategy of offering p_{θ_2}. By definition, $p_{\theta_2} > \min\{p_{\theta_1}, \theta_2\}$. Let $\sigma_s^*(\cdot) \equiv \left\{\sigma_{s,\theta}^*(\cdot)\right\}_{\theta\in\Theta}$, where

$$\sigma_{s,\theta}^*(p_b) \equiv \begin{cases} 1, & \text{if } p_b \leq \theta_1, \text{ or } p_b \in (\theta_1, \theta_2] \text{ and } \theta = \theta_2, \\ \max\left\{p_b, 1+\theta_1 - p_b\right\}, & \text{if } p_b \in (\theta_1, \theta_2] \text{ and } \theta = \theta_1, \\ \max\left\{p_b, 1+\theta_2 - p_b\right\}, & \text{if } p_b > \theta_2. \end{cases} \tag{3.1}$$

Intuitively, $\sigma_{s,\theta}^*(\cdot)$ stands for type-θ seller's counteroffer as a function of the buyer's offer p_b. We adopt the convention that if a type θ accepts the buyer's offer, then his counteroffer is p_b.

In order to understand the expression for $\sigma_{s,\theta}^*$, we explain the intuition behind $\max\left\{p_b, 1+\theta_i - p_b\right\}$. Recall that in a reputational bargaining game where it is common knowledge that $\theta = \theta_i$, for any pair of offers p_b and p_s with $\theta_i < p_b < p_s < 1$, the seller will concede at rate

$$\lambda_s \equiv \frac{r(1-p_s)}{p_s - p_b}, \tag{3.2}$$

and the buyer will concede at rate

$$\lambda_b^i \equiv \frac{r(p_b - \theta_i)}{p_s - p_b}. \tag{3.3}$$

As the probability of commitment types ε vanishes, the player with a lower concession rate will concede at time 0 with probability close to 1. According to (3.2) and (3.3), when θ_i is common knowledge, players will concede at the same rate when the seller offers $1+\theta_i-p_b$. This implies that the seller can secure a price of approximately $\max\left\{p_b, 1+\theta_i - p_b\right\}$ by *either* accepting the buyer's offer *or* counteroffering something slightly below $1+\theta_i-p_b$ and inducing the buyer to concede with probability close to 1 at time 0. Let

$$\pi^* \equiv \min\left\{1, \frac{p_{\theta_2} - \theta_2}{\min\{p_{\theta_1}, \theta_2\} - \theta_1}\right\}, \tag{3.4}$$

which by definition is strictly positive. In addition, $\pi^* < 1$ *if and only if*

$$\theta_2 - \theta_1 > \frac{1 - \theta_2}{2}, \tag{3.5}$$

that is, the difference between θ_1 and θ_2 is large relative to the surplus generated by the high-cost type. Recall that the sets of commitment types are rich when $\nu \to 0$. Theorem 1 shows that for generic π, all equilibria of the reputational bargaining game converge to the same limit point as ν and ε go to 0.[7] It also characterizes the welfare properties of the unique limiting equilibrium.

Theorem 1. *There exists at least one equilibrium of the reputational bargaining game with exogenous production costs. Suppose* $\pi \in \Delta\{\theta_1, \theta_2\}$ *satisfies* $\pi(\theta_1) \notin \{0, \pi^*, 1\}$*. For every* $\eta > 0$*, there exists* $\bar{\nu} > 0$ *such that when* $\nu < \bar{\nu}$*, there exists* $\bar{\varepsilon}_\nu > 0$ *such that for every* $\varepsilon \in (0, \bar{\varepsilon}_\nu)$ *and every equilibrium* $(\sigma_s, \sigma_b, \tau_s, \tau_b)$ *under* (ε, ν)*:*

1. *If* $\pi(\theta_1) < \pi^*$*, then* σ_b *is* η*-close to* $\overline{\sigma}_b$ *and* σ_s *is* η*-close to* σ_s^* *on the equilibrium path. The expected welfare loss from delay is less than* η *conditional on every* $\theta \in \Theta$*.*

2. *If* $\pi(\theta_1) > \pi^*$*, then* σ_b *is* η*-close to* $\underline{\sigma}_b$ *and* σ_s *is* η*-close to* σ_s^* *on the equilibrium path. Conditional on* $\theta = \theta_1$*, the expected welfare loss from delay is less than* η*. Conditional on* $\theta = \theta_2$*, the buyer's equilibrium payoff is* 0 *and the expected welfare loss from delay is* η*-close to*

$$(1 - \theta_2)\left\{1 - \frac{\max\left\{p_{\theta_1}, 1 - (\theta_2 - \theta_1)\right\} - \theta_1}{1 - \theta_1}\right\}. \tag{3.6}$$

The proof is in Section 4.1 with some details relegated to the Online Appendix. According to Theorem 1, the qualitative features of the limiting equilibrium depend on (i) the difference $\theta_2 - \theta_1$ between the two production costs, and (ii) the distribution π over production costs. In particular,

1. When the difference between θ_1 and θ_2 is small in the sense that θ_1 and θ_2 violate (3.5), the buyer offers a high price p_{θ_2} and the seller accepts immediately. The same limiting equilibrium arises when θ_1 and θ_2 satisfy (3.5) and the low type occurs with probability less than π^*.

[7] Throughout the paper, we measure the distance between two distributions (e.g., two mixed strategies) using the Prokhorov metric, defined in Billingsley (2013a). Intuitively, two distributions μ and μ' are close if for every Borel set A, the value of $\mu(A)$ is close to that of $\mu'(A')$ for some small neighborhood A' of A.

2. When θ_1 and θ_2 satisfy (3.5) and $\pi(\theta_1) > \pi^*$, the buyer offers a low price $\min\{p_{\theta_1}, \theta_2\}$. The high type demands the entire surplus 1 and the buyer concedes after some delay, which leads to an expected welfare loss of (3.6). The low type trades immediately either by accepting the buyer's offer or by offering $1 - (\theta_2 - \theta_1)$, depending on the comparison between p_{θ_1} and θ_2.

Theorem 1 suggests that costly delay occurs in equilibrium *if and only if* the difference between the two production costs is large enough and the seller is likely to have a low production cost. Our inefficient bargaining result stands in contrast to the efficiency results in reputational bargaining games without private information about payoffs (Kambe 1999, Abreu and Gul 2000), as well as those in games where one of the players has private information about their discount rate (Abreu, Pearce and Stacchetti 2015). Those papers show that when there is a rich set of *stationary* commitment types and the probability of commitment types vanishes, players will trade at the strongest type's (e.g., the most patient type) Rubinstein bargaining price with almost no delay.

We argue that inefficient delay arises whenever the buyer (i.e., the *uninformed player*) uses her offer to *screen* the seller, that is, to induce sellers with different costs to demand different prices. Screening is *feasible* when players can build reputations and the uninformed player faces uncertainty about her opponent's cost. Screening is *profitable* for the uninformed player when both the probability of the low type and the difference between the two costs are large enough.

To elaborate, we start from an auxiliary game where *both types of the seller are required to demand the same price.* If $\theta_2 < p_b < p_s < 1$, then the seller concedes at rate $\frac{r(1-p_s)}{p_s-p_b}$, with the high type starting to concede only after the low type has finished conceding. The buyer first concedes at rate $\frac{r(p_b-\theta_1)}{p_s-p_b}$ and then concedes at a lower rate $\frac{r(p_b-\theta_2)}{p_s-p_b}$ after the high type starts to concede. As $\varepsilon \to 0$, the buyer spends most of her time conceding at rate $\frac{r(p_b-\theta_2)}{p_s-p_b}$, so her average concession rate is close to $\frac{r(p_b-\theta_2)}{p_s-p_b}$. As in Abreu and Gul (2000), each player faces a trade-off between demanding more surplus and increasing their average concession rate relative to their opponent's. As long as both types of the seller counteroffer the same price, it is optimal for the buyer to offer p_{θ_2} by which she can secure payoff of $1-p_{\theta_2}$ regardless of the seller's offer. This leads to the efficient equilibrium.

However, instead of offering a high price p_{θ_2} and inducing both types of the seller to trade immediately, the buyer can also *screen* the seller by offering a low price that belongs to $(\theta_1, \theta_2]$. To see why different types of the seller will counteroffer different prices, note that type θ_2 obtains a negative payoff from conceding, so he can *credibly commit* not to concede. Due to his commitment power, the usual trade-off between demanding more surplus and increasing his concession rate is no longer relevant, which motivates him to demand a high price. Type θ_1 receives a strictly positive payoff from conceding, so he faces a trade-off between (i) demanding more surplus, (ii) increasing his concession rate, and (iii) pooling with the high type. When the high type's demand increases, it is more costly for the low type to imitate the high type since doing so will further lower his concession rate. When the high type's demand is close to the entire surplus, the low type prefers to offer a low price even at the expense of separating from the high type.

What will the high type offer *in equilibrium* after the buyer offers $p_b \in (\theta_1, \theta_2]$? An important observation is that for every $p_s, p_s' \in \mathbf{P}_s$ with $p_s > p_s'$, offering p_s leads to a higher expected price but a longer expected delay. As a result, the highest offer made by the low type cannot be greater than the lowest offer made by the high type. Let $\overline{p}_s \equiv \max\{\mathbf{P}_s \setminus \{1\}\}$, which is the seller's highest commitment demand below 1. If the high type demands anything p_s that is below $\overline{p}_s$, then the low type will never demand anything more than p_s. This leads to a contradiction since the buyer will concede immediately after the seller demands $\overline{p}_s$, which is a profitable deviation for the low type. In Section 4.1, we rule out equilibria where the high type demands $\overline{p}_s$. This implies that in all equilibria, the high type demands the entire surplus and never concedes to the buyer. The buyer concedes after some delay in order to discourage the low type from imitating the high type.

In summary, the buyer faces a trade-off when she chooses between making a screening offer $p_b \in (\theta_1, \theta_2]$ and a pooling offer p_{θ_2}: Screening reduces the surplus she can extract from the high type but may lower the price she pays to the low type. The latter is true if and only if the difference between θ_1 and θ_2 is large enough. This is because when θ_1 and θ_2 are too close, every screening offer $p_b \in (\theta_1, \theta_2]$ is too low relative to the Rubinstein bargaining price of the low type p_{θ_1}. If $\theta_2 + p_{\theta_2} \leq 2p_{\theta_1}$ or equivalently $\theta_2 - \theta_1 \leq \frac{1-\theta_2}{2}$, then for any $p_b \in (\theta_1, \theta_2]$, the low type can counteroffer something greater than p_{θ_2} and induce the buyer to concede almost immediately, in which case screening is unprofitable for the buyer. This explains the logic behind (3.5). When $\theta_2 - \theta_1 > \frac{1-\theta_2}{2}$, (3.4) is the probability of the low type under which the buyer's benefit from screening equals her cost of screening, so the buyer prefers to screen the seller if and only if $\pi(\theta_1) > \pi^*$.

Expected Delay & Welfare: We pin down the expected delay and social welfare in the inefficient equilibrium using the two types of the seller's incentive constraints. Formally, the expected delay is $\pi(\theta_2)\Big\{1 - \mathbb{E}[e^{-r\tau_b}|\theta = \theta_2]\Big\}$ and the expected welfare loss from delay is $\pi(\theta_2)(1-\theta_2)\Big\{1 - \mathbb{E}[e^{-r\tau_b}|\theta = \theta_2]\Big\}$, both are *decreasing* functions of $\mathbb{E}[e^{-r\tau_b}|\theta = \theta_2]$. In equilibrium, the low-cost type

θ_1 cannot find it profitable to demand 1, which implies an upper bound for $\mathbb{E}[e^{-r\tau_b}|\theta=\theta_2]$:

$$\underbrace{(1-\theta_1)\mathbb{E}[e^{-r\tau_b}|\theta=\theta_2]}_{\text{type }\theta_1\text{'s payoff from demanding 1}} \leq \underbrace{\max\{p_{\theta_1}, 1-\theta_2+\theta_1\}-\theta_1}_{\text{type }\theta_1\text{'s equilibrium payoff}}. \tag{3.7}$$

The high-cost type θ_2 cannot find it profitable to demand $p_s \approx 1$ and then wait for the buyer to concede. Recall the definition of λ_b^1 in (3.3), which is the buyer's concession rate when the low-type seller is conceding. When players' offers are p_b and p_s, let T_1 be the time it takes for type θ_1 to finish conceding and let c_b be the probability with which the buyer concedes at time 0, both of which depend on the buyer's posterior belief. Type θ_2 has no incentive to deviate to $p_s \approx 1$ when

$$\underbrace{(1-\theta_2)\mathbb{E}[e^{-r\tau_b}|\theta=\theta_2]}_{\text{type }\theta_2\text{'s equilibrium payoff}} \geq \underbrace{(p_s-\theta_2)\left(c_b+(1-c_b)\left(1-e^{-(r+\lambda_b^1)T_1}\right)\frac{\min\{p_{\theta_1},\theta_2\}-\theta_1}{p_s-\theta_1}\right)}_{\text{type }\theta_2\text{'s payoff from deviating to }p_s\approx 1}. \tag{3.8}$$

This leads to a lower bound for $\mathbb{E}[e^{-r\tau_b}|\theta=\theta_2]$. In particular, we show in Section 4.1 that, as $p_s \to 1$ and $\varepsilon \to 0$, the right-hand-side of (3.8) is bounded below by

$$\frac{\max\{p_{\theta_1}, 1-\theta_2+\theta_1\}-\theta_1}{1-\theta_1}(1-\theta_2), \tag{3.9}$$

which is attained when the buyer assigns zero probability to the seller having a high cost after observing offers (p_b, p_s). The upper and the lower bounds for $\mathbb{E}[e^{-r\tau_b}|\theta=\theta_2]$ coincide in the limit, which implies that these two bounds pin down the value of $\mathbb{E}[e^{-r\tau_b}|\theta=\theta_2]$.

The equilibrium value of $\mathbb{E}[e^{-r\tau_b}|\theta=\theta_2]$ implies that, compared to the efficient equilibrium, the high-cost seller's payoff is weakly greater in the inefficient equilibrium. Hence, the buyer obtains a higher payoff from screening at the expense of the low-cost seller. That is to say, the low-cost seller not only bears the cost of delay but is also expropriated by the buyer.

Comparative Statics: We apply Theorem 1 to examine how the expected welfare loss and the expected delay of reaching agreement depend on the primitives, such as the distribution of the seller's production cost $\pi(\theta_1)$, his production cost under the new technology θ_1, and that under the default technology θ_2. As in Theorem 1, we focus on the limit where $(\varepsilon, \nu) \to (0,0)$. We start from the effect of an increase in the fraction (or the probability) of low-cost seller:

Corollary 1. *Both the expected welfare loss and the expected delay are zero when $\pi(\theta_1) < \pi^*$, and are strictly decreasing in $\pi(\theta_1)$ when $\pi(\theta_1) \in (\pi^*, 1)$.*

This follows directly from Theorem 1. Intuitively, bargaining is efficient in the limit when $\pi(\theta_1) < \pi^*$. When $\pi(\theta_1) > \pi^*$, bargaining is inefficient only when the seller's production cost is θ_2, and conditional on $\theta = \theta_2$, the expected welfare loss is independent of $\pi(\theta_1)$.

Corollary 2. *Both the expected delay and the expected welfare loss are weakly decreasing in θ_1.*

Intuitively, improving the efficiency of the new technology (i.e., a decrease in θ_1) has two effects, both of which lead to a longer delay. First, a lower θ_1 makes screening more profitable for the buyer, which expands the range of π under which the buyer prefers to make the screening offer. Second, when $\pi(\theta_1) > \pi^*$, the expected delay after the high type offers 1 weakly increases as θ_1 decreases, and strictly increases whenever $\theta_2 \leq p_{\theta_1}$. This is driven by the two incentive constraints that pin down the expected delay: the low type's incentive constraint leads to a lower bound on the expected delay and the high type's incentive constraint leads to an upper bound. According to (3.7) and (3.8), as θ_1 decreases, the low type's payoff from deviation increases and the high type's payoff from deviation decreases. Hence, the expected delay that satisfies both incentive constraints increases.

Corollary 3. *The expected delay is weakly increasing in θ_2. The expected welfare loss is weakly increasing in θ_2 when $\theta_2 \in (\theta_1, p_{\theta_1})$ and is weakly decreasing in θ_2 when $\theta_2 \in (p_{\theta_1}, 1)$.*

Intuitively, improving the efficiency of the default technology (i.e., a decrease in θ_2) has two effects. First, a lower θ_2 makes screening less profitable, which reduces the range of π under which the buyer prefers to make the screening offer. This decreases the expected delay as well as the expected welfare loss from delay. However, there is another effect, namely, a lower θ_2 increases the surplus from trading with type θ_2, which makes each unit of delay more costly in terms of social welfare. Hence, players will reach an agreement sooner when the default technology becomes more efficient, and the expected welfare loss also decreases if and only if θ_2 is lower than p_{θ_1}.

Remarks: We conclude this section with several remarks. First, what will happen when π satisfies $\pi(\theta_1) = \pi^*$? Although the buyer will be indifferent between the pooling offer p_{θ_2} and her optimal screening offer $p_b \in (\theta_1, \theta_2]$ in the limit where $\varepsilon \to 0$, she will have a strict preference under a generic ε. In fact, the cutoff belief π^* will play a crucial role when we analyze the reputational bargaining game with endogenous technology adoption. By choosing an investment probability that is close to π^*, the seller makes the buyer indifferent between p_{θ_2} and her optimal screening offer. The buyer's mixing probabilities are pinned down by the seller's indifference condition at the adoption stage.

Second, in the inefficient equilibrium of our model, the buyer screens the seller by offering a price such that only the low type has an incentive to concede. In contrast, when the only uncertainty is about a player's discount rate, as in APS, there is no offer under which some type has a strict incentive to concede while other types have no incentive to concede. This explains why the uninformed player cannot induce different types of the informed player to offer different prices in APS, but she can do that in our model.

Third, in our model, inefficient delay occurs only when the cost difference between adjacent types is large enough. We generalize this finding in Section 5, showing that inefficient delay occurs *if and only if* there exist two adjacent types whose cost difference is large enough. This contrasts to the efficiency results in Fanning (2022), which require a rich set of production costs. His richness assumption fits applications where the heterogeneity in production cost is driven by differences in human capital, which are continuous in nature. However, when cost heterogeneity arises from differences in production technologies, it is natural to assume that there is a gap between the production costs. This is because decisions on whether to adopt a new technology are usually indivisible, for example, it might be infeasible to adopt half of the technology.

3.2 Reputational Bargaining with Endogenous Technology Adoption

This section analyzes the reputational bargaining game in which the seller's production cost is *endogenously* determined by his adoption decision before the bargaining stage and the buyer cannot observe whether he has adopted. Recall the definition of π^* in (3.4) and that $\pi^* < 1$ if and only if (θ_1, θ_2) satisfies (3.5). If (θ_1, θ_2) also satisfies a stronger condition that $p_{\theta_1} < \theta_2$, then $\pi^* = \frac{1-\theta_2}{1-\theta_1}$.

We state the interesting parts of our equilibrium characterization as our main result, Theorem 2. A full description of the limiting equilibria can be found in Figure 1, with details in Section 4.2.

Theorem 2. *There exists at least one equilibrium of the reputational bargaining game with endogenous technology adoption. For every $\eta > 0$, there exists $\bar{\nu} > 0$ such that when $\nu < \bar{\nu}$, there exists $\bar{\varepsilon}_\nu > 0$ such that for every $\varepsilon \in (0, \bar{\varepsilon}_\nu)$:*

1. *Suppose (θ_1, θ_2) violates (3.5). In every equilibrium, the expected delay is less than η, and the adoption probability is no more than η if $c > \theta_2 - \theta_1$ and is no less than $1 - \eta$ if $c < \theta_2 - \theta_1$.*

2. *Suppose (θ_1, θ_2) satisfies (3.5). There exists an open interval $(\underline{c}, \overline{c}) \subset \left(\frac{\theta_2-\theta_1}{2}, \theta_2 - \theta_1\right)$ such that for every $c \in (\underline{c}, \overline{c})$, there exists an equilibrium where the adoption probability is within an η-neighborhood of π^* and the expected delay is bounded above zero.*

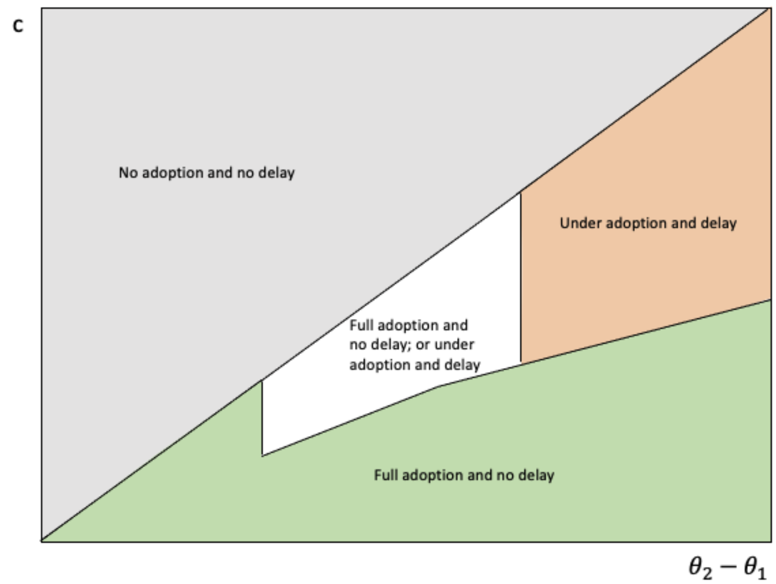

Figure 1: Limiting equilibria in the reputational bargaining with endogenous technology adoption

3. *Suppose (θ_1, θ_2) satisfies $p_{\theta_1} < \theta_2$. If $c \in \left(\frac{\theta_2-\theta_1}{2}, \theta_2 - \theta_1\right)$, then in all equilibria, the adoption probability is within an η-neighborhood of π^* and the expected delay is bounded above zero.*

The proof is in Section 4.2. Theorem 2 implies that in the limit where the set of commitment types is rich and the probability of commitment types is close to 0, the seller's adoption decision can be socially inefficient *if and only if* the social benefit from adoption $\theta_2 - \theta_1$ is large enough. Under a stronger condition that θ_2 is greater than the Rubinstein bargaining price under the low production cost $p_{\theta_1} \equiv \frac{1+\theta_1}{2}$, inefficient adoption and costly delay occur in *all* equilibria. It also implies that inefficient adoption occurs *if and only if* there are significant delays in bargaining.

In order to understand Theorem 2, we start from a heuristic explanation based on our characterization result for reputational bargaining games with an *exogenous* cost distribution (Theorem 1). Then we point out a contradiction that results from this line of reasoning and explain why Theorem 1 cannot be directly applied to settings where the cost distribution is *endogenous*.

Fix any $\pi \in \Delta\{\theta_1, \theta_2\}$ that has full support with $\pi(\theta_1) \neq \pi^*$. Theorem 1 implies that as $\varepsilon \to 0$, the difference between the low-cost seller's equilibrium payoff and that of the high-cost seller's is $\theta_2 - \theta_1$ in every efficient equilibrium and is approximately $(\theta_2 - \theta_1)\alpha$ in every inefficient equilibrium where

$$\alpha \equiv \begin{cases} \frac{1}{2} & \text{if } p_{\theta_1} < \theta_2 \\ \frac{1-\theta_2}{1-\theta_1} & \text{if } p_{\theta_1} \geq \theta_2 \text{ and } (\theta_1, \theta_2) \text{ satisfies (3.5).} \end{cases}$$

Intuitively, in every efficient equilibrium, both types of the seller trade immediately at the same price, in which case the seller captures all the gains from adoption. In every inefficient equilibrium, the high type's payoff is approximately the same as his payoff in the efficient equilibria, so the low type not only bears the welfare losses from delay but also transfers some of his gains to the buyer.

Fix (θ_1, θ_2) and any adoption cost c that is *strictly* between $(\theta_2 - \theta_1)\alpha$ and $\theta_2 - \theta_1$. In equilibrium, the seller cannot adopt with zero probability since the buyer will offer the high price p_{θ_2}, in which case the seller's gain from adoption is $\theta_2 - \theta_1$, providing him a strict incentive to adopt. He cannot adopt with probability 1 since the buyer will then offer p_{θ_1} and type θ_2 seller's payoff is at least $\frac{1-\theta_2}{2}$ when he demands $p_s \approx 1$ and waits for the buyer to concede. The seller's gain from adoption is close to $\frac{\theta_2 - \theta_1}{2}$, in which case he has no incentive to adopt. The seller also cannot adopt with any interior probability that is not π^*, since his gain from adoption will not equal his cost of adoption.

The above logic suggests that in every equilibrium, the seller must adopt with probability π^*. However, when (θ_1, θ_2) violates (3.5), or equivalently $\pi^* = 1$, *all* equilibria are efficient in the game with an exogenous cost distribution, so there does not seem to exist an adoption probability that makes the seller indifferent between adopting and not adopting as long as $(\theta_2 - \theta_1)\alpha < c < \theta_2 - \theta_1$.

The above contradiction is driven by the different orders of limits in Theorems 1 and 2, making Theorem 1 inapplicable to settings where π is endogenous. Specifically, Theorem 1 characterizes the set of equilibria for a *fixed cost distribution* π in the limit where $\varepsilon \to 0$. The same order of limit applies to other results in reputational bargaining with private information about payoffs such as those in Abreu, Pearce and Stacchetti (2015) and Fanning (2022). But in the game with *endogenous* technology adoption, π depends on the probability of commitment types ε. Hence, it could be the case that the probability that $\theta = \theta_2$ also vanishes as $\varepsilon \to 0$.[8]

This is indeed what happens when (θ_1, θ_2) violates (3.5) and $c \in \left(\frac{\theta_2 - \theta_1}{2}, \theta_2 - \theta_1\right)$. We show that for every $\xi > 0$, there exists $\overline{\varepsilon} > 0$ such that the conclusion of Theorem 1 applies for all $\varepsilon < \overline{\varepsilon}$ and π with $\pi(\theta_1) \in [0, 1 - \xi]$. However, the qualitative features of the equilibrium are different for any fixed $\varepsilon > 0$ as $\pi(\theta_1) \to 1$. Although for any fixed $\pi \in \Delta\{\theta_1, \theta_2\}$, the buyer will offer a high price p_{θ_2} and will trade with both types of the seller immediately as $\varepsilon \to 0$, she has an incentive

[8] A related issue appears in Gul (2001) who studies Coasian bargaining with endogenous investments. Unlike Gul, Sonnenschein and Wilson (1986) who fix the distribution of values and then send the frequency of offers to infinity, the informed player's investment probability depends on the frequency of offers in Gul (2001).

to offer a low price p_{θ_1} for any small but fixed ε as $\pi(\theta_1)$ goes to 1. In response to the buyer's offer, the high-cost seller raises the offer and trades with delay, and the low-cost seller accepts the offer with probability close to 1 and pools with the high-cost seller with probability close to 0. The expected delay for the high-cost seller is pinned down by the seller's indifference condition at the adoption stage. This delay has negligible payoff consequences from an ex ante perspective since it is proportional to the probability of $\theta = \theta_2$ and $\pi(\theta_2)$ is arbitrarily close to 0 as $\varepsilon \to 0$.

When $c \in \left(\frac{\theta_2-\theta_1}{2}, \theta_2 - \theta_1\right)$ and (θ_1, θ_2) satisfies not only (3.5) but also a stronger condition that $p_{\theta_1} < \theta_2$, one can no longer sustain an approximately efficient outcome where the seller adopts the technology with probability close to 1. This is because, when the buyer offers p_{θ_1}, the seller has no incentive to concede when his cost is θ_2. As a result, the seller can secure payoff $\frac{1-\theta_2}{2}$ by not adopting the technology and demanding something close to 1. This guaranteed payoff $\frac{1-\theta_2}{2}$ is strictly greater than his payoff from adopting the technology and accepting the buyer's offer p_{θ_1}, which contradicts the hypothesis that the seller adopts the technology with probability close to 1. Therefore, in every equilibrium, the seller will adopt with probability bounded below 1. As $\varepsilon \to 0$, the equilibrium adoption probability is close to π^*, since it is the only adoption probability that can make the buyer indifferent between making the screening offer p_{θ_1} and the pooling offer p_{θ_2}.

When $c \in \left(\frac{\theta_2-\theta_1}{2}, \theta_2 - \theta_1\right)$ and (θ_1, θ_2) satisfies (3.5) but $p_{\theta_1} \geq \theta_2$, there exist inefficient equilibria where the seller adopts with probability close to π^* since Theorem 1 applies uniformly to all π with $\pi(\theta_1)$ bounded below 1. However, there are also efficient equilibria where the seller adopts with probability close to 1. We explain in detail why there are multiple limit points in Section 4.2.

The above explanation also sheds light on why inefficient adoption occurs *if and only if* there are significant delays in bargaining. Intuitively, if the equilibrium adoption decision is bounded away from efficiency, then it cannot be the case that both types of the seller trade immediately. This is because otherwise, both types must trade at the same price and the seller can capture all the gains from adoption, providing him an incentive to make the efficient adoption decision. If the adoption decision is approximately efficient, then the probability with which the seller does not adopt must be arbitrarily close to zero.[9] Since inefficient delay cannot occur when the seller has a low cost, the welfare losses from delay must be negligible from an ex ante perspective.

[9] In our model, inefficient adoption is caused by the hold-up problem. Therefore, the seller will not adopt when his adoption cost is greater than the social benefit from adoption. This implies that inefficiency can only take the form of under-adoption. Thus, the only relevant case to consider is the one in which the seller's adoption cost is strictly less than the social benefit from adoption but he does not adopt with probability bounded above 0.

Comparative Statics: We apply Theorem 2 to study the effects of an *increase in the benefit from adoption* (i.e., a decrease in θ_1) or a *decrease in the cost of adoption* (i.e., a decrease in c) on the equilibrium probability of adoption and on the expected delay of reaching agreement. Similar to Section 3.1, we focus on the limiting scenario where $(\varepsilon, \nu) \to (0,0)$.

One challenge comes from the multiplicity of limiting equilibria when (θ_1, θ_2) satisfies (3.5) but $p_{\theta_1} > \theta_2$, that is, the social benefit from adoption $\theta_2 - \theta_1$ is intermediate. We provide sufficient conditions under which an increase in the benefit from adoption *decreases* the probability of adoption, as well as conditions under which an increase in the benefit from adoption *increases* the expected delay. The conclusions we obtain are robust to the selection of equilibria. Formally, we measure the expected delay according to $1 - \mathbb{E}[e^{-r\min\{\tau_s, \tau_b\}}]$, where τ_s and τ_b are players' concession times.

Corollary 4. *For every θ_1, θ_2, and c that satisfy*

$$\theta_2 - \theta_1 > \frac{1-\theta_2}{2}, \text{ and } \max\Big\{\frac{1}{2}, \frac{1-\theta_2}{1-\theta_1}\Big\}(\theta_2 - \theta_1) < c < \theta_2 - \theta_1,$$

and every $\widehat{\theta}_1 < \theta_1$ that satisfies $\frac{1+\widehat{\theta}_1}{2} < \theta_2$ and $\widehat{\theta}_1 \in (\theta_2 - 2c, \theta_2 - c)$. There exist $\overline{\varepsilon} > 0$ and $\overline{\nu} > 0$ such that for every $\varepsilon < \overline{\varepsilon}$ and $\nu < \overline{\nu}$,

1. *The probability of adoption in any equilibrium under $(\theta_1, \theta_2, c, \varepsilon, \nu)$ is strictly greater than the probability of adoption in any equilibrium under $(\widehat{\theta}_1, \theta_2, c, \varepsilon, \nu)$.*

2. *The expected delay in any equilibrium under $(\theta_1, \theta_2, c, \varepsilon, \nu)$ is strictly less than the expected delay in any equilibrium under $(\widehat{\theta}_1, \theta_2, c, \varepsilon, \nu)$.*

We depict the complete comparative statics with respect to θ_1 and c in Figures 2a and 2b, where the white region represents parameter values under which there are multiple limiting equilibria. Corollary 4 implies that when the production cost under the new technology decreases from θ_1 to $\widehat{\theta}_1$, i.e., adoption becomes more socially beneficial, the probability of adoption decreases as long as $\theta_2 - \widehat{\theta}_1$ is intermediate: It is large enough so that the buyer has an incentive to screen the seller, but is not too large relative to the adoption cost c so that the seller has no incentive to adopt if he knew that the buyer will offer $p_{\widehat{\theta}_1}$. It also implies that a decrease from θ_1 to $\widehat{\theta}_1$ when $\theta_2 - \widehat{\theta}_1$ is intermediate can also lead to longer delays, which leads to further efficiency losses.

If the cost of adoption decreases from c to $\widehat{c}$, except for the parameter values under which there are multiple limit points, the probability of adoption weakly increases, which follows from Theorem 2. The effects on the expected delay is ambiguous. This is because there is no delay when $c > \theta_2 - \theta_1$

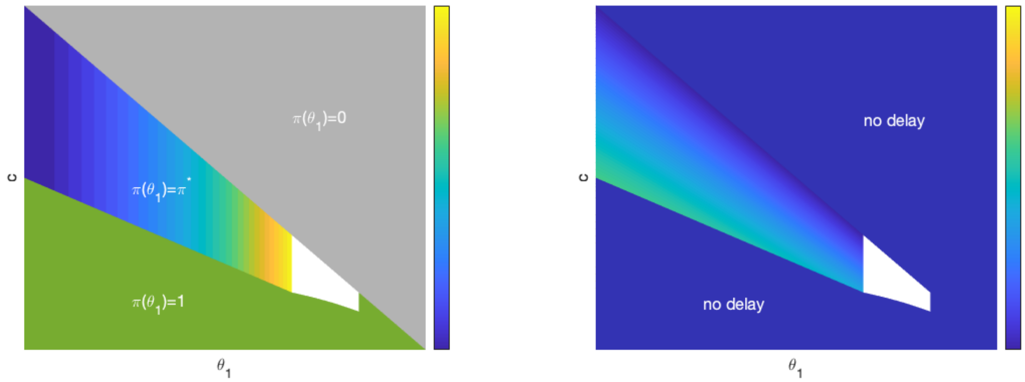

(a) Adoption probability (b) Expected delay

Figure 2: Comparative statics on the equilibrium outcomes. The white region represents parameter configurations under which there are multiple limiting equilibria. In the region in which the unique limiting equilibrium is inefficient, the values of $\pi(\theta_1)$ and $1-\mathbb{E}[e^{-r\min\{\tau_s,\tau_b\}}]$ are depicted in panels (a) and (b), respectively, in ascending order according to the color bar on the right of the figure. In the remaining regions, the adoption probability is fixed at the efficient level, which equals 0 when $c > \theta_2-\theta_1$, and equals 1 when $c < \theta_2-\theta_1$, and the expected delay is zero in the limit.

or when $c < \frac{\theta_2-\theta_1}{2}$, in which cases the seller either never adopts the new technology or adopts the new technology for sure, so there is no delay in trade. In contrast, when $c \in (\frac{\theta_2-\theta_1}{2}, \theta_2-\theta_1)$, there are inefficient equilibria in which the seller adopts with probability strictly between 0 and 1 and the buyer offers p_{θ_1} with positive probability, leading to significant delays in equilibrium.

4 Proofs of Theorems 1 and 2

We prove Theorem 1 in Section 4.1 and Theorem 2 in Section 4.2. Our proofs can be modified to establish the general results in Section 5, which focus on environments where the seller chooses between more than two production technologies with different costs of production.

4.1 Proof of Theorem 1

Our proof proceeds in four steps. First, we describe the equilibrium in the war-of-attrition game under (p_b, p_s). Next, we use the continuation payoffs in the war-of-attrition game to characterize the seller's equilibrium offer after observing the buyer's offer p_b. Then, we use the buyer's sequential rationality to show that she offers either p_{θ_2} or $\min\{p_{\theta_1}, \theta_2\}$, and which one she offers is determined by the comparison between $\pi(\theta_1)$ and the cutoff π^*. These three steps together establish uniqueness of the equilibrium limit point. In Online Appendix A, we establish the existence of equilibrium.

Fix $\pi \in \Delta\{\theta_1, \theta_2\}$. Let $(\sigma_b, \sigma_s, \tau_b, \tau_s)$ be an equilibrium strategy profile of the bargaining game, where σ_b is the distribution of the buyer's initial offer, σ_s is the distribution of the seller's initial offer as a function of his production cost and the buyer's initial offer, and (τ_b, τ_s) is the distribution of players' concession times. We sometimes abuse notation by using $\sigma_b(p_b)$ and $\sigma_s(p_s|\theta, p_b)$ to denote the probability with which σ_b and $\sigma_s(\theta)(p_b)$ assign to offers p_b and p_s, respectively.[10] Let P_b denote the support of σ_b. Let

$$\hat{\varepsilon}_b(p_b) = \frac{\varepsilon \mu_b(p_b)}{\varepsilon \mu_b(p_b) + (1-\varepsilon)\sigma_b(p_b)} \tag{4.1}$$

$$\hat{\varepsilon}_s(p_b, p_s) = \frac{\varepsilon \mu_s(p_s)}{\varepsilon \mu_s(p_s) + (1-\varepsilon)[\pi\{\theta_1\}\sigma_s(p_s|\theta_1, p_b) + \pi\{\theta_2\}\sigma_s(p_s|\theta_2, p_b)]} \tag{4.2}$$

$$\hat{\pi}_j(p_b, p_s) = \frac{(1-\varepsilon)\pi\{\theta_j\}\sigma_s(p_s|\theta_j, p_b)}{\varepsilon \mu_s(p_s) + (1-\varepsilon)[\pi\{\theta_1\}\sigma_s(p_s|\theta_1, p_b) + \pi\{\theta_2\}\sigma_s(p_s|\theta_2, p_b)]}, \quad \text{for every } j \in \{1, 2\}. \tag{4.3}$$

Intuitively, $\hat{\varepsilon}_b(p_b)$ and $\hat{\varepsilon}_s(p_b, p_s)$ are the probabilities with which the buyer and the seller, respectively, are commitment types conditional on their offers (p_b, p_s), and $\hat{\pi}_j(p_b, p_s)$ is the probability assigned to the seller being the rational type and having production cost θ_j.

First, we characterize the equilibrium in the continuation game after players offer $(p_s, p_b) \in \mathbf{P}_s \times \mathbf{P}_b$ with $1 > p_s > p_b > \theta_1$. Fix (p_b, p_s) and the resulting $(\hat{\varepsilon}_b, \hat{\varepsilon}_s, \hat{\pi}_1, \hat{\pi}_2)$. We denote the resulting continuation game by $\Gamma(p_b, p_s, \hat{\varepsilon}_b, \hat{\varepsilon}_s, \hat{\pi})$, and a pair of equilibrium strategies for the buyer and the seller of type θ by $\tau_b \in \Delta(\mathbb{R}_+ \cup \{\infty\})$ and $\tau_s : \Theta \to \Delta(\mathbb{R}_+ \cup \{\infty\})$, respectively. Let

$$m \equiv \max\{j \in \{1, 2\} : \theta_j < p_b\},$$

which is well defined given that $p_b > \theta_1$. Recall that $\lambda_s \equiv \frac{r(1-p_s)}{p_s - p_b}$ is the seller's concession rate. For every $j \in \{1, ..., m\}$, recall that $\lambda_b^j \equiv \frac{r(p_b - \theta_j)}{p_s - p_b}$ is the buyer's concession rate that keeps type θ_j seller indifferent between conceding and waiting. For convenience, let $\lambda_b^{m+1} = \hat{\pi}_{m+1} = 0$.

If the seller of type θ_j doesn't concede at time zero with positive probability and his strategy is to concede at a constant rate equal to λ_s over the interval (T^{j-1}, T^j), with $0 = T^0 \leq T^1 \leq ... \leq T^m$, then the probability with which the buyer's posterior belief assigns to the following event:

- the seller is either committed or has a production cost strictly above θ_j

[10] As will become clear later, this probability is always well defined, since in equilibrium the buyer's and the seller's strategies over bargaining postures are supported on a subset of, respectively, $\mathbf{P}_b$ and $\mathbf{P}_s$ which are finite sets.

reaches 1 at time:

$$T_s^j \equiv \frac{-\log(\hat{\varepsilon}_s + \sum_{i>j} \hat{\pi}_i)}{\lambda_s}. \tag{4.4}$$

Likewise, if the buyer doesn't concede at time zero and concedes at rate λ_b^j over the time interval (T^{j-1}, T^j), then the buyer finishes conceding at time:

$$T_b \equiv \frac{-\log(\hat{\varepsilon}_b) - \sum_{j=1}^{m-1}(\lambda_b^j - \lambda_b^{j+1})T^j}{\lambda_b^m}.$$

In equilibrium, both players must finish conceding at the same time. Therefore, one of them concedes with positive probability at time 0 as long as $T_b \neq T_s^m$. Let

$$L \equiv \frac{-\lambda_s \log \hat{\varepsilon}_b}{-\sum_{j=1}^{m}(\lambda_b^j - \lambda_b^{j+1})\log(\hat{\varepsilon}_s + \hat{\pi}_{j+1})}. \tag{4.5}$$

One can verify that $L < 1$ if and only if $T_b < T_s^m$. Hence, the seller concedes with positive probability at time 0 if and only if $L < 1$ and the buyer concedes with positive probability at time 0 if and only if $L > 1$. We refer to the player who concedes at time zero with strictly positive probability as the *weak* player. In order to derive the probability with which the weak player concedes to their opponent at time zero, let

$$\hat{c}_s^i \equiv 1 - \left(\hat{\varepsilon}_s^{-\lambda_s} \prod_{j=i}^{m} (\hat{\varepsilon}_s + \hat{\pi}_{j+1})^{\lambda_b^j - \lambda_b^{j+1}}\right)^{1/\lambda_b^i} \quad \text{for every } i \in \{1, ..., m\}$$

$$\hat{c}_b = 1 - \hat{\varepsilon}_b \exp\left\{\sum_{j=1}^{m} \lambda_b^j (T_s^j - T_s^{j-1})\right\}.$$

Let $j^* \equiv \min\{j \in \{1, ..., m\} : \hat{c}_s^j < \sum_{i \leq j} \hat{\pi}_i\}$. Suppose first that the buyer is the weak player. Then, $\hat{c}_b$ is the probability with which the buyer concedes at time zero so that the seller's belief that the buyer is the commitment type reaches one at time T_s^m. Likewise, if the seller is the weak player and $j^* = 1$, then $\hat{c}_s^1$ is the probability with which the type θ_1 seller concedes at time 0 so that the buyer's belief that the seller is either committed or that $\theta \geq p_b$ reaches one at time T_b. However, if $j^* > 1$, it is not enough to have the rational type-θ_1 seller conceding with probability one in period zero to make both players finish conceding at the same time. In those cases, we need to have all types strictly below θ_{j^*} conceding at time 0 with probability one and possibly type θ_{j^*} doing so as well with positive probability. As a result, when the seller is the weak player, his concession probability at time 0 equals $\hat{c}_s^{j*}$. Lemma 4.1 summarizes the above findings:

Lemma 4.1. *Fix offers (p_b, p_s) with $1 > p_s > p_b > \theta_1$. In any equilibrium of $\Gamma(p_b, p_s, \hat{\varepsilon}_b, \hat{\varepsilon}_s, \hat{\pi})$, the buyer concedes with positive probability at time zero if and only if $L > 1$ and the seller concedes with positive probability at time 0 if and only if $L < 1$. Players' concession probabilities at time 0 are $c_b \equiv \max\{0, \hat{c}_b\}$ and $c_s \equiv \max\{0, \hat{c}_s^{j^*}\}$, respectively.*

A formal proof of Lemma 4.1 can be found in APS, which we omit in order to avoid repetition.

Next, consider the continuation game after no player has conceded at time 0. In equilibrium, the seller of type θ_j, with $j \in \{j^*, ..., m-1\}$, finishes conceding at time

$$T^j = T_s^j + \frac{\log(1 - c_s)}{\lambda_s}. \tag{4.6}$$

In addition, the rational types of the buyer and the seller finish conceding at the same time, given by:

$$T^m \equiv \min\left\{\frac{-\log(\hat{\varepsilon}_b) - \sum_{j=j^*}^{m-1}(\lambda_b^j - \lambda_b^{j+1})T_s^j}{\lambda_b^m}, T_s^m\right\}. \tag{4.7}$$

Lemma 4.2 completes the characterization of players' strategies in the war-of-attrition game:

Lemma 4.2. *In every equilibrium of the war-of-attrition game $\Gamma(p_b, p_s, \hat{\varepsilon}_b, \hat{\varepsilon}_s, \hat{\pi})$ in which $p_b > \theta_1$ and $p_s < 1$, players' equilibrium concession times τ_b and $\tau_s(\theta)$ must satisfy:*

1. *For every $j \in \{j^*, ..., m\}$, the buyer concedes at rate λ_b^j when $t \in (T^{j-1}, T^j)$ with $T^{j^*-1} = 0$.*

2. *The seller with production cost $\theta \in \{\theta_{j^*}, ..., \theta_m\}$ concedes at rate λ_s when $t \in (T^{j-1}, T^j)$ with $T^{j^*-1} = 0$.*

3. *The seller never concedes if his production cost is strictly greater than θ_m.*

Next, we characterize players' concession probabilities at time 0 in the limit where $\varepsilon \to 0$. Formally, consider an infinite sequence $\{\varepsilon^k\}_{k=0}^{+\infty}$ satisfying $\varepsilon^k \to 0$ as $k \to \infty$. Let (σ_b^k, σ_s^k) be players' equilibrium bargaining strategies when the ex ante probability of commitment types is ε^k. Without loss of generality, we focus on the case where (σ_b^k, σ_s^k) converges to $(\sigma_b^\infty, \sigma_s^\infty)$.[11] Let $(\hat{\varepsilon}_b^k, \hat{\varepsilon}_s^k, \hat{\pi}^k)$ be given by (4.1), (4.2) and (4.3) using $(\varepsilon^k, \sigma_b^k, \sigma_s^k)$, and let $\lim_{k\to\infty} \hat{\pi}_j^k = \hat{\pi}_j^\infty$ for every $j \in \{1, 2\}$ and $\hat{\varepsilon}_i^\infty \equiv \lim_{k\to\infty} \hat{\varepsilon}_i^k$ for every $i \in \{b, s\}$.

[11] This is because otherwise, we can apply the Helly's selection theorem (Billingsley, 2013b), that $\Delta[0, 1]$ is sequentially compact in the topology of weak convergence, and find a converging subsequence and focus on that subsequence.

Lemma 4.3. *Suppose $\{\varepsilon^k\}_{k=1}^{\infty}$ is such that $\varepsilon^k \to 0$ as $k \to \infty$. Let $(c_b^k, c_s^k)_{k=1}^{\infty}$ be given according to Lemma 4.1 in the game $\Gamma(p_b, p_s, \hat{\varepsilon}_b^k, \hat{\varepsilon}_s^k, \hat{\pi}^k)$ with $\theta_1 < p_b < p_s < 1$, and let $(c_b^{\infty}, c_s^{\infty})$ be the limit as $k \to \infty$. Then*

(1) If $\sigma_s^{\infty}(p_s|\theta_1, p_b) + \sigma_s^{\infty}(p_s|\theta_2, p_b) > 0$ and $\lambda_b^2 > \lambda_s$, then $c_s^{\infty}(p_b, p_s) = 1$.

(2) If $\sigma_b^{\infty}(p_b) > 0$, $\hat{\pi}_2^{\infty}(p_s, p_b) > 0$, and $\lambda_s > \lambda_b^2$ or $p_b \leq \theta_2$, then $c_b^{\infty}(p_b, p_s) = 1$.

(3) If $\sigma_b^{\infty}(p_b) > 0$, and $\hat{\varepsilon}_s^{\infty}(p_b, p_s) > 0$ or $\lambda_s > \lambda_b^1$, then $c_b^{\infty}(p_b, p_s) = 1$.

One consequence of Lemma 4.3 is that as long as the buyer assigns strictly positive probability to $\theta = \theta_2$, the identity of the player who concedes at time 0 is determined by the comparison of concession rates between the buyer and the seller with a high production cost.

Next, we use the above results to derive players' equilibrium choices of initial offers in the limit where $\varepsilon \to 0$. Let $\overline{p}_\theta(p_b)$ be the supremum of the support of $\sigma_s(\cdot|p_b, \theta)$. Lemma 4.4 states a key property of the seller's equilibrium strategy.

Lemma 4.4. *For every $\eta > 0$, there exists $\overline{\varepsilon} > 0$ such that, if $\varepsilon < \overline{\varepsilon}$, then $\sigma_s(\cdot|p_b, \theta_2)$ assigns probability less than η to offers below $\overline{p}_{\theta_1}(p_b)$, for any p_b in the support of $\sigma_b(p_b)$.*

Intuitively, Lemma 4.4 says that following any of the buyer's on-path offer, the high-type seller's offer is weakly greater than the low-type seller's offer with probability 1.

Proof. Fix any p_b in the support of $\sigma_b(p_b)$. Without loss of generality, we focus on the case where $p_b \in \mathbf{P}_b \cap (\theta_1, p_{\theta_2}]$, since all other offers are strictly dominated for the buyer. Let p_s be an offer in the support of $\sigma_s(\cdot|p_b, \theta_1)$, and suppose by contradiction that there is $p'_s < p_s$ such that $\sigma_s(\cdot|p_b, \theta_2)$ assigns probability bounded away from zero to p'_s. First, suppose that $p_b > \theta_2$, and let $p_2(p_b) = \max\{p \in \mathbf{P}_s : p \leq 1 + \theta_2 - p_b\}$. If $p'_s > p_2(p_b)$, by part (1) of Lemma 4.3, the high type has to concede immediately to p_b after offering p'_s, and therefore his payoff is $p_b - \theta_2$. If otherwise $p'_s \leq p_2(p_b)$, by part (2) of Lemma 4.3, the buyer concedes with probability close to one after the seller offers p'_s, and thus the high type's payoff is approximately $p'_s - \theta_2$. Let

$$P(p'_s, p_b) \equiv \begin{cases} p_b, & \text{if } p'_s > p_2(p_b) \\ p'_s, & \text{if } p'_s \leq p_2(p_b). \end{cases}$$

We can write the high type's payoff as $P(p'_s, p_b) - \theta_2$.

Suppose $p'_s \neq p_2(p_b)$, and thus $P(p'_s, p_b) < p_2(p_b)$. We argue that it must be the case that type θ_1 offers $p_2(p_b)$ with positive probability, and that type θ_2 does so with probability close to zero. Suppose by way of contradiction that this is not the case. According to parts (2) and (3) of Lemma 4.3, the buyer would concede with probability close to 1 after the seller offers $p_2(p_b)$, which implies that type θ_2 would profit from deviating from p'_s to $p_2(p_b)$. This leads to a contradiction.

For the low type to be willing to offer $p_2(p_b)$ instead of p'_s, it must be the case that $c_b p_2(p_b) + (1 - c_b)p_b - \theta_1 \geq P(p'_s, p_b) - \theta_1$, where c_b is the probability with which the buyer concedes at time 0 after the seller offers $p_2(p_b)$. The high type's payoff from deviating to $p_2(p_b)$ and waiting for the buyer to concede is approximately

$$\left\{c_b + (1 - c_b)\frac{p_b - \theta_1}{p_2(p_b) - \theta_1}\right\}\Big(p_2(p_b) - \theta_2\Big) \geq \frac{P(p'_s, p_b) - \theta_1}{p_2(p_b) - \theta_1}\Big(p_2(p_b) - \theta_2\Big) > P'(p'_s, p_b) - \theta_2.$$

This leads to a contradiction, since the high type strictly benefits from deviating to $p_2(p_b)$. It remains to consider the possibility that $p'_s = p_2(p_b)$. If so, the buyer concedes almost immediately by part (2) of Lemma 4.3. Type θ_1 has to concede immediately when his offer p_s is strictly greater than $p_2(p_b)$. Hence, type θ_1's payoff is $p_b - \theta_1$. This leads to a contradiction, since type θ_1 can profitably deviate by offering $p_2(p_b)$ instead of something strictly greater than $p_2(p_b)$.

Second, suppose that $p_b \in (\theta_1, \theta_2]$. Let $\bar{p}_s = \max\{\mathbf{P}_s \setminus \{1\}\}$. Since $p'_s < p_s \leq 1$, the buyer concedes immediately following p'_s by part (2) of Lemma 4.3. Therefore, type θ_2's payoff is approximately $p'_s - \theta_2$. It cannot be that $p'_s = \bar{p}_s$, since then the low type has to concede to p_b when he makes an offer above $\bar{p}_s$, and therefore, receives a strictly higher payoff by deviating to p'_s. Next, if $p'_s < \bar{p}_s$, then it must be the case that type θ_1 offers $\bar{p}_s$ with strictly positive probability, and that type θ_2 does so with probability close to zero. This is because otherwise, the buyer will concede at time 0 with probability 1 after the seller offers $\bar{p}_s$, in which case type θ_2 has a strictly profitable deviation. This requires that $c_b\bar{p}_s + (1 - c_b)p_b - \theta_1 \geq p'_s - \theta_1$. By deviating to $\bar{p}_s$ and waiting for the buyer to concede, type θ_2's expected payoff is at least

$$\left(c_b + (1 - c_b)\frac{p_b - \theta_1}{\bar{p}_s - \theta_1}\right)(\bar{p}_s - \theta_2) \geq \frac{p'_s - \theta_1}{\bar{p}_s - \theta_1}(\bar{p}_s - \theta_2) > p'_s - \theta_2$$

as $\varepsilon \to 0$. This leads to a contradiction. □

We now describe the seller's equilibrium strategy σ_s. First, we show that when the buyer offers $p_b \in (\theta_2, p_{\theta_2}]$, both types of the seller will offer the same price that is approximately $p_2(p_b)$. Suppose by way of contradiction that type θ_2 offers some $p_s < p_2(p_b)$. According to Lemma 4.4, type θ_1's offer is weakly lower than p_s with probability one. This implies that, if type θ_2 were to deviate to $p_2(p_b)$, then the buyer will concede with probability close to 1 at time 0 by Lemma 4.3. This leads to a contradiction. Similarly, if type θ_2 offers $p_s > p_2(p_b)$ with probability bounded away from zero, she has to concede to p_b at time zero, and can therefore do strictly better by offering $p_2(p_b)$ instead. Since type θ_2 makes an offer close to $p_2(p_b)$ with probability close to one, it immediately follows that type θ_1 finds it optimal to do so as well.

If the buyer offers anything greater than p_{θ_2}, then her payoff is approximately $1 - p_b$. This is because the seller is in a weak bargaining position if he offers anything greater and would therefore have to concede with probability close to 1 in the limit where $\varepsilon \to 0$. Combining this with the above arguments, it follows that any offer $p_b \in \mathbf{P}_b$ such that $p_b > \theta_2$ with $\sigma_b(p_b)$ bounded away from zero must be arbitrarily close to p_{θ_2}, which is the price that maximizes the buyer's payoff $1 - \max\{p_b, 1 - p_2(p_b)\}$ over $p_b \in \mathbf{P}_b \cap (\theta_2, 1]$.

Next, we derive the seller's counteroffer and bound the buyer's equilibrium payoff when she offers $p_b \in \mathbf{P}_b \cap (\theta_1, \theta_2]$. Fix $p_b \in \mathbf{P}_b \cap (\theta_1, \theta_2]$. First, consider the case in which $\sigma_b(p_b)/\varepsilon$ is bounded above zero, so the seller's belief after observing the buyer offering p_b assigns probability bounded below 1 to the buyer being committed. We show that type θ_2's response is to demand the entire surplus with probability one. To see this, first note that it cannot be that he demands $p_s < \bar{p}_s = \max\{\mathbf{P}_s \setminus \{1\}\}$ with positive probability. If he did, then Lemmas 4.3 and 4.4 would again imply that he can profitably deviate to $\bar{p}_s$ which prompts an immediate concession from the buyer. Second, if he demands $\bar{p}_s$ in equilibrium, then Lemma 4.3 implies that the buyer would concede immediately, and therefore both seller types would strictly benefit from pooling at this offer. But then, since ν is small, the buyer's payoff when she offers p_b is close to zero. This contradicts $\sigma_b(p_b) > 0$, since the buyer can ensure a strictly positive payoff by offering p_{θ_2}.

Therefore, type θ_2 demands the entire surplus after the buyer offers $p_b \leq \theta_2$. However, it cannot be that type θ_1 pools with type θ_2 by demanding 1 as well. If type θ_1 demands 1 as well, the fact that $p_b > \theta_1$ implies that, in equilibrium, the buyer would wait until type θ_1 concedes before making a concession, in which case it is unprofitable for type θ_1 to demand 1. Therefore, type θ_1's offer is different from type θ_2's, so it is optimal for him to demand $\max\{p_b, p_1(p_b)\}$, where $p_1(p_b) \equiv \max\{p'_s \in \mathbf{P}_s : p'_s \leq 1 - (p_b - \theta_1)\}$.

Next, if p_b is such that $\sigma_b(p_b)/\varepsilon \approx 0$, then the seller's belief after observing the buyer offers p_b will assign probability close to 1 to the buyer being committed. This will lead to immediate concession by type θ_1, and thus a payoff of approximately $\pi(\theta_1)(1-p_b)$ for the buyer. In summary, the buyer's payoff when she offers $p_b \in \mathbf{P}_b \cap (\theta_1, \theta_2]$ is approximately $\pi(\theta_1)(1 - \max\{p_b, p_1(p_b)\})$, which is maximized at $p_b \approx \min\{p_{\theta_1}, \theta_2\}$.

Hence, there are two candidate offers that the buyer will make with probability bounded above zero in equilibrium: p_{θ_2} or $\min\{p_{\theta_1}, \theta_2\}$. Moreover, if $\pi(\theta_1) < \pi^*$, the buyer strictly prefers offering p_{θ_2} over $\min\{p_{\theta_1}, \theta_2\}$, and thus limiting equilibrium strategies in this case are characterized by the first part in Theorem 1. The resulting equilibrium outcome is approximately efficient, with an agreement being reached with negligible delay at a price arbitrarily close to p_{θ_2}.

If $\pi(\theta_1) > \pi^*$, the buyer strictly prefers to make the screening offer $\min\{p_{\theta_1}, \theta_2\}$, and therefore, the equilibrium strategies must be close to those described in part two of Theorem 1. If $p_{\theta_1} \leq \theta_2$, the buyer offers p_{θ_1} and type θ_1 accepts. Otherwise, the buyer offers θ_2 in equilibrium, after which type θ_1 raises the price to $p_1(\theta_2) \approx 1 - (\theta_2 - \theta_1) > \theta_2$ and triggers a war-of-attrition.

However, Lemma 4.3 shows that the expected delay in the resulting war-of-attrition vanishes as $\varepsilon \to 0$ and thus the equilibrium outcome, conditional on the seller's cost is θ_1, is approximately efficient. If the seller's cost is θ_2, he will respond to the buyer's equilibrium offer by demanding the entire surplus 1. In order to deter a deviation from the low type, the buyer must wait a considerable amount of time before conceding to this demand. As argued in Section 3.1, the incentive constraints of the seller pin down the expected welfare loss from delay to be approximately given by (3.6). To complete the argument provided there, we show how to derive (3.9) from (3.8). In order to avoid the buyer from conceding immediately after the seller demands $\overline{p}_s$, which would in turn give rise to a profitable deviation for the seller, it must be that type θ_1 demands $\overline{p}_s$ with positive probability (by Lemma 4.3). As a result, the low type's incentive compatibility requires that $c_b\overline{p}_s + (1-c_b)\min\{p_{\theta_1}, \theta_2\} - \theta_1 \approx \max\{p_{\theta_1}, 1 - \theta_2 + \theta_1\} - \theta_1$. Plugging this into (3.8) and using the fact that $T_1 \to +\infty$ as $\varepsilon \to 0$, we obtain (3.9).

We just showed that every equilibrium must satisfy the properties stated in Theorem 1. In order to complete the proof of Theorem 1, it remains to establish the existence of at least one equilibrium. We relegate the existence proof to Online Appendix A.

4.2 Proof of Theorem 2

First, we study the equilibria of a reputational bargaining game in which ν and ε are fixed and the probability with which $\theta = \theta_2$ goes to zero. We characterize the limiting equilibrium strategies of the low-cost seller and the buyer. Since the probability of the high type vanishes, these strategies are sufficient to pin down limiting equilibrium outcomes. Let $\sigma_{b,1}$ be the buyer's strategy that assigns probability 1 to $p_{\theta_1} \equiv \frac{1+\theta_1}{2}$. Let

$$\sigma_{s,1} \equiv \begin{cases} 1, & \text{if } p_b \leq \theta_1 \\ \max\{p_b, 1+\theta_1 - p_b\}, & \text{if } p_b > \theta_1. \end{cases}$$

Lemma 4.5 characterizes limiting equilibrium strategies of the buyer as well as the low-type seller, for an exogenous sequence of π with $\pi(\theta_1)$ converging to 1.

Lemma 4.5. *For every $\eta > 0$, there exist $\bar{\varepsilon}, \bar{\nu} > 0$ such that when $\varepsilon < \bar{\varepsilon}$ and $\nu < \bar{\nu}$, there exists $\bar{\pi}_{\varepsilon,\nu} \in (0, 1)$ such that for every $\pi(\theta_1) \in (\bar{\pi}_{\varepsilon,\nu}, 1]$: σ_b is η-close to $\sigma_{b,1}$, $\sigma_s(\theta_1)$ is η-close to $\sigma_{s,1}$ on the equilibrium path, and the expected delay of reaching an agreement is no more than η.*

Proof. Fix some small $\varepsilon > 0$ and $\nu > 0$. First, we describe the low-type seller's offer following any offer p_b in the support of the buyer's strategy. Without loss of generality, we can restrict attention to $p_b \in \mathbf{P}_b \cap (\theta_1, p_{\theta_2}]$, since offers that do not belong to this set are strictly suboptimal for the buyer. Suppose first that $p_b \in \mathbf{P}_b \cap (p_{\theta_1}, p_{\theta_2}]$. Regardless of the strategy of the high-type seller, for a fixed small value of ε and $\pi(\theta_1)$ that is sufficiently close to 1, Lemma 4.1 implies that type-θ_1 seller will concede at time 0 with probability close to 1 if he demands anything greater than p_b with probability bounded away from zero. Therefore, the buyer's payoff when she offers $p_b > p_{\theta_1}$ is arbitrarily close to $1 - p_b$ when $\pi(\theta_1)$ is close enough to 1.

If $p_b \in \mathbf{P}_b \cap (\theta_1, p_{\theta_1}]$, then it must be the case that type-θ_1 seller responds by offering $p_1(p_b)$ with probability close to one. Such an offer leads the buyer to concede at time 0 with probability close to 1 when ε is small, so any offer below $p_1(p_b)$ is strictly dominated. If he were to offer $p'_s > p_1(p_b)$ with probability bounded above zero, then for a fixed small value of ε, if $\pi(\theta_1)$ is sufficiently close to 1, the low-type seller will concede at time 0 with positive probability given the conclusion of Lemma 4.1. This implies that his payoff when he offers p'_s is approximately $p_b - \theta_1$, which is strictly lower than his payoff from offering $p_1(p_b)$.

Therefore, when $\pi(\theta_1)$ is sufficiently close to 1, the buyer's payoff from any offer $p_b \in \mathbf{P}_b \cap (\theta_1, p_{\theta_2}]$ is approximately $1 - \max\{p_b, p_1(p_b)\}$, which is maximized at $p_b \approx p_{\theta_1}$. This implies that in equilibrium, the buyer will offer p_{θ_1} with probability close to 1 when $\pi(\theta_1)$ is close enough to 1. The above arguments then imply that, on the equilibrium path, the type-θ_1 seller will counteroffer p_{θ_1} with probability close to 1. This in turn implies that expected delay is negligible. □

We apply the conclusions of Lemma 4.5 and Theorem 1 to describe players' limiting equilibrium strategies in the bargaining game with endogenous technology adoption. Throughout, we use V_θ to denote the equilibrium payoff type θ net of the adoption costs, and we use $\pi(\theta_1)$ to denote the seller's equilibrium adoption probability. The following series of lemmas provides necessary conditions for the limiting equilibria under every parameter configuration.

First, consider the case in which $c > \theta_2 - \theta_1$, in another word, the cost of adoption strictly exceeds the social benefit from adoption. Obviously, the seller adopts with zero probability in every equilibrium. This in turn implies that the buyer has no incentive to offer anything strictly below θ_2. As a result, she will offer $p_{\theta_2} \equiv \frac{1+\theta_2}{2}$ in equilibrium, which is summarized as Lemma 4.6.

Lemma 4.6. *If $c > \theta_2 - \theta_1$, then for every $\eta > 0$, there exists $\bar{\nu} > 0$ such that when $\nu < \bar{\nu}$, there exists $\bar{\varepsilon}_\nu > 0$ such that for every $\varepsilon \in (0, \bar{\varepsilon}_\nu)$, the adoption probability and the expected delay are less than η in any equilibrium.*

Proof. Suppose $c > \theta_2 - \theta_1$. Suppose by way of contradiction that the seller adopts with probability greater than η. If $\pi(\theta_1) < \pi^*$, then the equilibrium characterization in Theorem 1 applies and there is almost immediate trade at a price arbitrarily close to p_{θ_2}. Thus, $V_{\theta_1} - c \approx p_{\theta_2} - \theta_1 - c < p_{\theta_2} - \theta_2 \approx V_{\theta_2}$, which contradicts $\pi(\theta_1) > 0$. If $\pi(\theta_1) > \pi^*$, then either the second part of Theorem 1 or Lemma 4.5 applies, and the low type trades immediately at a price p^*, where $p^* = \max\{p_{\theta_1}, 1 + \theta_1 - \theta_2\}$ if $\pi(\theta_1)$ is bounded away from 1, and $p^* = p_{\theta_1}$ if $\pi(\theta_1)$ is sufficiently close to 1. However, the seller can always deviate to not adopt and trade immediately at the price p^*, which yields a payoff of $p^* - \theta_2 > p^* - \theta_1 - c$, a contradiction. □

In the rest of this proof, we assume that $c < \theta_2 - \theta_1$. Lemma 4.7 examines the case in which $p_{\theta_1} < \theta_2$ and $c \in \left(\frac{\theta_2 - \theta_1}{2}, \theta_2 - \theta_1\right)$. The first condition states that the gap between θ_2 and θ_1 is large enough so that the buyer benefits from screening the seller under certain values of $\pi(\theta_1)$. The second condition implies that the adoption cost is neither too high nor too low, so as to ensure that the seller is willing to mix at the adoption stage with probabilities that make the buyer indifferent between the screening offer p_{θ_1} and the pooling offer p_{θ_2}. The buyer's mixing probabilities over p_{θ_1} and p_{θ_2} are chosen in order to make the seller indifferent at the adoption stage. We show in Lemma 4.7 that the limiting equilibrium outcome is unique under these two conditions.

Lemma 4.7. *If $p_{\theta_1} < \theta_2$ and $c \in \left(\frac{\theta_2 - \theta_1}{2}, \theta_2 - \theta_1\right)$, then for every $\eta > 0$, there exists $\bar{\nu} > 0$ such that when $\nu < \bar{\nu}$, there exists $\bar{\varepsilon}_\nu > 0$ such that for every $\varepsilon \in (0, \bar{\varepsilon}_\nu)$, the adoption probability is η-close from π^* and the expected delay is bounded above zero.*

Proof. Suppose that $p_{\theta_1} < \theta_2$ and $c \in \left(\frac{\theta_2-\theta_1}{2}, \theta_2 - \theta_1\right)$. Suppose by contradiction that $\pi(\theta_1)$ is bounded away from π^*. If $\pi(\theta_1) > \pi^*$, then Lemma 4.5 and Theorem 1 imply that the seller's payoff when he adopts is $V_{\theta_1} - c \approx p_{\theta_1} - \theta_1$. If he instead chooses the default technology θ_2, by demanding the entire surplus after the buyer offers $p_{\theta_1} < \theta_2$, he can ensure a payoff of approximately $\frac{1-\theta_2}{2}$. This is strictly greater than $p_{\theta_1} - \theta_1 - c$, whenever $c > \frac{\theta_2-\theta_1}{2}$, which contradicts $\pi(\theta_1) > 0$.

Next, suppose that $\pi(\theta_1) < \pi^*$. Then, our characterization in Theorem 1 implies that $V_{\theta_1} - c \approx p_{\theta_2} - \theta_1 - c > p_{\theta_2} - \theta_2$, where the inequality follows from $c < \theta_2 - \theta_1$. This contradicts $\pi(\theta_1) < 1$.

Thus, in any equilibrium, the seller must adopt with probability close to π^*. This adoption rule guarantees that the buyer is indifferent between offering p_{θ_1} and p_{θ_2}. After the buyer offers p_{θ_2}, trade happens with almost no delay at this price. After she offers p_{θ_1}, there is trade with negligible delay conditional on the seller being low type, and there is an expected delay of approximately $\frac{1}{2}$ (see (3.6)) when the seller's type is high.

In order to ensure that the seller is indifferent at the adoption stage, it must be that the buyer's strategy over bargaining postures satisfies

$$\rho^*(p_{\theta_2} - \theta_1) + (1-\rho^*)(p_{\theta_1} - \theta_1) - c \approx \rho^*(p_{\theta_2} - \theta_2) + (1-\rho^*)\frac{1}{2}(1-\theta_2) \iff \rho^* \approx \frac{2c - (\theta_2 - \theta_1)}{\theta_2 - \theta_1}$$

Where ρ^* is the probability that the buyer offers (approximately) p_{θ_2}. Observe that $\rho^* \in (0,1)$ if and only if $c \in \left(\frac{\theta_2-\theta_1}{2}, \theta_2 - \theta_1\right)$. This implies that an equilibrium with adoption probability close to π^* cannot be sustained if $p_{\theta_1} < \theta_2$ and the adoption cost is outside of this region. This fact will be invoked in the proof of Lemma 4.9. The expected delay is approximately

$$(1-\pi^*)(1-\rho^*)\frac{1}{2} \approx \frac{\theta_2 - \theta_1 - c}{1 - \theta_1} > 0.$$

□

Lemma 4.7 establishes the third part of Theorem 2. To complete the proof of the second part, it remains to consider the case in which (3.5) is satisfied but $p_{\theta_1} > \theta_2$, or equivalently $\theta_2 - \theta_1 \in \left(\frac{1-\theta_2}{2}, 1 - \theta_2\right)$, and the adoption cost intermediate. We characterize limiting equilibria in this case in Lemma 4.8.

Lemma 4.8. *If (θ_1, θ_2) satisfies (3.5), $p_{\theta_1} > \theta_2$, and $c \in \left(\frac{(1-\theta_2)(\theta_2-\theta_1)}{1-\theta_1}, \theta_2 - \theta_1\right)$, then for every $\eta > 0$, there exists $\bar{\nu} > 0$ such that when $\nu < \bar{\nu}$, there exists $\bar{\varepsilon}_\nu > 0$ such that for every $\varepsilon \in (0, \bar{\varepsilon}_\nu)$, in any equilibrium,*

- *either the adoption probability is η-close to π^* and the expected delay is bounded above zero,*
- *or the adoption probability is greater than $1-\eta$ and the expected delay is less than η.*

Proof. Suppose that (θ_1, θ_2) satisfies (3.5), $p_{\theta_1} > \theta_2$, and $c \in \left(\frac{(1-\theta_2)(\theta_2-\theta_1)}{1-\theta_1}, \theta_2 - \theta_1\right)$. Suppose first that $\pi(\theta_1) < \pi^*$. Then, by Theorem 1, $V_{\theta_1} - c \approx p_{\theta_2} - \theta_1 - c > p_{\theta_2} - \theta_2 \approx V_{\theta_2}$, which contradicts $\pi(\theta_1) < 1$. Next, suppose that $\pi(\theta_1) \in (\pi^*, 1)$ bounded away from both endpoints. Theorem 1 implies that $V_{\theta_1} - c \approx (1-\theta_2) - c < \frac{1-\theta_2}{1-\theta_1}(1-\theta_2) \approx V_{\theta_2}$, where the inequality follows from the assumption that $c > \frac{(1-\theta_2)(\theta_2-\theta_1)}{1-\theta_1}$. This is in contradiction with the hypothesis that $\pi(\theta_1) > 0$.

Thus, the equilibrium adoption probability must be close to either π^* or 1. If it is approximately π^*, then it must be that the buyer mixes between offering θ_2 and p_{θ_2} in a way that makes the seller indifferent at the adoption stage. Similar to Lemma 4.7, this requires:

$$\rho^*(p_{\theta_2} - \theta_1) + (1-\rho^*)(1-\theta_2) - c \approx \rho^*(p_{\theta_2} - \theta_2) + (1-\rho^*)\frac{1-\theta_2}{1-\theta_1}(1-\theta_2)$$
$$\iff \rho^* \approx \frac{(1-\theta_1)c - (1-\theta_2)(\theta_2-\theta_1)}{(\theta_2-\theta_1)^2}$$

With $\rho^* \in (0,1) \iff c \in \left(\frac{(1-\theta_2)(\theta_2-\theta_1)}{1-\theta_1}, \theta_2 - \theta_1\right)$, as assumed in the lemma. Therefore, the expected delay is approximately

$$(1-\pi^*)(1-\rho^*)\frac{1-\theta_2}{1-\theta_1} \approx \frac{(3\theta_2 - 1 - 2\theta_1)(\theta_2 - \theta_1 - c)}{2(\theta_2-\theta_1)^2} > 0.$$

Next, consider the case in which $\pi(\theta_1)$ is arbitrarily close to 1. By Lemma 4.5, the buyer offers p_{θ_1} and $V_{\theta_1} \approx p_{\theta_1} - \theta_1$. Conditional on not adopting, the seller counteroffers $p_2(p_{\theta_1})$. In order to deter the low type from deviating from p_{θ_1} to $p_2(p_{\theta_1}) > p_{\theta_1}$, it must be that the buyer concedes with zero probability after the seller offers $p_2(p_{\theta_1})$. This condition pins down the (very small) probability that the low type offers $p_2(p_{\theta_1})$ in equilibrium, which we denote by β.

Let $T_1 \in \mathbb{R}_+$ denote the time that the low-type finishes conceding after players' offers are $(p_{\theta_1}, p_2(p_{\theta_1}))$. As $\varepsilon \to 0$, type-θ_2 seller's equilibrium payoff, denoted by V_{θ_2}, is approximately

$$V_{\theta_2} \approx (1 - e^{-(r+\lambda_b^1)T_1})\frac{(1-p_{\theta_1})^2}{1+\theta_2 - p_{\theta_1} - \theta_1} \leq \frac{(1-p_{\theta_1})^2}{1+\theta_2 - p_{\theta_1} - \theta_1}$$

Since the seller must be indifferent between adopting and not adopting the new technology, the time at which type θ_1 finishes conceding, T_1, must be such that $V_{\theta_1} - c = V_{\theta_2}$. In equilibrium,

$$T_1 = \frac{-\log\left(\frac{\varepsilon\mu_s(p_2(p_{\theta_1}))+(1-\varepsilon)(1-\pi(\theta_1))}{\varepsilon\mu_s(p_2(p_{\theta_1}))+(1-\varepsilon)(1-\pi(\theta_1)+\pi(\theta_1)\beta)}\right)}{\lambda_s}$$

The condition can be satisfied by choosing $\pi(\theta_1)$ accordingly. Since $c > \frac{(1-\theta_2)(\theta_2-\theta_1)}{1-\theta_1} > \frac{(1-p_{\theta_1})^2}{1+\theta_2-p_{\theta_1}-\theta_1}$, it must be that T_1 is bounded above, for if not the seller would strictly prefer to not adopt. This in turn requires that $\pi(\theta_1) \approx 1$, but not equal to 1.

On the other hand, the type-θ_2 seller can always ensure a payoff of $p_{\theta_1} - \theta_2$, which implies that $V_{\theta_2} \geq p_{\theta_1} - \theta_2 \approx V_{\theta_1} - (\theta_2 - \theta_1)$. Therefore, if moreover $c < \theta_2 - \theta_1$, there exists an equilibrium in which the seller adopts the technology with probability 1 under these parameter values.

Finally, because the seller adopts with probability close to 1 and there is negligible delay conditional on the seller adopting, the expected delay is less than η when ε is small enough. □

Lemma 4.8 has two implications. First, as stated in Theorem 2, when (3.5) holds and $\theta_2 < p_{\theta_1}$, there is an open set of production costs, given by $\left(\frac{(1-\theta_2)(\theta_2-\theta_1)}{1-\theta_1}, \theta_2 - \theta_1\right) \subset \left(\frac{\theta_2-\theta_1}{2}, \theta_2 - \theta_1\right)$ such that there exists an equilibrium with inefficient adoption and delay in bargaining. This equilibrium shares the same features as the unique equilibrium found in Lemma 4.7.

Second, there is equilibrium multiplicity in this region. In particular, there also exists an approximately efficient equilibrium where the seller adopts with probability close to one and delay is negligible. What makes it possible to sustain the efficient outcome when $\theta_2 < p_{\theta_1}$ and the cost of adoption is high (e.g., it is arbitrarily close to $\theta_2 - \theta_1$)? The answer is that, whenever the buyer expects the seller to adopt with probability close to one, her optimal strategy is to offer p_{θ_1}, as she would in the game in which the seller's type were known to be equal to θ_1. Because $\theta_2 < p_{\theta_1}$, the seller cannot commit to wait indefinitely if he chooses to not adopt. Thus, conditional on not adopting, the seller will concede eventually, and the amount of delay with which he reaches an agreement ensures that he is willing to adopt with probability close to 1. Moreover, the fact that $\pi(\theta_1)$ is close to one ensures that this delay can be sustained as an outcome of the war of attrition. However, this reasoning does not apply when $\theta_2 > p_{\theta_1}$. This is because, conditional on not adopting, the high-type seller is able to commit to never concede after the buyer offers p_{θ_1}. This will in turn drive up his non-adoption payoff, giving rise to a profitable deviation from $\pi(\theta_1) > 0$.

Finally, we consider the case in which either (3.5) is violated or the cost of adoption is sufficiently low. In this case, we have efficient investment and almost no delay in bargaining.

Lemma 4.9. *If the parameters of the model satisfy (3.5) and* $c < \max\left\{\frac{1}{2}, \frac{1-\theta_2}{1-\theta_1}\right\}(\theta_2 - \theta_1)$*, or if (3.5) is violated and* $c < \theta_2 - \theta_1$*, then for every* $\eta > 0$*, there exists* $\bar{\nu} > 0$ *such that when* $\nu < \bar{\nu}$*, there exists* $\bar{\varepsilon}_\nu > 0$ *such that for every* $\varepsilon \in (0, \bar{\varepsilon}_\nu)$*, the adoption probability is greater than* $1 - \eta$ *and the expected welfare loss from delay is less than* η*.*

Proof. Suppose first that (3.5) is violated and $c < \theta_2 - \theta_1$. If the seller adopts with probability strictly less than one, then by Theorem 1 we have that $V_{\theta_1} - c \approx p_{\theta_2} - \theta_1 - c > p_{\theta_2} - \theta_2 \approx V_{\theta_2}$, where the inequality follows from the assumption that $c < \theta_2 - \theta_1$. This contradicts $\pi(\theta_1) < 1$.

Next, suppose that (3.5) is satisfied and $c < \max\left\{\frac{1}{2}, \frac{1-\theta_2}{1-\theta_1}\right\}(\theta_2 - \theta_1)$. If $\pi(\theta_1) < \pi^*$, then $V_{\theta_1} - V_{\theta_2} = \theta_2 - \theta_1 - c$, and the fact that $c < \theta_2 - \theta_1$ implies that the seller would strictly prefer to adopt, contradicting $\pi(\theta_1) < 1$. Moreover, our proofs of Lemmas 4.7 and 4.8 showed that an inefficient equilibrium with adoption probability close to π^* exists only if $c > \max\left\{\frac{1}{2}, \frac{1-\theta_2}{1-\theta_1}\right\}(\theta_2 - \theta_1)$, so this type of equilibrium is ruled out under the asusmption in Lemma 4.9.

If otherwise $\pi(\theta_1) \in (\pi^*, 1)$ is bounded away from both endpoints, then by Theorem 1 we have that $V_{\theta_1} - c \approx \max\{p_{\theta_1}, 1 + \theta_1 - \theta_2\} - \theta_1 - c > \max\left\{\frac{1}{2}, \frac{1-\theta_2}{1-\theta_1}\right\}(1 - \theta_2) \approx V_{\theta_2}$, where the inequality follows from the fact that $c < \max\left\{\frac{1}{2}, \frac{1-\theta_2}{1-\theta_1}\right\}(\theta_2 - \theta_1)$. Thus, the seller must adopt with probability arbitrarily close to 1 in this case as well.

Finally, by Lemma 4.5, the outcome is approximately efficient conditional on the seller adopting, and thus the expected delay is arbitrarily small as ε vanishes. □

Lemmas 4.6 and 4.9 together imply the first part of Theorem 2. The above argument establishes the common properties of all equilibria. What remains to be shown is the existence of at least one equilibrium, which is relegated to Online Appendix B.

5 Extension: Choosing Between Multiple New Technologies

This section extends our theorems to settings where the seller chooses a production technology from $\{1, 2, ..., n\}$ before bargaining with the buyer, where θ_j stands for the production cost of technology j and c_j stands for the cost of adopting technology j. Let $\Theta \equiv \{\theta_1, ..., \theta_n\}$ and $C \equiv \{c_1, ..., c_n\}$.

We assume that $0 < \theta_1 < ... < \theta_n < 1$ and $c_1 > ... > c_n = 0$. This implies that (i) there exists a default technology θ_n that is costless to adopt, (ii) all new technologies $\theta_1, ..., \theta_{n-1}$ are costly to adopt but lead to lower production costs compared to the default one, and (iii) technologies that have higher adoption costs have lower production costs. The first two assumptions are without loss of generality. The third assumption rules out technologies that are strictly dominated, which will never be adopted in any equilibrium. We focus on the case where there is a unique *socially efficient technology* and that the efficient technology is not the default one, that is, there exists $j^o < n$ such that $\{j^o\} = \arg\min_{k\in\{1,2,\dots,n\}} \Big\{\theta_k + c_k\Big\}$.

First, we consider a reputational bargaining game where the distribution over production cost $\pi \in \Delta(\Theta)$ is exogenous. Theorem 3 characterizes players' equilibrium strategies in the limit as ν and ε go to zero. Let $\sigma^*_{b,i} \in \Delta[0,1]$ denote the buyer's strategy of offering $\min\{p_{\theta_i}, \theta_{i+1}\}$. Let

$$\sigma^*_{s,\theta}(p_b) \equiv \begin{cases} 1, & \text{if } p_b \leq \theta, \\ \max\Big\{p_b, 1+\theta_j - p_b\Big\}, & \text{if } p_b > \theta \text{ and } \theta_j = \max\{\hat{\theta} \in \Theta : p_b > \hat{\theta}\} \end{cases} \tag{5.1}$$

be a strategy for type θ. That is, for every $p_b > \theta_1$ and $\theta_j = \max\{\hat{\theta} \in \Theta : p_b > \hat{\theta}\}$, $\sigma^*_s \equiv (\sigma^*_{s,\theta})_{\theta\in\Theta}$ prescribes all types with production cost strictly greater than θ_j to demand the entire surplus 1, and all types with production cost no more than θ_j to offer a price under which the buyer and the seller have the same concession rate when the seller's production cost is known to be θ_j.

Theorem 3 characterizes the unique limiting equilibrium of the reputational bargaining game with an exogenous cost distribution under the generic conditions that

$$\arg\max_{i\in\{1,\dots,n\}} \ \pi[\theta_1, \theta_i]\Big(\min\{p_{\theta_i}, \theta_{i+1}\} - \theta_i\Big)$$

is a singleton with its unique element denoted by i^*, and that the cost distribution π is interior.

Theorem 3. *There exists at least one equilibrium of the reputational bargaining game with exogenous production costs. Suppose $\pi \in \Delta(\Theta)$ is such that $\pi(\theta) > 0$ for all $\theta \in \Theta$. For every $\eta > 0$, there exists $\bar{\nu} > 0$ such that when $\nu < \bar{\nu}$, there exists $\bar{\varepsilon}_\nu > 0$ such that for every $\varepsilon \in (0, \bar{\varepsilon}_\nu)$ and every equilibrium $(\sigma_s, \sigma_b, \tau_s, \tau_b)$ under (ε, ν),*

1. *σ_b is η-close to σ_{b,i^*} and σ_s is η-close to σ^*_s on the equilibrium path.*

2. *Conditional on $\theta \leq \theta_{i^*}$, the expected welfare loss from delay is less than η.*

3. *Conditional on $\theta > \theta_{i^*}$, the buyer's equilibrium payoff is 0 and the expected welfare loss from delay is η-close to*

$$(1-\theta)\left\{1 - \frac{\max\Big\{p_{\theta_{i^*}}, 1-(\theta_{i^*+1} - \theta_{i^*})\Big\} - \theta_{i^*}}{1-\theta_{i^*}}\right\}. \tag{5.2}$$

The proof is in Online Appendix C, which is similar to the one for Theorem 1. The structure of the equilibrium is analogous to the two-type case analyzed in Theorem 1. By making an offer p_b that belongs to $(\theta_i, \theta_{i+1}]$ with $i \in \{1, ..., n-1\}$, the buyer is able to screen the seller by providing incentives to all types with cost weakly lower than θ_i to trade with negligible delay, and all types with cost strictly greater than θ_i to separate and demand the entire surplus. Using the same arguments as those in Section 3, the optimal way in which the buyer can screen types with cost no more than θ_i is by offering $\min\{p_{\theta_i}, \theta_{i+1}\}$. This ensures the buyer a payoff of $\pi[\theta_1, \theta_i](\min\{p_{\theta_i}, \theta_{i+1}\} - \theta_i)$.

By construction, the buyer will screen the seller in equilibrium if and only if $i^* < n$, in which case she will offer a price weakly below θ_n instead of offering $p_{\theta_n} \equiv \frac{1+\theta_n}{2}$. Conditional on the buyer offering $\min\{p_{\theta_{i^*}}, \theta_{i^*+1}\} \leq \theta_n$, there will be inefficient delay whenever the seller's production cost satisfies $\theta > \theta_{i^*}$. This is because delay is necessary in order to satisfy the low-cost types' incentive constraints, which is in turn necessary for all types with cost greater than θ_{i^*} demanding the entire surplus to be an equilibrium outcome. The expected delay in (5.2) is pinned down by the conditions which ensure that, after the buyer offers $\min\{p_{\theta_{i^*}}, \theta_{i^*+1}\}$: (i) type θ_{i^*} doesn't benefit from deviating to demanding 1, and (ii) type θ_{i^*+1} doesn't profit from deviating to making an offer slightly below 1 and waiting for the buyer to concede.

In the complementary case in which $i^* = n$, the buyer cannot profit from screening and as a result, she prefers to make the pooling offer p_{θ_n}. As in Section 3, this leads to an efficient outcome in which trade happens immediately at price approximately p_{θ_n}.

Next, we describe the limiting equilibria in the game with endogenous investment. Under the assumption that $j^o < n$, the investment and the bargaining outcome may be inefficient provided that the buyer finds it beneficial to screen the seller. To formalize this, we spell out the condition on the cost-gap under which the buyer may benefit from screening the seller, *under some distribution of cost types*. By our previous analysis, this holds whenever

$$\theta_n - \theta_1 > \frac{1-\theta_n}{2}. \tag{5.3}$$

Condition (5.3) generalizes (3.5) to the setting with more than two production technologies. As we explain below, if it is violated, the buyer will never benefit from screening the seller by making an offer below θ_{j^o}, since such an offer is dominated by p_{θ_n}.

Theorem 4. *There exists at least one equilibrium of the reputational bargaining game with endogenous technology adoption. For every $\eta > 0$, there exist $\overline{\nu} > 0$ and $\overline{\varepsilon} > 0$ such that in every* *equilibrium where* $\nu < \overline{\nu}$ *and* $\varepsilon < \overline{\varepsilon}$*,*

1. *If (5.3) is violated, the seller adopts* θ_{j^o} *with probability greater than* $1-\eta$*, and the expected welfare loss from delay is less than* η*, in any equilibrium.*

2. *If (5.3) is satisfied, there exists an open set of adoption costs such that there exists an equilibrium where the seller adopts* θ_{j^o} *with probability bounded below one and the expected delay is bounded above zero.*

3. *If* $p_{\theta_{j^o}} < \theta_n$*, there exists an open set of adoption costs such that, in all equilibria, the seller adopts* θ_{j^o} *with probability bounded below one and the expected delay is bounded above* 0*.*

The proof is in Online Appendix D, which is similar to the one for Theorem 2. In order to understand the statement of Theorem 4, consider first the case in which (5.3) is violated. This condition ensures that, in any equilibrium, the buyer will make an offer equal to the Rubinstein bargaining price under the highest-cost type in the support of the seller's adoption decision. In particular, let $\pi \in \Delta(\Theta)$ be the distribution over production costs that the seller chooses in equilibrium. We relabel the elements of Θ so that $\text{supp}(\pi) = \{\hat{\theta}_1, ..., \hat{\theta}_m\}$, and let $\hat{c}_j$ denote the adoption cost associated with $\hat{\theta}_j$. First, note that the negation of (5.3) implies that for all $i \in \{1, ..., m-1\}$

$$\hat{\theta}_{i+1} \leq \theta_1 + 1 - p_{\theta_n} < \hat{\theta}_i + 1 - p_{\hat{\theta}_i} = p_{\hat{\theta}_i}.$$

Therefore, if the buyer decides to screen the seller with production cost no more than $\hat{\theta}_i$, then she will offer $\hat{\theta}_{i+1}$. Furthermore, if the buyer screens the seller by offering $\hat{\theta}_{i+1}$ for some $i \in \{1, ..., m-1\}$, then she obtains a payoff of

$$\pi[\hat{\theta}_1, \hat{\theta}_i](\hat{\theta}_{i+1} - \hat{\theta}_i) \leq \pi[\hat{\theta}_1, \hat{\theta}_i](1 - p_{\theta_n}) \leq \pi[\hat{\theta}_1, \hat{\theta}_i](1 - p_{\hat{\theta}_m}) < 1 - p_{\hat{\theta}_m}$$

This implies that any such offer is strictly dominated by offering $p_{\hat{\theta}_m}$. Given that the equilibrium price offered by the buyer equals $p_{\hat{\theta}_m}$, the seller's equilibrium payoff is $p_{\hat{\theta}_m} - \hat{\theta}_m - \hat{c}_m \leq p_{\hat{\theta}_m} - \theta_{j^o} - c_{j^o}$, with strict inequality if $\hat{\theta}_m \neq \theta_{j^o}$. This implies that it is optimal for the seller to adopt the socially efficient technology with probability close to 1.

Conversely, condition (5.3) is sufficient for inefficiencies to arise in equilibrium under an open set of production costs. In particular, if (5.3) is satisfied and the adoption costs are such that $j^o = 1$, we can construct an equilibrium where the seller mixes between adopting θ_{j^o} and keeping the default technology θ_n, in an analogous manner as in Section 3.2. To do this, the seller's adoption strategy must be such that the buyer is indifferent between offering p_{θ_n} and $\min\{p_{\theta_{j^o}}, \theta_n\}$, which is ensured by $\pi(\theta_{j^o}) \approx \frac{p_{\theta_n}-\theta_n}{\min\{p_{\theta_{j^o}},\theta_n\}-\theta_{j^o}}$. Observe that (5.3) guarantees that $\pi(\theta_{j^o}) < 1$.

The discussions above imply that the seller with production cost θ_n trades with delay in equilibrium when the buyer offers $\min\{p_{\theta_{j^o}}, \theta_n\}$. Moreover, if $c_{j^o} > \max\left\{\frac{1}{2}, \frac{1-\theta_n}{1-\theta_{j^o}}\right\}(\theta_n - \theta_{j^o})$, then there exists a mixed strategy over bargaining postures for the buyer that assigns probability to p_{θ_n} and $\min\{p_{\theta_{j^o}}, \theta_n\}$ that guarantees that the seller is indifferent between choosing θ_n and θ_{j^o}. An additional condition, which we derive in the Online Appendix, ensures that he doesn't benefit from deviating to an alternative technology $\theta \notin \{\theta_{j^o}, \theta_n\}$. Thus, an equilibrium with inefficient investment and bargaining delay exists under an open set of adoption costs when (5.3) is satisfied.

If the stronger condition $p_{\theta_{j^o}} < \theta_n$ is satisfied, then any equilibrium must be inefficient if $c_{j^o} \in \left(\frac{\theta_n-\theta_{j^o}}{2}, \theta_n - \theta_{j^o}\right)$. This is because, if he invests efficiently, the seller's equilibrium payoff is $p_{\theta_{j^o}} - \theta_{j^o} - c_{j^o}$. Given that $\theta_n > p_{\theta_{j^o}}$, he can deviate to θ_n and demand the entire surplus, which ensures a payoff of approximately $\frac{1-\theta_n}{2}$. This is strictly larger than $p_{\theta_{j^o}} - \theta_{j^o} - c_{j^o}$ whenever $c_{j^o} > \frac{\theta_n-\theta_{j^o}}{2}$.

6 Concluding Remarks

We study a reputational bargaining model where a seller's production cost is determined endogenously by his technology adoption decision before bargaining with a buyer. Due to players' incentives to build reputations, they are reluctant to revise their offers in the bargaining stage.

In contrast to the case studied by Gul (2001) in which the uninformed player frequently revise their offers, we show that when players can establish reputations for being obstinate, the seller's adoption decision may be bounded away from efficiency even when his adoption decision cannot be observed by the buyer. Moreover, inefficient adoption and costly delay arise in equilibrium if and only if the social benefit from adopting the new technology is large enough.

Our results suggest an explanation for the under-adoption of cost-saving technologies in the case where producers know the effectiveness of the new technologies, their adoption decisions cannot be directly observed by their trading partners, but they are reluctant to adopt socially efficient production technologies due to the fear of expropriation by their trading partners. We conclude with some comments on the robustness of our result once we vary some of our modeling assumptions.

Timing of Offers: Our baseline model assumes that the uninformed buyer makes their offer before the informed seller does, which is consistent with the modeling assumptions in APS and Fanning (2022). In an earlier working paper version, we study the case where players make their initial offers *simultaneously* and obtain similar results: In the game with an exogenous distribution over production costs, bargaining is efficient in all limiting equilibria when the difference between adjacent types' production costs is small, and efficient and inefficient limiting equilibria co-exist when the difference between adjacent types' production costs is large. The intuition is that when players make offers simultaneously, the seller does not know whether the buyer will make a screening offer or a pooling offer, which explains why both can arise in equilibrium. But in the game with endogenous technology adoption, there is a unique limiting equilibrium, in which the seller's adoption decision is socially inefficient if and only if the benefit from adoption is large enough.

Our main results partially extend to a model where the order with which players make offers is endogenous and, as in Kambe (1999) and Wolitzky (2012), each player becomes committed with positive probability after making their initial offer. In this game, there exists an equilibrium where the buyer makes an offer before the seller does, and players' equilibrium strategies coincide with those in the baseline model. Nevertheless, there also exist other equilibria due to the seller's incentive to signal his production cost. In particular, the off-path belief about the seller's cost has a significant effect on players' incentives when the seller can make an offer before the buyer does.

In the setting where players may only become committed before making an offer, the endogenous timing of offers that arises in equilibrium will depend on the assumptions about when the commitment types make their offer. This is because a rational player will time her offer so as to mimic some commitment type. In the presence of temporal commitment types that may make an offer with delay, APS show that the equilibrium is qualitatively different from the one that obtains when all commitment types make their offers at time zero. In particular, they show that the strong type of informed player (the high-cost seller in our model and the more patient player in theirs) would use delay to signal their type. Solving for the equilibrium under temporal commitment types is beyond the scope of this paper, but their findings seem to suggest that endogeneizing the timing of offers in this way would lead to different equilibrium outcomes.

Bargaining Power: In contrast to Gul (2001) in which the uninformed player makes all the offers and hence has all the bargaining power, we study a reputational bargaining model in which both players have bargaining power, determined by the ratio of their discount rates. In particular, a player has more bargaining power when they are more patient relative to their opponent.

Our baseline model assumes that players share the same discount rate and hence have equal bargaining power. By extending our analysis to the case where players have *different discount rates*, we can examine the effects of bargaining power on the seller's adoption decision and on the expected delay in reaching agreements.

Suppose the buyer's discount rate is r_b and the seller's discount rate is r_s. When the buyer's value for the object is 1 and the seller's cost is θ, the equilibrium price in the Rubinstein bargaining game is $p_\theta \equiv \frac{r_b}{r_s+r_b} + \frac{r_s}{r_s+r_b}\theta$. Since the seller obtains a fraction $\frac{r_b}{r_b+r_s}$ of the total surplus $1-\theta$, his bargaining power is $\frac{r_b}{r_b+r_s}$ and the buyer's bargaining power is $\frac{r_s}{r_b+r_s}$.

Suppose first that r_b/r_s is small enough that $p_{\theta_1} < \theta_2$, which is the case in which the buyer's bargaining power is relatively high. Then, if $\pi(\theta_1)$ is sufficiently high, the limiting equilibrium in the game with exogenous production costs is inefficient, and features the buyer offering p_{θ_1} and trading with delay. Otherwise, the limiting equilibrium is efficient. When the seller's adoption decision is endogenous, there will be interior adoption and strictly positive delay in equilibrium if and only if $c \in \left(\frac{r_b}{r_b+r_s}(\theta_2-\theta_1), \theta_2-\theta_1\right)$. In particular, if $c < \theta_2 - \theta_1$ so that adoption is optimal and r_b/r_s is arbitrarily small, the limiting equilibrium will be inefficient except for an interval of adoption costs of vanishing measure.

As in the baseline model, the intuition behind the bargaining inefficiencies comes from the uninformed buyer's incentive to screen the informed seller. Screening is more attractive for the buyer when she has more bargaining power: the price cut that she obtains from making the screening offer is approximately $\frac{r_s}{r_b+r_s}(\theta_2-\theta_1)$, which decreases with r_b/r_s.

Conversely, in the case in which r_b/r_s is high enough so that $p_{\theta_1} > \theta_2$, the welfare properties of the limiting equilibrium will hinge on the size of $\theta_2 - \theta_1$ in a similar way as in Theorems 1 and 2. Specifically, if $\theta_2 - \theta_1 < \frac{r_b}{r_b+r_s}(1-\theta_2)$, then the unique limiting equilibrium is efficient, both under an exogenous distribution of production cost and under endogenous distributions of production costs. Otherwise, there always exists an efficient limiting equilibrium, but an inefficient equilibrium with under adoption and delay may also exist under intermediate values of the adoption cost.

This highlights a stark contrast when we compare equilibrium welfare in the extreme cases in which one of the player's discount rate is arbitrarily higher relative to their opponent's. In order to see this, suppose that $c < \theta_2 - \theta_1$ (otherwise, the equilibrium is always efficient in the limit). If the buyer is arbitrarily more patient than the seller, then the unique limiting equilibrium features under adoption and delay for almost all values of c. If the seller is arbitrarily more patient than the buyer, then an efficient limiting equilibrium always exists, although another inefficient equilibrium may arise under certain parameter configurations.

References

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

[2] Abreu, Dilip and David Pearce (2007) "Bargaining, Reputation, and Equilibrium Selection in Repeated Games with Contracts," *Econometrica*, 75(3), 653-710.

[3] Abreu, Dilip, David Pearce and Ennio Stacchetti (2015) "One Sided Uncertainty and Delay in Reputational Bargaining," *Theoretical Economics*, 10(3), 719-773.

[4] Ali, Akhter and Awudu Abdulai (2010) "The Adoption of Genetically Modified Cotton and Poverty Reduction in Pakistan," *Journal of Agricultural Economics*, 61(1), 175-192.

[5] Ausubel, Lawrence and Raymond Deneckere (1989) "Reputation in Bargaining and Durable Goods Monopoly," *Econometrica*, 57(3), 511-531.

[6] Baliga, Sandeep and Tomas Sjöström (2023) "Long Wars," Working Paper.

[7] Billingsley, Patrick (2013a) *Convergence of Probability Measures*. John Wiley and Sons.

[8] Billingsley, Patrick (2013b) *Probability and Measure*. John Wiley and Sons.

[9] Che, Yeon-Koo and József Sákovics (2004) "A Dynamic Theory of Holdup," *Econometrica*, 72(4), 1063-1103.

[10] Compte, Olivier and Philippe Jehiel (2002) "On the Role of outside Options in Bargaining with Obstinate Parties," *Econometrica*, 70(4), 1477-1517.

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

[12] Dilmé, Francesc and Daniel Garrett (2022) "A Dynamic Theory of Random Price Discounts," Working Paper.

[13] Dutta, Rohan (2022) "Bargaining as a Struggle Between Competing Attempts at Commitment," Working Paper.

[14] Ekmekci, Mehmet and Hanzhe Zhang (2022) "Reputational Bargaining with Ultimatum Opportunities," Working Paper.

[15] Falck-Zepeda, José Benjamin, Greg Traxler, and Robert Nelson (2000) "Surplus Distribution from the Introduction of a Biotechnology Innovation," *American Journal of Agricultural Economics*, 82(2), 360-369.

[16] Fanning, Jack (2018) "No Compromise: Uncertain Costs in Reputational Bargaining," *Journal of Economic Theory*, 175, 518-555.

[17] Fanning, Jack (2022) "Outside Options, Reputations, and the Partial Success of the Coase Conjecture," Working Paper.

[18] Feinberg, Yossi and Andrzej Skrzypacz (2005) "Uncertainty about Uncertainty and Delay in Bargaining," *Econometrica*, 73(1), 69-91.

[19] Fuchs, William and Andrzej Skrzypacz (2010) "Bargaining with Arrival of New Traders," *American Economic Review*, 100(3), 802-836.

[20] Grossman, Sanford and Oliver Hart (1986) "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," *Journal of Political Economy*, 94(4), 691-719.

[21] Gul, Faruk (2001) "Unobservable Investment and the Hold-Up Problem," *Econometrica*, 69(2), 343-376.

[22] Kambe, Shinsuke (1999) "Bargaining with Imperfect Commitment," *Games and Economic Behavior*, 28(2), 217-237.

[23] Kim, Kyungmin (2009) "The Coase Conjecture with Incomplete Information on the Monopolist's Commitment," *Theoretical Economics*, 4(1), 17-44.

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

[25] Ortner, Juan (2017) "Durable Goods Monopoly with Stochastic Costs," *Theoretical Economics*, 12, 817-861.

[26] Pei, Harry and Maren Vairo (2021) "Reputational Bargaining with Unknown Values," arXiv:2201.01827.

[27] Qaim, Matin and Alain de Janvry (2003) "Genetically Modified Crops, Corporate Pricing Strategies, and Farmers' Adoption: The Case of BT Cotton in Argentina," *American Journal of Agricultural Economics*, 85(4), 814-828.

[28] Rubinstein, Ariel (1982) "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50(1), 97-109.

[29] Strulovici, Bruno (2017) "Contract Negotiation and the Coase Conjecture," *Econometrica*, 85(2), 585-616.

[30] Williamson, Oliver (1975) "Markets and Hierarchies: Analysis and Antitrust Implications," New York, Macmillan Publishers.

[31] Wolitzky, Alexander (2012) "Reputational Bargaining With Minimal Knowledge of Rationality," *Econometrica*, 80(5), 2047-2087.

[32] Wolitzky, Alexander (2018) "Learning from Others' Outcomes," *American Economic Review*, 108(10), 2763-2801.