Reputational Bargaining and Inefficient Technology Adoption*

Harry Pei[†] Maren Vairo[‡]

October 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 cost of production. We show that even when the buyer *cannot* observe the seller's adoption decision, players' reputational incentives can lead to inefficient under-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. Our result implies that an increase in the benefit from adoption may lower the probability of adoption and that the seller's opportunity to adopt a cost-saving technology may lower social welfare.

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 he knows that the gain from adoption outweighs its cost, he 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 that bargaining is efficient 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

*We thank Dilip Abreu, Sandeep Baliga, Francesc Dilmé, Mehmet Ekmekci, Jack Fanning, Richard Holden, Ruitian Lang, Lucas Maestri, Juan Ortner, Karthik Sastry, Bruno Strulovici, Chris Udry, Alex Wolitzky, and Hanzhe Zhang for helpful comments. We thank the 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

as his opponents can *frequently revise their offers* and *cannot observe* how much he has invested.

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 the investment decision cannot be observed by others, players' reputational incentives can lead to *inefficient investments* and *costly delays in reaching agreements*.[1] We also show that under-investment occurs in equilibrium *if and only if* the social benefit from investment is large enough.

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.

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.[2] With positive probability, each player is one of *a rich set of commitment types* who makes an exogenous offer and never concedes. With complementary probability, they are rational and decide what to offer and when to concede in order to maximize their discounted payoff. In consistent with the reputational bargaining literature pioneered by Kambe (1999) and Abreu and Gul (2000), we focus on the limiting case where the probability of commitment types goes to zero.

Theorem 1 characterizes 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 players only have private information about their discount rates.

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] When the seller's investment can be *observed* by the buyer in a reputational bargaining game, there will be no delay in bargaining but the seller's adoption decision will be *socially inefficient*.

[2] The uninformed player making the offer first is standard in the reputational bargaining literature which is also assumed in Abreu, Pearce and Stacchetti (2015) and Fanning (2023)'s contemporary work. If the informed player makes their offer first, then their signaling incentives lead to multiple equilibria, which we leave for future research.

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 reputation building. After the buyer offers a price *between the two production costs*, i.e., making a *screening offer*, the high-cost seller's payoff from conceding is negative, so he can credibly commit not to concede. Due to the high type's commitment power, he does not need to increase his speed of reputation building, which 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. We show that in *every* equilibrium, the high type *demands the entire surplus* and the low type demands strictly less. In order to satisfy the low type's incentive constraint, players trade with delay after the seller demands the entire surplus.

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 equals *the high type's Rubinstein bargaining price* in equilibrium) 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, after which she will lose all her surplus when she faces the high 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. Intuitively, when the cost difference is small, all the screening offers are much lower than 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 since she will end up paying a higher price to the low type.

Our main result, 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 results imply that when the adoption decision is unobservable, inefficient adoption and costly delay 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 *lower* the probability of adoption and the seller's opportunity to adopt a cost-saving technology may *lower* social welfare.

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. A useful observation is 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 interesting case where the adoption cost is *strictly 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 and will provide the buyer a strict incentive to offer low prices.

Therefore, 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 so 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. With positive probability, the seller does not adopt and the buyer makes the screening offer, in which case players will reach an agreement with delay. The equilibrium adoption decision is *socially inefficient* since it is efficient to adopt yet the adoption probability is bounded below one.

In the complementary scenario where the social benefit from adoption is small, Theorem 1 seems to suggest that the buyer prefers to make the pooling offer regardless of the adoption probability. If this is the case, then the seller's private gain from adoption equals the social benefit. When the social benefit from adoption is strictly greater than the cost of adoption, there seems to be 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 offers a low price, the low-cost seller accepts the offer with probability close to 1, and the high-cost seller counteroffers a higher price and trades with delay.

In summary, compared to games with exogenous cost distributions, in the game with endogenous adoption, delay *always* arises *conditional on* the seller having a high production cost. However, the *expected* welfare loss from delay is close to 0 when the adoption probability is close to 1.

Our result suggests an explanation for the under-adoption of cost-saving technologies, which is documented in agriculture and healthcare. Wolitzky (2018) provides an alternative explanation based on social learning, which fits applications where producers *do not know* the effectiveness of the new technology. Our theory fits when (i) the producers *know* that the new technology can effectively lower cost from earlier adoption results, (ii) it is hard for buyers to observe the producers' adoption decisions, but (iii) the producers are reluctant to adopt due to the fear of being held-up.

One example that fits is the under-adoption of Bt cotton. It is well-known among farmers that Bt cotton can reduce insecticide applications which can lower the cost per unit yield (Qaim and de Janvry 2003). Although the *crop* of Bt cotton is more resistant to pests compared to that of traditional cotton, it is hard for buyers to distinguish between the two since (i) the crops have similar appearances and (ii) both crops lead to the same final product, i.e., cotton. Some farmers in Pakistan sampled by Ali and Abdulai (2010) are major landholding households, who may have bargaining power. However, the adoption rate is low even among those households, e.g., it is only 62% in the Punjab province of Pakinstan (Ali and Abdulai 2010). 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 given that (i) the adopted farmers are more likely to be members of organizations that have more bargaining power, and (ii) the farmers' share of surplus is much lower than that in countries that have higher adoption rates (Falck-Zepeda, et al. 2000).

We discuss our contributions to the related literature in the remainder of this section. Section 2 sets up the baseline model in which the seller chooses between two production technologies. The main results are stated in Section 3 and are shown in Section 4. Section 5 extends our results to settings where the seller chooses between three or more technologies. Section 6 concludes and discusses cases in which players have different discount rates or different commitment probabilities.

Hold-Up Problem: Grossman and Hart (1986) assume that investments are *observable* and that bargaining is efficient. They show that a player's investment will be bounded below the socially optimum unless they have all the bargaining power. Gul (2001) assumes that investments are *unobservable*. He shows that players will invest efficiently even when they have no bargaining power, as long as their opponent can frequently revise their offers.

We incorporate the possibility that players might be reluctant to revise their offers due to their reputation concerns. We show that the hold-up problem re-emerges even when the investment decision is *unobservable*. Our result implies that the *absolute magnitude* of the social benefit from investment has a significant effect on players' investment incentives. This is complementary to the existing theories on the hold-up problem in which a player's investment incentive depends only on the *ratio* between the social benefit from investment and their cost of investment.[3]

Our results also apply to other forms of *selfish* investments, such as the seller deciding whether to *divest* and become less efficient and the buyer deciding whether to increase their valuation. When investments are *cooperative* in the sense of Che and Hausch (1999), such as the seller's investment increases the buyer's value, the seller has no incentive to invest when his investment is unobservable, but has a positive incentive to invest when it is observable. This stands in contrast to the case studied in our paper where non-observability leads to stronger investment incentives.

Reputational Bargaining: To the best of our knowledge, the existing works on reputational bargaining focus on the case in which the distribution over players' preferences is *exogenous*. We take a first step to analyze the case where the distribution over a seller's production cost is *endogenous*. We explain why the existing characterization results on reputational bargaining with *exogenous* type distributions *cannot* be applied to settings with an *endogenous* type distribution.

We also incorporate *heterogeneity* in players' costs in the reputational bargaining model of Abreu and Gul (2000).[4] In contrast to Abreu, Pearce and Stacchetti (2015, or APS) who analyze the case where one player has private information about their discount rate, we show that heterogeneity in costs enables the uninformed player to *screen* the informed player via prices, leading to delays.[5]

A contemporary work of Fanning (2023) also incorporates heterogeneity in players' values (or equivalently, costs). It shows that bargaining is efficient when (i) no rational type is indifferent between accepting and rejecting any commitment type's offer, and (ii) the value (or cost) is drawn from a rich set. His first requirement violates our richness assumption on the set of commitment types, which we discuss in Section 2.1. Our inefficient adoption result requires the existence of two adjacent types with sufficiently different production costs, which violates his richness assumption on the set of values. Which modeling assumption on the set of values/costs fits depends on the application. Our assumption fits when the heterogeneity in production costs is driven by *differences in production technologies*. This is because adoption decisions are usually *indivisible* and adopting

[3] In Che and Sákovics (2004), investment incentives depend on the difference between the social surplus and the cost of investment, whereas investment incentives depend on the *absolute magnitude* of social surplus in our model.

[4] 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.

[5] APS also consider the case in which some commitment types play *non-stationary strategies*. Since our motivation is to revisit the hold-up problem when players are unwilling to revise their offers, we assume that all commitment types demand the same price over time, which is also assumed in Abreu and Gul (2000) and Fanning (2023).

an innovative technology may significantly lower the cost of production.

Bargaining with Incomplete Information: We focus on private value settings as in Gul, Sonnenschein and Wilson (1986) where there is a gap between players' values. In contrast to their canonical result that there is no delay as the bargaining friction vanishes, we show that players' reputation concerns can cause delays. Our result stands in contrast to the inefficiency results that are driven by interdependent values (Deneckere and Liang 2006, Baliga and Sjöström 2023), no gap between players' values (Ausubel and Deneckere 1989), costly concessions (Dutta 2022), risk aversion (Dilmé and Garrett 2022), and the arrival of new traders (Fuchs and Skrzypacz 2010).

Ortner (2017) shows that when a seller's cost may decrease over time, the bargaining outcome is efficient if and only if the buyers' 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 inefficiencies in our model are driven by players' incentives to build reputations for being obstinate, which differs from the mechanism behind Ortner's result.

2 The 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 we normalize 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 the buyer *cannot* observe. 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, with $0 < \theta_1 < \theta_2 < 1$. We extend our results to settings with three or more production technologies in Section 5.

In the second stage, the buyer offers a price $p_b \in [0, 1]$. The seller either accepts the offer and sells at price p_b, or rejects the offer and makes a counteroffer $p_s \in (p_b, 1]$, 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 share the same discount rate $r > 0$. If trade happens at time $\tau \in [0, +\infty)$ and 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 buyer-commitment-type is characterized by $p_b \in \mathbf{P}_b \subset [0,1]$, who offers p_b and never accepts prices greater than p_b. Each seller-commitment-type is characterized by $p_s \in \mathbf{P}_s \subset [0,1]$, who offers p_s and never accepts prices 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 committed, which we assume have full support.

We assume that $\mathbf{P}_b$ and $\mathbf{P}_s$ are compact and countable, $1 \in \mathbf{P}_s$, and $\sup \mathbf{P}_s \backslash \{1\} = 1$, that is, the seller *can* build a reputation for demanding the entire surplus and also for demanding something less than but close to the entire surplus.[6] Section 2.1 discusses this assumption in detail. Let

$$\nu \equiv \inf \left\{ \overline{\nu} > 0 \middle| \text{ for every } p \in [0,1],\ (p-\overline{\nu}, p+\overline{\nu}) \cap \mathbf{P}_s \neq \emptyset \text{ and } (p-\overline{\nu}, p+\overline{\nu}) \cap \mathbf{P}_b \neq \emptyset \right\}. \quad (2.1)$$

Intuitively, ν measures the *richness* of the sets of commitment types. In consistent with Abreu and Gul (2000), our analysis focuses on the case in which both ν and ε are close to 0, that is, each player has a rich set of commitment types 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).

2.1 Discussions on the Modeling Assumptions

We take a reputational bargaining approach since our motivation is to revisit the hold-up problem when players may not want to change their bargaining postures due to their reputation concerns. Compared to bargaining models with incomplete information but *without* commitment types such as Gul, Sonnenschein and Wilson (1986) and Gul (2001), reputational bargaining models lead to sharp predictions on players' behaviors and welfare even when each player has some bargaining

[6] Our requirement is satisfied, for example, when there exists $\nu \in (0,1)$ such that the set of seller-commitment-types *contains* 1 and $p^j \equiv 1-(1-\nu)^j$ for every $j \in \mathbb{N}$. We allow the existence of other commitment types.

power, which sounds more realistic than assuming that one player *cannot* make any offer.[7]

We assume that the *uninformed* buyer makes their offer before the *informed* seller. The uninformed party making the first offer is a standard assumption in reputational bargaining models with one-sided uncertainty about payoffs, which is also assumed in APS and the contemporary work of Fanning (2023). Our assumption is plausible when a firm procures products from its upstream supplier, in which the firm sets a price first before negotiating with the supplier. The case where the informed player making the first offer is left for future work.

We assume that every obstinate type demands the same price over time. This is a standard assumption, which is also assumed in Kambe (1999), Abreu and Gul (2000), Ekmekci and Zhang (2022), and Fanning (2018, 2023). Our assumption is motivated by a practical concern that once a player changes their demand, it might be hard for them to convince others that they are obstinate.

We assume that the set of commitment types is *rich* in the sense that (i) there exists a commitment type who demands something close to p for every $p \in [0, 1]$, and (ii) there exist seller-commitment-types who demand the entire surplus as well as seller-commitment-types whose demands are less than but close to the entire surplus. Our first requirement is standard in the literature, which is also assumed in Abreu and Gul (2000) and Fanning (2023).

As for our second requirement, although Abreu and Gul (2000) assume that all commitment types demand something strictly between 0 and 1, their main result that players will trade immediately at the Rubinstein bargaining price *extends* to the case where the set of commitment types satisfies our richness requirements. We also establish the *existence* of equilibrium in Online Appendices A and B in our setting where the number of commitment types is countably infinite.

Our second requirement is violated in Fanning (2023) who assumes that no rational type is indifferent between accepting and rejecting any offer made by any commitment type. His assumption rules out commitment types who demand the entire surplus, under which he shows that players will reach an agreement without any delay. The motivation for our requirement is twofold.

First, an economic interpretation for *a rich set of commitment types* is that a player can build a reputation for being obstinate as long as they demand the same thing over time, *regardless of what their demand is.* Although we require the set of commitment types to be countable in order to circumvent measurability issues, our results are robust when we include *any* additional commitment

[7] It is well-known that bargaining models where the informed player can make offers are not tractable to analyze. As a result, most of the existing works focus on the case where the uninformed player makes all the offers (e.g., Gul, Sonnenschein and Wilson 1986, Gul 2001). An exception is Gerardi, Hörner and Maestri (2014) that characterizes the set of equilibrium payoffs when the informed player makes *all* the offers. We are unaware of any paper that analyzes models *without* any commitment type where both the informed and uninformed player can make offers.

type, as long as that commitment type demands *the same price* over time.

Second, our results are in the spirit of Fudenberg and Levine (1989) and Abreu and Gul (2000) that *ideally*, reputation results should apply *as long as a certain set of commitment types occur with positive probability*, even when there might be other commitment types as well.[8] When the cost difference between the two types of the seller is large enough, we show that equilibria with delay *exist* as long as there *exists* a seller-commitment-type who demands the entire surplus, and that delay will occur in *all* equilibria when there also exist a sequence of seller-commitment-types whose demands approach the entire surplus. Admittedly, our results do not allow for non-stationary commitment types, but they are robust against the presence of other *stationary* commitment types. Our motivation for focusing on stationary commitment types is explained earlier in this section.

3 Analysis & Results

Section 3.1 analyzes a benchmark where the buyer can observe the seller's adoption decision, in which case she will know the seller's production cost. Section 3.2 analyzes a reputational bargaining game in which the seller's cost is drawn from an *exogenous* distribution and the buyer does *not* know the realized cost. Theorem 1 identifies conditions under which players will reach an agreement after a significant delay as well as the sources of inefficiencies. Section 3.3 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 Benchmark: Adoption Decision is Observable

Suppose the buyer *can observe* the seller's adoption decision, that is, the buyer knows θ. Proposition 3 in Abreu and Gul (2000) implies that as $\varepsilon \to 0$, players will trade with almost no delay at a 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 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

[8]For example, Fudenberg and Levine (1989) show that the patient player can secure their Stackelberg payoff as long as there *exists* a commitment type who takes the Stackelberg action. Abreu and Gul (2000) show that players obtain their Rubinstein bargaining payoffs as long as there *exist* commitment types who demand those payoffs.

sum of these payoffs equals the social surplus from trade $1-\theta$. The seller's gain from adoption is

$$(p_{\theta_1}-\theta_1)-(p_{\theta_2}-\theta_2)=\frac{\theta_2-\theta_1}{2}. \tag{3.1}$$

Hence, the seller will adopt only 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 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 is efficient but the adoption decision is inefficient.

3.2 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, Theorem 1, characterizes the common properties of *all* equilibria as the probability of commitment types vanishes and the set of commitment types satisfy our richness assumption.

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 $\sigma_s^*(\cdot) \equiv \big\{\sigma_{s,\theta}^*(\cdot)\big\}_{\theta\in\Theta}$, where

$$\sigma_{s,\theta}^*(p_b) \equiv \begin{cases} 1, & \text{if } p_b \leq \theta_1, \text{ or } p_b \in (\theta_1,\theta_2] \text{ and } \theta=\theta_2, \\ \max\big\{p_b, 1+\theta_1-p_b\big\}, & \text{if } p_b \in (\theta_1,\theta_2] \text{ and } \theta=\theta_1, \\ \max\big\{p_b, 1+\theta_2-p_b\big\}, & \text{if } p_b > \theta_2. \end{cases} \tag{3.2}$$

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 $\max\big\{p_b, 1+\theta_i-p_b\big\}$. 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

$$\lambda_s \equiv \frac{r(1-p_s)}{p_s-p_b}, \tag{3.3}$$

and the buyer will concede at rate

$$\lambda_b^i \equiv \frac{r(p_b-\theta_i)}{p_s-p_b}. \tag{3.4}$$

These are the rates that make the other player indifferent between conceding and not conceding.

Proposition 3 in Abreu and Gul (2000) implies that 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. If θ_i is common knowledge, then (3.3) and (3.4) imply that 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

$$\pi^* \equiv \min\left\{1, \frac{p_{\theta_2}-\theta_2}{\min\{p_{\theta_1},\theta_2\}-\theta_1}\right\}, \tag{3.5}$$

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

$$\theta_2-\theta_1 > \frac{1-\theta_2}{2}, \tag{3.6}$$

that is, the difference between θ_1 and θ_2 is large relative to the surplus generated by the high-cost type. Theorem 1 shows that for generic π, all equilibria of the bargaining game converge to the same limit point when the sets of commitment types are rich and players are committed with probability close to 0.[9] It also characterizes the welfare properties of the unique limiting equilibrium.

Theorem 1. *For every $\pi \in \Delta\{\theta_1,\theta_2\}$, there exists at least one equilibrium. Suppose in addition that π 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*

$$(1-\theta_2)\left\{1-\frac{\max\left\{p_{\theta_1},1-(\theta_2-\theta_1)\right\}-\theta_1}{1-\theta_1}\right\}. \tag{3.7}$$

The proof is in Section 4.1 with some details relegated to Online Appendix A. According to Theorem 1, the qualitative features of the limiting equilibrium depend on (i) the difference $\theta_2-\theta_1$

[9] 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.

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 (3.6), 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 (3.6) and the low type occurs with probability less than π^*.

2. When θ_1 and θ_2 satisfy (3.6) 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. This leads to an expected welfare loss of (3.7). 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 delays arise 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 payoff information (Kambe 1999 and Abreu and Gul 2000), as well as those in reputational bargaining games with one-sided private information about payoffs, but either the private information is about the discount rate (e.g., APS) or the set of commitment types violates our richness assumption (e.g., Fanning 2023). We discuss those models by the end of this section.

We argue that inefficient delays arise 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 for him, which motivates him to demand a high price. Since type θ_1 receives a strictly positive payoff from conceding, 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 or equal 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 demand *in equilibrium* after the buyer offers $p_b \in (\theta_1, \theta_2]$? Suppose by way of contradiction that he demands $p_s < 1$ with positive probability. The low type must offer every price in $\mathbf{P}_s \cap (p_s, 1)$ with positive probability. This is because otherwise, the high type will find it strictly profitable to deviate to one of these prices instead of offering p_s. In equilibrium, for every pair of prices offered by the seller with positive probability, offering the higher price will lead to a longer expected delay but a higher expected trading price. This implies that when the low type has an incentive to demand a higher price p'_s, the high type should have no incentive to demand a lower price p_s. This leads to a contradiction. Lemma 4.1 uses this logic to show that in every equilibrium, the high type must demand the entire surplus following any screening offer.

Therefore, 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 offer 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 (3.6). When $\theta_2 - \theta_1 > \frac{1-\theta_2}{2}$, (3.5) 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 make the screening offer if and only if $\pi(\theta_1) > \pi^*$.

A natural question following Theorem 1 is that what will happen when $\pi(\theta_1) = \pi^*$? 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 for one of these offers 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, which we will discuss in Section 3.3.

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 $\mathbb{E}[e^{-r\tau_b}|\theta=\theta_2]$:

$$\underbrace{(1-\theta_1)\mathbb{E}[e^{-r\tau_b}|\theta=\theta_2]}_{\text{type } \theta_1\text{'s payoff from demanding } 1} \leq \underbrace{\max\{p_{\theta_1}, 1-\theta_2+\theta_1\}-\theta_1}_{\text{type } \theta_1\text{'s equilibrium payoff}}. \tag{3.8}$$

When players offer p_b and p_s, let T_1 denote the time it takes for type θ_1 to finish conceding and let c_b denote the probability with which the buyer concedes at time 0, both of which depend on the buyer's posterior belief about the seller's type. Type θ_2 cannot profit from demanding p_s that is less than but close to 1 and then waiting for the buyer to concede, which implies that

$$\underbrace{(1-\theta_2)\mathbb{E}[e^{-r\tau_b}|\theta=\theta_2]}_{\text{type } \theta_2\text{'s equilibrium payoff}} \geq \underbrace{(p_s-\theta_2)\left(c_b+(1-c_b)\Big(1-e^{-(r+\lambda_b^1)T_1}\Big)\frac{\min\{p_{\theta_1},\theta_2\}-\theta_1}{p_s-\theta_1}\right)}_{\text{type } \theta_2\text{'s payoff from deviating to } p_s\approx 1}, \tag{3.9}$$

where λ_b^1 is defined in (3.4). This leads to a lower bound on $\mathbb{E}[e^{-r\tau_b}|\theta=\theta_2]$. We show in Section 4.1 that as $p_s\to 1$ and $\varepsilon\to 0$, the right-hand-side of (3.9) is at least

$$\frac{\max\{p_{\theta_1}, 1-\theta_2+\theta_1\}-\theta_1}{1-\theta_1}(1-\theta_2). \tag{3.10}$$

This lower bound is attained when the buyer assigns zero probability to the seller having a high cost after observing offers p_b and p_s. The upper and the lower bounds for $\mathbb{E}[e^{-r\tau_b}|\theta=\theta_2]$ coincide in the limit, which pin down the limiting value of $\mathbb{E}[e^{-r\tau_b}|\theta=\theta_2]$. Our calculation 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. The low-cost seller not only bears the loss from 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 π, 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 examining the effect of an increase in the fraction of sellers with a low production cost.

Corollary 1. *Both the expected welfare loss and the expected delay are zero when $\pi(\theta_1) \in [0, \pi^*)$, but are strictly positive and are strictly decreasing in $\pi(\theta_1)$ when $\pi(\theta_1) \in (\pi^*, 1)$.*

Corollary 1 suggests that the expected welfare loss from delay is maximized when $\pi(\theta_1)$ is slightly above π^*. Intuitively, bargaining is efficient when the low type occurs with probability no more than π^*. When $\pi(\theta_1)$ is above π^*, inefficient delay occurs only when the seller has a high production cost θ_2, and conditional on $\theta = \theta_2$, the expected welfare loss from delay is independent of $\pi(\theta_1)$. Next, we examine the effect of an increase in the production cost of the high-cost type.

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 expected 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 (3.8) and (3.9), 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. Next, we examine the effect of an increase in the production cost of the high-cost type.

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}.

Remark: We show that bargaining is inefficient when (i) the set of commitment types is *rich* and (ii) the cost difference between the two types of the seller is large enough. These lessons extend to the case where the seller has three or more rational types. The presence of bargaining inefficiencies stands in contrast to APS and Fanning (2023). APS assume that players only have private information about their discount rate, in which case there is no offer under which some type has a strict incentive to concede while others have no incentive to concede. This explains why the uninformed player cannot induce different types of the informed player to make different offers.

Fanning (2023) studies reputational bargaining in which the buyer has private information about her value. He assumes that no rational type is indifferent between accepting and rejecting any commitment type's offer. This violates our richness assumption on the set of commitment types since it rules out seller-commitment types that demand the entire surplus. Intuitively, when it is *infeasible* for the seller to build a reputation for demanding the entire surplus, the buyer *cannot* screen the seller in equilibrium since it is not optimal for her to concede after delay when the seller will never concede. The motivation for our richness assumption can be found in Section 2.1.

3.3 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 (3.5) and that $\pi^* < 1$ if and only if (θ_1, θ_2) satisfies (3.6). 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 characterization as Theorem 2. A description of the limiting equilibria can be found in Lemmas 4.7-4.10 in Section 4.2, which we depict graphically in Figure 1.

Theorem 2. *There exists at least one equilibrium. 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 (3.6). In every equilibrium, the expected delay is less than η, and the adoption probability is less than η if $c > \theta_2 - \theta_1$ and is more than $1 - \eta$ if $c < \theta_2 - \theta_1$.*

2. *Suppose (θ_1, θ_2) satisfies (3.6). There exists an open interval $(\underline{c}, \bar{c}) \subset \left(\frac{\theta_2-\theta_1}{2}, \theta_2 - \theta_1\right)$ such that for every $c \in (\underline{c}, \bar{c})$, there exists an equilibrium where the adoption probability is within an η-neighborhood of π^* and the expected delay is bounded above 0.*

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 0.*

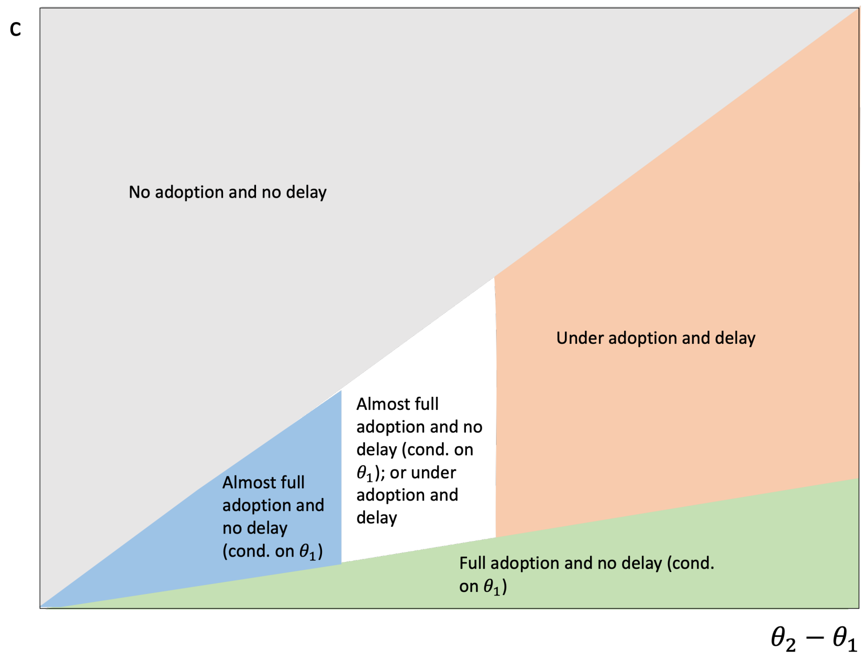

Figure 1: Limiting equilibria in the reputational bargaining game with endogenous technology adoption. Aside from the white region, there is a unique limiting equilibrium. In the white region, there are two limiting equilibria, one of them has efficient investment and negligible delay in reaching agreement, another one has inefficient investment and significant delay in reaching agreement.

The proof is in Section 4.2, with some technical details relegated to Online Appendix B. Theorem 2 shows 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 as long as the adoption cost c is between half of the social benefit from adoption $\frac{\theta_2-\theta_1}{2}$ and the entire social benefit from adoption $\theta_2 - \theta_1$.

Theorem 2 has three implications. First, inefficient adoption occurs in equilibrium if and only if *in expectation*, there are significant delays in reaching agreement. This stands in contrast to the benchmark scenario where the buyer observes the seller's adoption decision, in which case there is negligible delay in reaching agreement but the seller's adoption decision is socially inefficient.

Second, an increase in the social benefit from adoption $\theta_2 - \theta_1$ can lower the probability of adoption. This is because inefficient adoption and costly delays arise only when the social benefit is large enough. Third, the seller's opportunity to adopt a cost-saving technology may not improve social welfare, and in some sense, it may even *lower* social welfare. This is because there will be no delay in bargaining when the seller cannot adopt but there might be costly delays when the seller can adopt. We formally state these results as Corollaries 4 and 5 later in this section.

One observation that is *not* explicitly stated is that when the parameter values belong to the blue region of Figure 1, costly delay will arise in *all* equilibria *conditional on* the seller having a high production cost. This is the case even when (θ_1, θ_2) violates (3.6), under which for *any exogenous full support distribution* over production costs, the buyer *strictly* prefers to offer the Rubinstein bargaining price of the high type $p_{\theta_2} \equiv \frac{1+\theta_2}{2}$ and there is almost no delay in reaching agreement.

To understand Theorem 2 as well as the above observation, we start from a heuristic explanation using our 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*.

First, fix any $\pi \in \Delta\{\theta_1, \theta_2\}$ that has full support and satisfies $\pi(\theta_1) \neq \pi^*$. Theorem 1 implies that as $\varepsilon \to 0$, in every efficient equilibrium, the difference between the low-cost type's equilibrium payoff and that of the high-cost type's is approximately $\theta_2 - \theta_1$, and in every inefficient equilibrium, this difference in equilibrium payoffs is approximately $(\theta_2 - \theta_1)\alpha$ 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 low type not only bears the welfare losses from delay but is also appropriated by the buyer.

Fix 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 (3.6), 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 when $(\theta_2 - \theta_1)\alpha < c < \theta_2 - \theta_1$.

The above contradiction is driven by the different orders of limits in Theorems 1 and 2, making

Theorem 1 and other existing results on reputational bargaining inapplicable to settings where π is endogenous. Specifically, Theorem 1 characterizes the set of equilibria under a *fixed cost distribution* π in the limit where $\varepsilon \to 0$. The same order of limit also applies to the results in APS and Fanning (2023). However, π depends on the probability of commitment types ε when adoption is *endogenous* and the probability with which the seller having a high production cost may also vanish as $\varepsilon \to 0$.

This is indeed what happens when (θ_1, θ_2) violates (3.6) 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 Theorem 1 applies for all $\varepsilon < \overline{\varepsilon}$ and π with $\pi(\theta_1) \in [0, 1-\xi]$. However, the qualitative features of the equilibria 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 strictly prefers the pooling offer p_{θ_2} as $\varepsilon \to 0$, she strictly prefers the screening offer p_{θ_1} for any small but fixed ε as $\pi(\theta_1)$ goes to 1.

In response to the buyer's offer, type θ_2 counteroffers a higher price and trades with delay, with the expected delay pinned down by the seller's indifference condition at the adoption stage, and type θ_1 accepts the offer with probability close to 1 and pools with type θ_2 with probability close to 0. Although there are significant delays conditional on θ_2, these delays have negligible payoff consequences from an ex ante perspective since the probability of type θ_2 vanishes 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 (3.6) 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 (3.6) 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.[10] 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.

Adoption Probability & Welfare: First, we provide sufficient conditions under which an increase in the social benefit from adoption lowers the probability of adoption and leads to a longer expected delay. Delay is measured by $1 - \mathbb{E}[e^{-r\min\{\tau_s,\tau_b\}}]$, where τ_s and τ_b are the times at which the rational types of the seller and the buyer, respectively, finish conceding when no player concedes with positive probability at time 0.

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

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

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 exists $\overline{\nu} > 0$ such that for every $\nu < \overline{\nu}$, there exists $\overline{\varepsilon}_\nu > 0$ such that if $\varepsilon < \overline{\varepsilon}_\nu$,

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 $(\widehat{\theta}_1, \theta_2, c, \varepsilon, \nu)$.*

The proof is in Online Appendix F. We depict the complete comparative statics on the adoption probability and the expected delay 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 a longer expected delay, which leads to further efficiency losses.

[10] In our model, inefficient adoption is caused by the hold-up problem, so the seller will not adopt when his adoption cost is greater than the social benefit. This implies that inefficiency can only take the form of under-adoption.

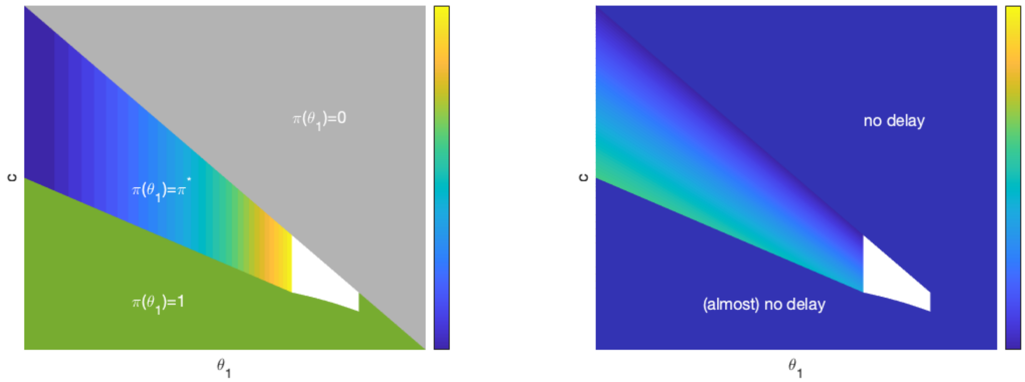

(a) Adoption probability

(b) Expected delay

Figure 2: Comparative statics on the equilibrium outcomes. The white region represents the set of parameter values under which there are multiple limiting equilibria. In regions where 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 vertical order according to the color bar on the right of the figure. That is, aside from the white region, a deeper color represents a lower adoption probability (in the left panel) and a lower expected delay (in the right panel) in the unique limiting equilibrium. In the remaining regions, the seller's adoption deicision is socially efficient: The adoption probability is 0 when $c > \theta_2 - \theta_1$, and is 1 when $c < \theta_2 - \theta_1$, and the expected delay is zero in the limit.

Next, we compare the equilibrium outcomes in our model to the ones in a benchmark where the seller has *no opportunity* to adopt the cost-saving technology. We provide sufficient conditions under which the seller's *opportunity* to adopt the technology does not increase social welfare, as well as conditions under which the opportunity to adopt leads to a strictly lower welfare.

Corollary 5. *For every $\eta > 0$, there exists $\overline{\nu} > 0$ such that for every $\nu < \overline{\nu}$, there exists $\overline{\varepsilon}_\nu > 0$ such that when $\varepsilon < \overline{\varepsilon}_\nu$,*

1. *If $\theta_2 - \theta_1 > 1 - \theta_2$ and $c \in \left(\frac{\theta_2-\theta_1}{2}, \theta_2 - \theta_1\right)$, then the equilibrium welfare when the seller is not allowed to invest is η-close to the equilibrium welfare when investing is allowed.*

2. *If $\theta_2 - \theta_1 \in \left(\frac{1-\theta_2}{2}, 1-\theta_2\right)$ and $c \in \left(\frac{1-\theta_2}{1-\theta_1}(\theta_2-\theta_1), \theta_2 - \theta_1\right)$, then the lowest equilibrium welfare when the seller cannot invest is strictly bounded above that when the seller can invest.*

The proof is in Online Appendix F. The intuition is that when the seller cannot adopt the technology, players will trade efficiently conditional on the seller having a high production cost. When the seller has an opportunity to adopt the cost-saving technology, it leads to uncertainty about the seller's production cost which will induce delays in reaching agreement. We show that the welfare losses from delay may completely offset the social benefit from adoption, and it might be strictly greater than the social benefit in some equilibria under some parameter values.

4 Proofs of Theorems 1 and 2

Before showing Theorem 1 in Section 4.1 and Theorem 2 in Section 4.2, we establish a lemma which is an implication of our richness assumption on the set of commitment types. This result applies to *all* equilibria, both in the case with an exogenous cost distribution and the case with endogenous technology adoption. Let $\hat{\varepsilon}_b(p_b)$ be the probability that the buyer is committed after she offers p_b.

Lemma 4.1. *Fix any equilibrium. For any p_b that satisfies $p_b \in \mathbf{P}_b \cap (\theta_1, \theta_2]$ and $\hat{\varepsilon}_b(p_b) < 1$, type-$\theta_2$ seller will demand 1 after observing the buyer offers p_b.*

Proof. Suppose by way of contradiction that type θ_2 offers $p_s < 1$ with positive probability after the buyer offers $p_b \in \mathbf{P}_b \cap (\theta_1, \theta_2]$. Our assumption that $\sup \mathbf{P}_s \setminus \{1\} = 1$ implies that $(p_s, 1) \cap \mathbf{P}_s$ is non-empty. Since $p_b \leq \theta_2$, type θ_1 must offer every $p'_s \in (p_s, 1) \cap \mathbf{P}_s$ with positive probability. This is because otherwise, the buyer's posterior assigns zero probability to type θ_1 after observing p'_s and will then concede immediately, in which case type θ_2 has a strict incentive to deviate to p'_s.

After the seller offers p_s and p'_s, respectively, let T and T' denote the times at which the rational-type buyer finishes conceding, let c_b and c'_b denote the buyer's concession probabilities at time 0, and let A and A' denote the discounted probability of trade conditional on the seller never concedes and the buyer being the rational type. On the one hand, type θ_2 weakly prefers p_s to p'_s, which implies that

$$(p_s - \theta_2)A \geq (p'_s - \theta_2)A'. \tag{4.1}$$

On the other hand, it is optimal for type-θ_1 seller to offer p'_s and to concede at time T', so he prefers this strategy to offering p_s and conceding at time T. This incentive constraint implies that:

$$e^{-rT'}\hat{\varepsilon}_b(p_b)(p_b - \theta_1) + (1 - \hat{\varepsilon}_b(p_b))A'(p'_s - \theta_1) \geq e^{-rT}\hat{\varepsilon}_b(p_b)(p_b - \theta_1) + (1 - \hat{\varepsilon}_b(p_b))A(p_s - \theta_1),$$

or

$$(e^{-rT'} - e^{-rT})\frac{\hat{\varepsilon}_b(p_b)}{1 - \hat{\varepsilon}_b(p_b)}(p_b - \theta_1) \geq A(p_s - \theta_1) - A'(p'_s - \theta_1). \tag{4.2}$$

We can bound the right-hand-side of (4.2) using (4.1), which gives:

$$A(p_s-\theta_1)-A'(p_s'-\theta_1)=A(p_s-\theta_2)-A'(p_s'-\theta_2)+(A-A')(\theta_2-\theta_1)\geq(A-A')(\theta_2-\theta_1). \quad (4.3)$$

Since $p_s<p_s'$, inequality (4.1) implies that $A>A'$. According to (4.2) and (4.3), $A>A'$ implies that $T'<T$, which further implies that $T>0$. As a result, type θ_1 must offer p_s with positive probability. This is because otherwise, the buyer will concede immediately following p_s, which contradicts our earlier conclusion that $T>0$. Therefore, type θ_1 must be indifferent between (i) offering p_s and conceding at time 0 and (ii) offering p_s' and conceding at time 0. This implies that $c_b(p_s-p_b)+p_b-\theta_1=c_b'(p_s'-p_b)+p_b-\theta_1$, or equivalently, $c_b(p_s-p_b)=c_b'(p_s'-p_b)$. Since $p_s'>p_s$, the fact that $c_b(p_s-p_b)=c_b'(p_s'-p_b)$ implies that $c_b\geq c_b'$.

Recall our formulas for players' concession rates. The buyer will concede at rate $\lambda_b=\frac{r(p_b-\theta_1)}{p_s-p_b}$ after the seller offers p_s and will concede at rate $\lambda_b'=\frac{r(p_b-\theta_1)}{p_s'-p_b}$ after the seller offers p_s'. Since $p_s'>p_s$, we have $\lambda_b>\lambda_b'$. The expressions for T and T' imply that

$$T=\frac{\log((1-c_b)/\hat{\varepsilon}_b(p_b))}{\lambda_b}<\frac{\log((1-c_b')/\hat{\varepsilon}_b(p_b))}{\lambda_b'}=T'.$$

This contradicts our earlier conclusion that $T'<T$. □

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 values 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 $\min\{p_{\theta_1},\theta_2\}$. Which one she offers depends on the comparison between $\pi(\theta_1)$ and the cutoff π^*. These three steps together establish uniqueness of the equilibrium limit point. We establish the existence of equilibrium in Online Appendix A.

Fix $\pi\in\Delta\{\theta_1,\theta_2\}$ and an equilibrium $(\sigma_b,\sigma_s,\tau_b,\tau_s)$. We abuse notation by using $\sigma_b(p_b)$ and $\sigma_s(p_s|\theta,p_b)$ to denote the probabilities with which σ_b and $\sigma_s(\theta)(p_b)$ assign to offers p_b and p_s, respectively. Let $\hat{\varepsilon}_b(p_b)$ and $\hat{\varepsilon}_s(p_b,p_s)$ be the probabilities with which the buyer and the seller, respectively, are commitment types after observing offers (p_b,p_s). Let $\hat{\pi}_j(p_b,p_s)$ be the probability that the seller is the rational type with production cost θ_j after observing offers (p_b, p_s). Formally,

$$\hat{\varepsilon}_b(p_b) \equiv \frac{\varepsilon\mu_b(p_b)}{\varepsilon\mu_b(p_b) + (1-\varepsilon)\sigma_b(p_b)} \tag{4.4}$$

$$\hat{\varepsilon}_s(p_b, p_s) \equiv \frac{\varepsilon\mu_s(p_s)}{\varepsilon\mu_s(p_s) + (1-\varepsilon)[\pi\{\theta_1\}\sigma_s(p_s|\theta_1, p_b) + \pi\{\theta_2\}\sigma_s(p_s|\theta_2, p_b)]} \tag{4.5}$$

$$\hat{\pi}_j(p_b, p_s) \equiv \frac{(1-\varepsilon)\pi\{\theta_j\}\sigma_s(p_s|\theta_j, p_b)}{\varepsilon\mu_s(p_s) + (1-\varepsilon)[\pi\{\theta_1\}\sigma_s(p_s|\theta_1, p_b) + \pi\{\theta_2\}\sigma_s(p_s|\theta_2, p_b)]}. \tag{4.6}$$

First, we characterize the equilibrium in the continuation game after players offer $(p_b, p_s) \in \mathbf{P}_b \times \mathbf{P}_s$ 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 type-θ seller 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 type-θ_j seller concedes at time 0 with 0 probability and concedes at rate 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 that is strictly above θ_j,

reaches 1 at time:

$$T_s^j \equiv \frac{-\log(\hat{\varepsilon}_s + \sum_{i>j} \hat{\pi}_i)}{\lambda_s}. \tag{4.7}$$

Likewise, if the buyer does not 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

$$L \equiv \frac{-\lambda_s \log \hat{\varepsilon}_b}{-\sum_{j=1}^{m} (\lambda_b^j - \lambda_b^{j+1}) \log(\hat{\varepsilon}_s + \hat{\pi}_{j+1})}. \tag{4.8}$$

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 0 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 0, let

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

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

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 0 so that the seller's belief that the buyer is the commitment type reaches 1 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 type θ_1 concedes at time 0 so that the buyer's belief that the seller is either committed or that $\theta \geq p_b$ reaches 1 at time T_b. However, if $j^* > 1$, it is not sufficient to have type-θ_1 seller conceding with probability 1 at time 0 to make both players finish conceding at the same time. In those cases, we need all types strictly below θ_{j^*} to concede at time 0 with probability 1 and possibly type θ_{j^*} to concede at time 0 with positive probability. As a result, when the seller is the weak player, his concession probability at time 0 equals $\hat{c}_s^{j*}$. We summarize these findings in Lemma 4.2:

Lemma 4.2. *Fix offers $(p_b, p_s) \in \mathbf{P}_b \times \mathbf{P}_s$ 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 4.2 can be found in APS, which we omit in order to avoid repetition.

Next, consider the continuation game when no player concedes at time 0. In equilibrium, for every $j \in \{j^*, ..., m-1\}$, type θ_j will finish conceding at time

$$T^j = T_s^j + \frac{\log(1-c_s)}{\lambda_s}. \tag{4.9}$$

In addition, the rational types of both players will finish conceding at the same time

$$T^m \equiv \min\left\{\frac{-\log(\hat{\varepsilon}_b) - \sum_{j=j^*}^{m-1}(\lambda_b^j - \lambda_b^{j+1})T_s^j}{\lambda_b^m}, T_s^m\right\}. \tag{4.10}$$

Lemma 4.3 characterizes players' equilibrium strategies in the war-of-attrition game:

Lemma 4.3. *In every equilibrium of the war-of-attrition game $\Gamma(p_b, p_s, \hat{\varepsilon}_b, \hat{\varepsilon}_s, \hat{\pi})$ with $(p_b, p_s) \in \mathbf{P}_b \times \mathbf{P}_s$ and $\theta_1 < p_b < p_s < 1$, the buyer's and the seller's concession times τ_b and $\tau_s(\theta)$ 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 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 $\{\varepsilon^k\}_{k=0}^{+\infty}$ that satisfies $\varepsilon^k \to 0$ 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 (4.4), (4.5) and (4.6) 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 $i \in \{b, s\}$.

Lemma 4.4. *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 4.2 in the game $\Gamma(p_b, p_s, \hat{\varepsilon}_b^k, \hat{\varepsilon}_s^k, \hat{\pi}^k)$ with $(p_b, p_s) \in \mathbf{P}_b \times \mathbf{P}_s$ and $\theta_1 < p_b < p_s < 1$, and let (c_b^∞, c_s^∞) be the limit of this sequence as $k \to \infty$. Then,*

1. *If $\hat{\varepsilon}_s^\infty(p_b, p_s) = 0$ and $\lambda_b^2 > \lambda_s$, then $c_s^\infty(p_b, p_s) = 1$.*

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

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

[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.

One consequence of Lemma 4.4 is that as long as the buyer assigns strictly positive probability to type θ_2, the identity of the player who concedes at time 0 is determined by the comparison of concession rates between the buyer and the type-θ_2 seller.

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 4.5 states a key property of the seller's equilibrium strategy.

Lemma 4.5. *For every $\eta > 0$, there exists $\overline{\varepsilon} > 0$ such that, for any p_b such that $\varepsilon/\sigma_b(p_b) < \bar{\varepsilon}$, $\sigma_s(\cdot|p_b, \theta_2)$ assigns probability less than η to offers below $\overline{p}_{\theta_1}(p_b)$.*

Intuitively, Lemma 4.5 says that following any offer made by the buyer with a probability that is bounded away from 0 or vanishes at a slower rate compared to ε, type θ_2's offer is weakly greater than type θ_1's offer with probability 1.

Proof. Fix $\varepsilon > 0$ and any p_b such that $\sigma_b(p_b) > 0$. 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. If $p_b \in (\theta_1, \theta_2]$, the result follows from Lemma 4.1.

Next, suppose that $p_b > \theta_2$, and let $p_2(p_b) \equiv \max\{p \in \mathbf{P}_s : p \leq 1 + \theta_2 - p_b\}$. The compactness of $\mathbf{P}_s$ ensures that this is always well-defined. Let p_s be an offer in the support of $\sigma_s(\cdot|p_b, \theta_1)$, and suppose that there is $p'_s < p_s$ such that $\sigma_s(p'_s|p_b, \theta_2) > \eta$.

Suppose $p_2(p_b) - p'_s > \eta$ for some $\eta > 0$. We argue that it must be the case that type θ_1 offers $p_2(p_b)$ with probability bounded above zero, and that type θ_2 does so with vanishing probability. Suppose by way of contradiction that this is not the case. According to Parts 2 and 3 of Lemma 4.4, for every $\delta > 0$, there is $\overline{\varepsilon}_\delta > 0$ such that $\varepsilon/\sigma_b(p_b) < \overline{\varepsilon}_\delta$ implies that the buyer's time-zero concession probability after the seller offers $p_2(p_b)$ is at least $1 - \delta$. Thus, type θ_2's payoff when he offers $p_2(p_b)$ is at least $(1 - \delta)(p_2(p_b) - \theta_2)$. Since $\delta > 0$ is arbitrary, we can take it to be sufficiently small in which case we get that the high type's payoff when offering $p_2(p_b)$ is strictly higher than the upper bound on his equilibrium payoff $p'_s - \theta_2$. This leads to a contradiction when $\varepsilon/\sigma_b(p_b) < \overline{\varepsilon}_\delta$.

On the other hand, by Part 2 of Lemma 4.4, $p'_s < p_2(p_b)$ implies that for every $\delta > 0$, there exists $\overline{\varepsilon}_\delta$ such that $\varepsilon/\sigma_b(p_b) < \overline{\varepsilon}_\delta$ implies that the buyer's time zero concession probability after the seller offers p'_s is at least $1 - \delta$. Let T_1 be the time at which the rational buyer finishes conceding against the low-type seller (as defined in (4.9)) after the seller offers $p_2(p_b)$, and let c_b be the associated time-zero concession probability for the buyer. Note that $T_1 \to +\infty$ as $\hat{\varepsilon}_b(p_b) \to 0$. For $\varepsilon/\sigma_b(p_b)$

sufficiently small, the low type's incentive to offer $p_2(p_b)$ instead of p'_s implies that

$$c_b(p_2(p_b) - \theta_1) + (1 - c_b)(p_b - \theta_1) \geq p'_s - \theta_1.$$

Type θ_2's payoff from deviating to $p_2(p_b)$ and waiting until T_1 for the buyer to concede is at least

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

As $\varepsilon/\sigma_b(p_b) \to 0$, type θ_2 strictly benefits from deviating to $p_2(p_b)$, which is a contradiction.

If $p'_s > p_2(p_b)$, then Part 1 of Lemma 4.4 implies that as $\varepsilon \to 0$, type θ_2 concedes with positive probability at time zero, and therefore his equilibrium payoff is $p_b - \theta_2$. An analogous argument to the one in the previous paragraph then implies that type θ_2 obtains a higher payoff from deviating to $p_2(p_b)$. It remains to consider the possibility that $p'_s = p_2(p_b)$. If this is the case, then the buyer concedes at time 0 with probability close to 1 by Part 2 of Lemma 4.4. 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 $p_2(p_b)$. □

We now describe the seller's equilibrium strategy σ_s. First, we show that when the buyer offers $p_b \in (\theta_2, p_{\theta_2}]$ such that $\varepsilon/\sigma_b(p_b)$ is arbitrarily close to zero, both types of the seller will offer the same price that is close to $p_2(p_b)$. Suppose by way of contradiction that type θ_2 offers some price p_s that is strictly bounded below $p_2(p_b)$. According to Lemma 4.5, 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 converging to 1 at time 0 by Lemma 4.4. This leads to a contradiction. Similarly, if type θ_2 offers $p_s > p_2(p_b)$ with probability bounded away from zero, he has to concede to p_b at time 0, but he obtains a strictly higher payoff by offering $p_2(p_b)$. It then follows that type θ_1 finds it optimal to pool with type θ_2 by offering $p_2(p_b)$.

If the buyer offers anything weakly greater than p_{θ_2}, then her payoff can be made arbitrarily close to $1 - p_b$ by taking ε to 0. This is because the seller is in a weak bargaining position if he offers anything greater and, by Lemma 4.4, would have to concede with probability converging to 1 as $\varepsilon \to 0$. Combining this with the above arguments, it follows that any offer $p_b \in \mathbf{P}_b$ such that $p_b > \theta_2$ and $\varepsilon/\sigma_b(p_b) < \bar{\varepsilon}$ must be arbitrarily close to p_{θ_2}, which is the price that maximizes the buyer's payoff $1 - \max\{p_b, p_2(p_b)\}$ over $p_b \in \mathbf{P}_b \cap (\theta_2, 1]$.

If the buyer offers $p_b \in \mathbf{P}_b \cap (\theta_1, \theta_2]$, then consider two cases. First, suppose the rational-type buyer offers p_b with zero probability in equilibrium. Then, the seller's posterior belief assigns probability 1 to the commitment type after observing p_b, after which type θ_1 will concede immediately and the buyer's payoff is at least $\pi(\theta_1)(1-p_b)$. Second, suppose the rational-type buyer offers p_b with positive probability in equilibrium, in which case Lemma 4.1 implies that type θ_2 seller will counteroffer 1 with probability 1. Next, we derive the low type's offer. If $\hat{\varepsilon}_b(p_b)$ is bounded above zero, the low type has to concede immediately following any offer in the support of his equilibrium strategy. The following lemma addresses the case where $\hat{\varepsilon}_b(p_b)$ vanishes with ε. Let $p_1(p_b) \equiv \max\{p \in \mathbf{P}_s : p \leq 1+\theta_1-p_b\}$.

Lemma 4.6. *For every $\eta > 0$, there exists $\bar{\varepsilon} > 0$ such that, for any $p_b \leq \theta_2$ such that $\varepsilon/\sigma_b(p_b) < \bar{\varepsilon}$, $\sigma_s(\cdot|p_b, \theta_1)$ is η-close to the Dirac measure on* $\max\{p_b, p_1(p_b)\}$.

Proof. If $\hat{\varepsilon}_b(p_b) \to 0$ as $\varepsilon \to 0$ and $p_b > p_1(p_b)$, then any offer above p_b induces type θ_1 to concede at time 0 with probability converging to 1. If $p_b \leq p_1(p_b)$, type θ_1 can induce the buyer to concede with probability arbitrarily close to 1 by offering $p_1(p_b)$. If he offers anything greater than $p_1(p_b)$ with probability bounded above 0, he will concede at time 0 with positive probability, which is weakly dominated by offering $p_1(p_b)$ (strictly so unless $p_b = p_1(p_b)$). Therefore, type θ_1 will offer $p_1(p_b)$ with probability converging to 1. □

Combining all cases, the buyer's payoff when she offers $p_b \in \mathbf{P}_b \cap (\theta_1, \theta_2]$ is bounded below by $\pi(\theta_1)(1-\max\{p_1(p_b), p_b\})$ and the lower bound is attained if the buyer offers p_b with non-vanishing probability in equilibrium. This expression is uniquely maximized on $p_b \in (\theta_1, \theta_2]$ at $\min\{p_{\theta_1}, \theta_2\}$.

Let $p_1^\nu = \arg\max_{p \in \mathbf{P}_b} (1-\max\{p_1(p_b), p_b\})$, and $p_2^\nu = \arg\max_{p \in \mathbf{P}_b} (1-\max\{p_2(p_b), p_b\})$. Our earlier arguments imply that for a fixed $\nu > 0$, as ε goes to zero, the buyer's equilibrium strategy is supported on some subset of $p_1^\nu \cup p_2^\nu$. Moreover, for $\nu > 0$ sufficiently small, if $\pi(\theta_1) < \pi^*$, the buyer strictly prefers offering $p \in p_2^\nu$ over $p \in p_1^\nu$, 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 offer $p \in p_1^\nu$. Therefore, the equilibrium strategies must converge to those described in part two of Theorem 1. If $p_{\theta_1} \leq \theta_2$, then in the limit, 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 $p_1(\theta_2) > \theta_2$ after which players play the war-of-attrition game. However, Lemma 4.4 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 3.2, the incentive constraints of the seller pin down the expected welfare loss from delay to be approximately given by (3.7). To complete the argument provided there, we show how to derive (3.10) from (3.9). In order to avoid the buyer from conceding immediately after the seller demands $p_s \in \mathbf{P}_s$ with $p_s > 1-\varepsilon$, which would in turn give rise to a profitable deviation for the seller, it must be that type θ_1 demands p_s with positive probability (by Lemma 4.4). As a result, type θ_1's incentive constraint requires that $c_b p_s + (1-c_b)\min\{p_{\theta_1}, \theta_2\} - \theta_1 \geq \max\{p_{\theta_1}, 1-\theta_2+\theta_1\} - \theta_1$. Plugging this into (3.9) and using the fact that $T_1 \to +\infty$ as $\varepsilon \to 0$, we obtain (3.10).

4.2 Proof of Theorem 2

This section establishes the common properties of all equilibria. We establish the existence of equilibrium in Online Appendix B. Throughout this proof, we use V_θ to denote the equilibrium payoff of type θ taking the adoption cost into account, and we use $\pi(\theta_1)$ to denote the seller's equilibrium adoption probability. The following series of lemmas provide necessary conditions for the limiting equilibria under every parameter configuration, which together establish Theorem 2.

First, consider the case in which $c > \theta_2 - \theta_1$, that is, the cost of adoption is strictly greater than the social benefit from adoption. We show that the seller adopts with zero probability in every equilibrium. This in turn implies that the buyer has no incentive to offer anything that is strictly lower than θ_2. As a result, she will offer approximately $p_{\theta_2} \equiv \frac{1+\theta_2}{2}$ in every equilibrium as $\varepsilon \to 0$.

Lemma 4.7. *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 is 0 and the expected delay is less than η in evrey equilibrium.*

Proof. Suppose by way of contradiction that the seller adopts with positive probability. His payoff in the bargaining stage after he adopts is $V_{\theta_1} \equiv \mathbb{E}[e^{-r\tau}(p-\theta_1)|\theta = \theta_1]$, where τ is the time of trade and p is the trading price. If the seller deviates to not adopting and uses type θ_1's strategy in the war-of-attrition game, then he can secure a payoff of $\mathbb{E}[e^{-r\tau}(p-\theta_2)|\theta = \theta_1]$. As a result,

$$V_{\theta_1} - c - V_{\theta_2} \leq \mathbb{E}[e^{-r\tau}|\theta = \theta_1](\theta_2 - \theta_1) - c \leq (\theta_2 - \theta_1) - c < 0,$$

which implies that the seller strictly prefers not to adopt. This leads to a contradiction. $\square$

The rest of this proof considers the case in which $c < \theta_2 - \theta_1$. Lemma 4.8 examines the subcase 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 is that the gap between θ_2 and θ_1 is large enough so that the buyer benefits from screening under certain values of $\pi(\theta_1)$. The second condition implies that the adoption cost is neither too high nor too low, which ensures 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. Lemma 4.8 characterizes the unique limiting equilibrium outcome under these two conditions.

Lemma 4.8. *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 to π^* and the expected delay is bounded above 0.*

Proof. Suppose by way of contradiction that $|\pi(\theta_1) - \pi^*| > \eta$. If $\pi(\theta_1) > \pi^*$, we show that the buyer's offer converges in distribution to p_{θ_1}. According to Lemma 4.1, if the buyer offers $p_b \leq \theta_2$ with positive probability, then type θ_2 will offer 1, and according to Lemma 4.6, there exist $\overline{\nu} > 0$ and $\overline{\varepsilon}_\nu > 0$ such that $\nu < \overline{\nu}$ and $\varepsilon < \overline{\varepsilon}_\nu$ implies that type θ_1 will counteroffer $\max\{p_b, p_1(p_b)\}$ with probability converging to 1 as $\varepsilon \to 0$. Therefore, if the buyer offers any $p_b \leq \theta_2$, then she can secure a payoff of $\pi(\theta_1)(1 - \max\{p_b, p_1(p_b)\})(1 - \eta)$, which, for η sufficiently small, is maximized by choosing $p_b^* \in \mathbf{P}_b \cap (p_{\theta_1} - \nu, p_{\theta_1} + \nu)$. If $\pi(\theta_1) > \pi^*$ and ν is sufficiently small, the resulting payoff is strictly greater than the buyer's highest payoff when she offers $p_b > \theta_2$, which is arbitrarily close to $1 - p_{\theta_2}$. Therefore, for every $\eta > 0$, there exist ν and ε close enough to 0 so that in every equilibrium, the buyer's offer will belong to $\mathbf{P}_b \cap (p_{\theta_1} - \nu, p_{\theta_1} + \nu)$ with probability more than $1 - \eta$.

If the seller instead chooses not to adopt and then demands $p_s \in \mathbf{P}_s$ that is arbitrarily close to 1 after the buyer offers $p_b^* < \theta_2$ and waits for the buyer to concede, then he can secure a payoff that converges to $\frac{1-\theta_2}{2}$ as $\varepsilon \to 0$. In order to see this, fix any $p_s \in \mathbf{P}_s$ that is close to 1. For any offer p_b^* that the buyer makes in equilibrium with probability bounded above zero, the posterior belief that the buyer is the commitment type is arbitrarily close to 0 when ε is sufficiently small. Therefore, the buyer will concede at time 0 with probability arbitrarily close to 1 unless type θ_1 offers p_s with positive probability. Type θ_1's incentive to offer p_s after the buyer offers p_b^* implies that

$$p_b^* - \theta_1 + c_b^*(\max\{p_b^*, p_1(p_b^*)\} - p_b^*) = e^{-rT}\hat{\varepsilon}_b(p_b^*)(p_b^* - \theta_1) + (1 - \hat{\varepsilon}_b(p_b^*))A(p_s - \theta_1),$$

where c_b^* is the probability that the buyer concedes at time 0 after the seller offers $\max\{p_b^*, p_1(p_b^*)\}$,[12] and A and T are the discounted concession probabilities of the rational-type buyer and the time at which the rational-type buyer finishes conceding following the seller's offer p_s.

Take the limit as $\varepsilon \to 0$ and $p_s \to 1$, the value of A converges to $\frac{\max\{p_b^*, p_1(p_b^*)\}-\theta_1}{1-\theta_1}$, which converges to $1/2$ as $\nu \to 0$. Thus, for every $\eta > 0$ and ν close enough to 0, there exists $\overline{\varepsilon}_\nu > 0$ such that $\varepsilon < \overline{\varepsilon}_\nu$ implies that type θ_2's payoff when he offers p_s is at least $\frac{1-\theta_2}{2} - \eta$. As $\eta \to 0$, this payoff lower bound is strictly greater than $p_{\theta_1} - \theta_1 - c$ whenever $c > \frac{\theta_2-\theta_1}{2}$. This contradicts $\pi(\theta_1) > 0$.

Next, suppose $\pi(\theta_1) < \pi^*$. Fix any small enough $\varepsilon > 0$, the buyer's incentive constraint requires that any price that she offers with probability bounded above 0 satisfies $p_b^* \in (p_{\theta_2} - \nu, p_{\theta_2} + \nu)$, and therefore $V_{\theta_1} - c$ converges to $p_{\theta_2} - \theta_1 - c$ and V_{θ_2} converges to $p_{\theta_2} - \theta_2$ as $\varepsilon \to 0$. Hence, for sufficiently small $\nu > 0$ and $\varepsilon > 0$, we obtain that $V_{\theta_1} - c > V_{\theta_2}$, where the inequality follows from $c < \theta_2 - \theta_1$. This contradicts $\pi(\theta_1) < 1$.

Therefore, in any equilibrium, the seller must adopt with probability close to π^*. As in the proof of Theorem 1, if ε is small enough, then the buyer's sequential rationality constraint requires that any price that she offers with probability bounded above 0 must be within an η-neighborhood of an element of $p_1^\nu \cup p_2^\nu$, where $p_i^\nu \equiv \arg\max_{p \in \mathbf{P}_b} (1 - \max\{p_i(p_b), p_b\})$. After the buyer offers $p \in p_2^\nu$, trade happens with negligible delay at this price. After she offers $p \in p_1^\nu$, there is trade with negligible delay conditional on the seller's cost being θ_1, and there is an expected delay converging to $\frac{1}{2}$ when the seller's cost being θ_2. Consequently, let ρ^* denote the limiting probability with which the buyer offers p_{θ_2}, type-θ_1 seller's payoff in the bargaining stage converges to $\rho^*(p_{\theta_2} - \theta_1) + (1-\rho^*)(p_{\theta_1} - \theta_1)$ and type θ_2's payoff converges to $p_{\theta_2} - \theta_2$. The seller's indifference condition at the adoption stage requires that

$$\rho^* = \frac{2c - (\theta_2 - \theta_1)}{\theta_2 - \theta_1}. \tag{4.11}$$

Note 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 converging to π^* cannot be sustained if $p_{\theta_1} < \theta_2$ and $c \notin \left(\frac{\theta_2-\theta_1}{2}, \theta_2 - \theta_1\right)$. This conclusion will be used in the proof of Lemma 4.10. Therefore, the expected delay in the limit where $\varepsilon \to 0$ for a fixed ν, and then taking the limit as $\nu \to 0$ is

$$(1 - \pi^*)(1 - \rho^*)\frac{1}{2} = \frac{\theta_2 - \theta_1 - c}{1 - \theta_1} > 0. \tag{4.12}$$

[12] If $\max\{p_b^*, p_1(p_b^*)\} = p_b^*$, then we set $c_b^* = 1$.

□

Lemma 4.8 establishes the third part of Theorem 2. In order to show the second part, we need to consider the case where (3.6) 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 is intermediate. Lemma 4.9 characterizes the limiting equilibria in this case.

Lemma 4.9. *If (θ_1, θ_2) satisfies (3.6), $p_{\theta_1} > \theta_2$, 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 in every equilibrium when $\varepsilon \in (0, \bar{\varepsilon}_\nu)$, 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. First, suppose that $\pi^* - \pi(\theta_1) > \eta$. Then, analogous to Lemma 4.8, type-θ_1 seller's equilibrium payoff converges to $p_{\theta_2} - \theta_1 - c$ and type-θ_2 seller's equilibrium payoff converges to $p_{\theta_2} - \theta_2$. Therefore, the seller strictly prefers to adopt the technology at cost $c \in \left(\frac{(1-\theta_2)(\theta_2-\theta_1)}{1-\theta_1}, \theta_2 - \theta_1\right)$. This contradicts our earlier conclusion that $\pi(\theta_1) < 1$.

Next, suppose that $\pi(\theta_1) \in (\pi^*, 1)$ and $\pi(\theta_1)$ is bounded away from both π^* and 1. Following the same argument as in the proof of Lemma 4.8, we know that the buyer will offer θ_2 with probability converging to 1 as $\varepsilon \to 0$. As a result, the limiting equilibrium payoff of type θ_1 equals $1-\theta_2-c$, and the limiting equilibrium payoff of type θ_2 equals $\frac{1-\theta_2}{1-\theta_1}(1-\theta_2)$. The assumption that $c > \frac{(1-\theta_2)(\theta_2-\theta_1)}{1-\theta_1}$ then implies that, for ε and ν sufficiently small, type θ_2's payoff is strictly more than that of type θ_1's. This contradicts our earlier hypothesis that $\pi(\theta_1) > \pi^* \geq 0$.

Therefore, the equilibrium adoption probability must converge to either π^* or 1 in the limit where $\varepsilon \to 0$. If the limit point is π^*, then it must be the case that the buyer mixes between offering θ_2 and p_{θ_2} in a way that makes the seller indifferent between adopting and not adopting the technology. As in the proof of Lemma 4.8, let ρ^* denote the limiting probability with which the buyer offers p_{θ_2}. The seller's incentive constraint at the adoption stage requires that in the limit as $\varepsilon \to 0$ for a fixed $\nu > 0$, and then taking the limit as $\nu \to 0$

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

or equivalently,

$$\rho^* = \frac{(1-\theta_1)c - (1-\theta_2)(\theta_2-\theta_1)}{(\theta_2-\theta_1)^2}. \tag{4.13}$$

With $\rho^* \in (0,1)$, we have $c \in \left(\frac{(1-\theta_2)(\theta_2-\theta_1)}{1-\theta_1}, \theta_2 - \theta_1\right)$. The discounted expected time at which players reach an agreement, in the limit as $\varepsilon \to 0$ for a fixed ν, and then taking the limit as $\nu \to 0$, is then given by

$$(1-\pi^*)(1-\rho^*)\frac{1-\theta_2}{1-\theta_1} = \frac{(3\theta_2 - 1 - 2\theta_1)(\theta_2 - \theta_1 - c)}{2(\theta_2-\theta_1)^2} > 0. \tag{4.14}$$

The above conditions must be satisfied in any equilibrium in which the seller's adoption probability is arbitrarily close to π^*. Next, we show that such an equilibrium exists. Let $\pi(\varepsilon,\nu) \in (0,1)$ (which is arbitrarily close to π^*) denote the adoption rate that makes the buyer exactly indifferent between offering $p_b \in p_1^\nu$ and $p_b \in p_2^\nu$. As we show in Online Appendix A, the continuation bargaining game with $\pi(\theta_1) = \pi(\varepsilon,\nu)$ has an equilibrium. Thus, it suffices to show that, given players' strategies in the bargaining game, the seller is indifferent at the adoption stage and hence that $\pi(\varepsilon,\nu)$ can be sustained in equilibrium. As argued in the previous paragraphs, when ε and ν are arbitrarily small, on-path bargaining strategies are arbitrarily close to: the buyer offers p_{θ_2} with probability $\rho^* \in [0,1]$ and θ_2 with probability $1-\rho^*$; after the buyer offers p_{θ_2}, both seller types accept, and after the buyer offers θ_2, the low type demands $1+\theta_1-\theta_2$ and the high type demands 1. Moreover, for small enough $\varepsilon > 0$ and $\nu > 0$, we know that if $\rho^* = 1$, then the seller's equilibrium payoff from adopting is strictly greater than his payoff from not adopting, while the opposite is true if $\rho^* = 0$. By continuity, there exists $\rho(\varepsilon,\nu) \in (0,1)$ such that the seller is indifferent between adopting and not adopting the technology, and therefore the adoption probability $\pi(\varepsilon,\nu)$ can be sustained as the seller's adoption strategy in an equilibrium of the game.

If the limit point for $\pi(\theta_1)$ is 1, then the same argument implies that the buyer's offer cannot converge to θ_2 as $\varepsilon \to 0$. This is because such an offer would give rise to a profitable deviation for the seller at the adoption stage. In order to rule out the buyer benefiting from offering θ_2, it must be the case that $1-\pi(\theta_1)$ converges to 0 *faster* than $\varepsilon \to 0$.

We now construct one such equilibrium, where in the limit the buyer offers p_{θ_1} and V_{θ_1} converges to $p_{\theta_1} - \theta_1$. Conditional on not adopting and the buyer offering $p_b \in (p_{\theta_1} - \nu, p_{\theta_1} + \nu) \cap \mathbf{P}_b$, the seller counteroffers $p_2(p_b)$ which is arbitrarily close to $1+\theta_2-p_{\theta_1}$ by making ν sufficiently small. In order to deter type θ_1 from deviating from offering $p_1(p_b)$ to offering some $p_2(p_b)$ that is strictly greater than $p_1(p_b)$, it must be that the buyer concedes with probability 0 after the seller offers $p_2(p_b)$. This condition pins down the probability with which type θ_1 offers $p_2(p_b)$ in equilibrium, which we denote by β.

Let $T_1 \in \mathbb{R}_+$ denote the time at which type θ_1 finishes conceding following offers p_b and $p_2(p_b)$. As $\varepsilon \to 0$, type-θ_2 seller's equilibrium payoff after the buyer offers $p_b \in (p_{\theta_1} - \nu, p_{\theta_1} + \nu) \cap \mathbf{P}_b$

converges to

$$(1-e^{-(r+\lambda_b^1)T_1})\frac{(1-p_b)(p_2(p_b)-\theta_2)}{p_2(p_b)-\theta_1} \leq \frac{(1-p_b)(p_2(p_b)-\theta_2)}{p_2(p_b)-\theta_1}. \tag{4.15}$$

As $\nu \to 0$, the right-hand-side of (4.15) converges to $\frac{(1-p_{\theta_1})^2}{1+\theta_2-p_{\theta_1}-\theta_1}$.

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

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

where $\lambda_s = \frac{r(1-p_2(p_b))}{p_2(p_b)-p_b}$. The above condition pins down the value of $\pi(\theta_1)$. 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 the case that T_1 is bounded above. This is because otherwise, the seller will strictly prefer not to adopt. This in turn requires that $\pi(\theta_1) \to 1$ as $\varepsilon \to 0$.

On the other hand, type-θ_2 seller can secure payoff $\mathbb{E}[p_b] - \theta_2$, which converges to $p_{\theta_1} - \theta_2$. Hence, the seller is indifferent between adopting and not adopting only if $c < \theta_2 - \theta_1$. Finally, since the seller's adoption probability converges to 1 and there is negligible delay conditional on the seller adopting the technology, the expected delay is less than η when ε is small enough. $\square$

Lemma 4.9 implies that first, as stated in Theorem 2, when (3.6) 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 significant delay in reaching agreement. This equilibrium shares the same features as the unique equilibrium characterized in Lemma 4.8.

Second, there are multiple limiting equilibria. When $\theta_2 < p_{\theta_1}$ and the cost of adoption is close to $\theta_2 - \theta_1$, we can sustain the approximately efficient equilibrium since whenever the buyer expects the seller to adopt with probability close to 1, her optimal strategy is to offer p_{θ_1}. Since $\theta_2 < p_{\theta_1}$, it must be the case that $\tau_s < +\infty$ if the seller does not adopt. Therefore, conditional on not adopting, the seller will concede in finite time and the expected delay can make him indifferent between adopting and not adopting. The fact that $\pi(\theta_1)$ is close to 1 ensures that this delay can be sustained as an outcome of the war-of-attrition game. However, this reasoning does not apply when $\theta_2 > p_{\theta_1}$. This is because conditional on not adopting, type θ_2 has no incentive to concede when the buyer offers p_{θ_1}. This leads to a profitable deviation for type θ_1 from offering p_{θ_1}.

Finally, we consider the case in which either (3.6) is violated or the cost of adoption is sufficiently low. We show that investment is efficient and there is almost no delay in reaching agreement.

Lemma 4.10. *If the parameters of the model satisfy (3.6) and* $c < \max\left\{\frac{1}{2}, \frac{1-\theta_2}{1-\theta_1}\right\}(\theta_2 - \theta_1)$*, or if (3.6) 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 at least* $1 - \eta$ *and the expected welfare loss from delay is no more than* η*.*

Proof. Suppose first that (3.6) is violated and $c < \theta_2 - \theta_1$. If the seller adopts with probability strictly less than one, then the buyer's incentive constraint implies that her equilibrium offer converges to p_{θ_2}. Therefore, type θ's payoff in the bargaining stage converges to $p_{\theta_2} - \theta$. The assumption that $c < \theta_2 - \theta_1$ then implies that the seller strictly prefers to adopt. This contradicts our earlier conclusion that $\pi(\theta_1) < 1$.

Next, suppose that (3.6) is satisfied and $c < \max\left\{\frac{1}{2}, \frac{1-\theta_2}{1-\theta_1}\right\}(\theta_2 - \theta_1)$. In the subcase where $\pi(\theta_1) < \pi^*$, the same argument as in the previous paragraph applies. Moreover, our proofs of Lemmas 4.8 and 4.9 imply that an inefficient equilibrium with limiting adoption probability equal to π^* exists only if $c > \max\left\{\frac{1}{2}, \frac{1-\theta_2}{1-\theta_1}\right\}(\theta_2 - \theta_1)$, which is ruled out under the assumption in Lemma 4.10. In the subcase where $\pi(\theta_1) \in (\pi^*, 1)$ and is bounded away from both 1 and π^*, then the buyer's incentive constraint requires that when $\varepsilon \to 0$, her offer in equilibrium belong to $\mathbf{P}_b \cap (\{\min\{\theta_2, p_{\theta_1}\} - \nu, \{\min\{\theta_2, p_{\theta_1}\} + \nu)$ with probability more than $1 - \eta$. As a result, the seller's payoff from adopting the technology converges to $\max\{p_{\theta_1}, 1 + \theta_1 - \theta_2\} - \theta_1 - c$ and his payoff from not adopting the technology converges to $\max\left\{\frac{1}{2}, \frac{1-\theta_2}{1-\theta_1}\right\}(1 - \theta_2)$. Hence, the limiting payoff gain from adoption is

$$\max\left\{\frac{1}{2}, \frac{1-\theta_2}{1-\theta_1}\right\}(\theta_2 - \theta_1) - c > 0,$$

which follows from the assumption that $c < \max\left\{\frac{1}{2}, \frac{1-\theta_2}{1-\theta_1}\right\}(\theta_2 - \theta_1)$. Hence, the seller's adoption probability also converges to 1 as $\varepsilon \to 0$.

Finally, Lemmas 4.2 and 4.3 imply that the outcome in the bargaining stage is efficient in the limit conditional on the seller's production cost is θ_1. This implies that the expected welfare loss from delay converges to 0 as ε vanishes. □

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. These assumptions are without loss of generality since a technology will never be adopted in any equilibrium if it costs more than a more efficient technology. We make a generic assumption that there is a unique *socially efficient technology* and focus on the interesting case that the socially efficient technology is not the default one. Formally, we assume that there exists $j^o < n$ such that $\{j^o\} = \arg\min_{k\in\{1,2,\ldots,n\}} \Big\{\theta_k + c_k\Big\}$.

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

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

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.

For any $i, j \in \{1, ..., n\}$ such that $i < j$, let $\pi[\theta_i, \theta_j]$ be the probability that $\theta \in [\theta_i, \theta_j]$. Theorem 3 characterizes the unique limiting equilibrium of the reputational bargaining game with an exogenous cost distribution under the generic conditions that

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

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. 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*

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

The proof is in Online Appendix D, which is similar to the one for Theorem 1 except that we need to establish a general version of Lemma 4.1 that allows for three or more production costs. The details can be found in Online Appendix C. According to Theorem 3, the qualitative features of the equilibrium in this general environment are similar to the ones in which there are two costs. 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 the proof of Theorem 1, the optimal way in which the buyer can screen types with cost no more than θ_i is by offering $p_b \in \mathbf{P}_b \cap (\min\{p_{\theta_i}, \theta_{i+1}\} - \nu, \min\{p_{\theta_i}, \theta_{i+1}\} + \nu)$. This ensures the buyer a payoff of $\pi[\theta_1, \theta_i](\min\{p_{\theta_i}, \theta_{i+1}\} - \theta_i)$ in the limit.

According to our theorem, the buyer will screen the seller in equilibrium if and only if $i^* < n$, in which case she will offer a price that is strictly lower than θ_n instead of offering a price that is close to $p_{\theta_n} \equiv \frac{1+\theta_n}{2}$. Conditional on the buyer offering p_b in a ν-neighborhood of $\min\{p_{\theta_{i^*}}, \theta_{i^*+1}\} \leq \theta_n$, 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. The expected delay in (5.2) is pinned down by the conditions that after the buyer offers $\min\{p_{\theta_{i^*}}, \theta_{i^*+1}\}$, (i) type θ_{i^*} does not benefit from demanding 1, and (ii) type θ_{i^*+1} does not 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 strictly prefers to make a pooling offer p_{θ_n}, after which all types of the seller accept immediately, than any screening offer. As in the two-cost environment, trade happens immediately at a price of approximately p_{θ_n}.

Next, we describe the limiting equilibria in the game with *endogenous* technology adoption. 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. We provide a condition on the gap between different types of the seller's production costs under which the buyer may benefit from making screening offers *under some distribution over cost types.* This is the case if and only if

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

Condition (5.3) generalizes (3.6) to the setting with more than two production technologies. As we explain below, if it is violated, then 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 E, 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 $\mathrm{supp}(\pi) = \{\hat{\theta}_1, ..., \hat{\theta}_m\}$, 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, the buyer will offer $\hat{\theta}_{i+1}$ when she screens the seller with production cost less than $\hat{\theta}_i$. If the buyer screens the seller by offering $\hat{\theta}_{i+1}$ for some $i \in \{1, ..., m-1\}$, then her payoff is

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

This implies that any such offer is strictly dominated by offering $p_{\hat{\theta}_m}$. Given that the equilibrium price offered by the buyer is arbitrarily close to $p_{\hat{\theta}_m}$, the seller's equilibrium payoff converges to $p_{\hat{\theta}_m} - \hat{\theta}_m - \hat{c}_m \leq p_{\hat{\theta}_m} - \theta_{j^o} - c_{j^o}$, 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 3.3. To do this, the seller's adoption strategy must be such that, in the limit, the buyer is indifferent between offering p_{θ_n} and $\min\{p_{\theta_{j^o}}, \theta_n\}$, which yields that $\pi(\theta_{j^o})$ converges to $\frac{p_{\theta_n} - \theta_n}{\min\{p_{\theta_{j^o}}, \theta_n\} - \theta_{j^o}}$. Condition (5.3) implies 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 does not 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 $p_{\theta_{j^o}} < \theta_n$, 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 limiting 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 that can be made arbitrarily close to $\frac{1-\theta_n}{2}$. 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: 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), 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.

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 the same bargaining power in the reputational bargaining game. 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 so that $p_{\theta_1} < \theta_2$, which is the case when the buyer's bargaining power is relatively high. If $\pi(\theta_1)$ is sufficiently high, then the limiting equilibrium in the game with exogenous production costs is inefficient, in which the buyer offers p_{θ_1} and trades with delay. Otherwise, the limiting equilibrium is efficient. When the seller's adoption decision is endogenous, the adoption probability is bounded below one and there is significant delay in reaching agreement *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 $c < \theta_2 - \theta_1$, in which case adopting the technology is socially optimal, and r_b/r_s is arbitrarily small, the limiting equilibrium is inefficient except for an interval of adoption costs of vanishing Lebesgue 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 reduction in price 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.

The above discussion highlights a stark contrast once we compare equilibrium welfare in the extreme cases where one of the player's discount rate is arbitrarily greater than their opponent's. In order to see this, let us focus on the case in which $c < \theta_2 - \theta_1$, that is, it is socially efficient for the seller to invest. If the buyer is arbitrarily more patient than the seller, then in the unique limiting equilibrium, there is under adoption and costly 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, David Pearce and Ennio Stacchetti (2015) "One Sided Uncertainty and Delay in Reputational Bargaining," *Theoretical Economics*, 10(3), 719-773.

[3] 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.

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

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

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

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

[8] Che, Yeon-Koo and Donald Hausch (1999) "Cooperative Investments and the Value of Contracting," *American Economic Review*, 89(1), 125-147.

[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 (2023) "Outside Options, Reputations, and the Partial Success of the Coase Conjecture," Working Paper.

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

[19] Fudenberg, Drew and David Levine (1989) "Reputation and Equilibrium Selection in Games with a Patient Player," *Econometrica*, 57, 759-778.

[20] Gerardi, Dino, Johannes Hörner and Lucas Maestri (2014) "The Role of Commitment in Bilateral Trade," *Journal of Economic Theory*, 154, 578-603.

[21] 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.

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

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

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

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

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

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