Smart Contracts and the Coase Conjecture*

Thomas Brzustowski†, Alkis Georgiadis-Harris‡ and Balázs Szentes§

July 23, 2021

Abstract

This paper reconsiders the problem of a durable-good monopolist who cannot make intertemporal commitments. The buyer's valuation is binary and his private information. The seller has access to dynamic contracts and, in each period, decides whether to deploy the previous period's contract or to replace it with a new one. The main result of the paper is that the Coase Conjecture fails: the monopolist's payoff is bounded away from the low valuation irrespective of the discount factor.

1 Introduction

The Coase Conjecture is a manifestation of the striking consequences of the lack of intertemporal commitment power: if the monopolist can post prices frequently, she clears the market quickly, at prices close to the lowest possible willingness-to-pay even when most consumers have high valuation for the good.[1] The goal of this paper is to examine the extent to which this conclusion is robust to considering more complex selling mechanisms than just price-posting. Motivated by *smart contracts* used in digital markets, we allow the seller to offer general dynamic contracts. Our main result is that if the monopolist has access to such contracts, the Coase Conjecture no longer holds.

In our model, there is a seller of a single good and a buyer. The buyer's valuation for the good is binary, high or low, and it is his private information[2]. We consider the case where the probability

*We are grateful to Daniel Garrett, Jacob Leshno, Niccolò Lomys, John Moore, Francesco Nava, Doron Ravid, Takuro Yamashita, and especially to Phil Reny for helpful comments. Our work was inspired by many conversations with Laura Doval and Vasiliki Skreta.

†Department of Economics, London School of Economics, London, UK. E-mail: t.p.brzustowski@lse.ac.uk.

‡Department of Economics, London School of Economics, London, UK. E-mail: a.georgiadis-harris@lse.ac.uk.

§Department of Economics, London School of Economics, London, UK. E-mail: b.szentes@lse.ac.uk.

[1] This phenomenon was first described by Coase (1972), and later formalized by Stokey (1982), Fudenberg, Levine and Tirole (1985), and Gul, Sonnenschein and Wilson (1986).

[2] We focus on the "gap case" and assume that the seller's production cost, normalized to be zero, is smaller than the low valuation. In the "no-gap case", Ausubel and Deneckere (1989) show that the Coase Conjecture fails even with posted prices.

of high valuation is large enough for the static monopoly price to be the high valuation. Time is discrete and both parties discount the future at the same rate. In the initial period, the seller offers a contract from a space described below. If the buyer accepts the contract, it determines the probabilities of trade and the transfers in subsequent periods until it is replaced. At the beginning of each period, the seller decides whether to proceed with the current contract or to void it and offer a new one.

We use the analogy of a mediator to describe a typical contract from our contract space. In each period, both the seller and the buyer may send messages to the mediator. In turn, the mediator sends private (and possibly public) signals to the contracting parties and implements an allocation. Perhaps the most notable feature of such a contract is that the mediator can posses information which the seller does not. When the seller abandons a contract, she loses that information. This feature is shown to be the driving force of our main result.

Some aspects of our contract space are reminiscent of the technologies developed in relation to the aforementioned smart contracts. First, smart contracts are automated in the sense that they execute trades without further consents from the contracting parties. Similarly, in our model, the allocation proposed by the mediator is implemented and cannot be renegotiated. Second, smart contracts in practice can be, and often are switched off just like the seller can abandon her current contract in our model. One of the reasons that contracts in digital markets are designed so that they can be switched off is to avoid the execution of unlawful transactions, for example, due to bankruptcy procedure against a contracting party.[3] Finally, we note that cryptographic encoding of a party's input can prevent the other contracting party to recover that input even if she has access to the contract's code. Such encoding also plays an important role in digital markets: smart contracts deployed on blockchain networks use cryptographically signed transactions.[4]

Since the seller's commitment power is limited, she may benefit from a small contract space. The reason is that removing contracts from the seller's action space makes the set of possible deviations shrink which, in turn, may enable the seller to stick with contracts which are advantageous from the ex ante perspectives. This can be seen most vividly by considering the scenario when each contract available to the seller specifies trading at the high valuation. In this case, the seller could achieve the full-commitment profit because, even though she maybe tempted to lower the price if there is no trade, she cannot do so. So, in order to model the consequences of limited commitment in a meaningful way, the contract space should be rich enough. To this end, we assume that the seller has access to all *simple and direct* contracts defined as follows. A contract is called simple and direct if the contract elicits the buyer's valuation in the initial period of its deployment and does not communicate with the contracting parties ever after. Our main result

[3]Another context in which a contracting party retains her right to void smart contracts is where the issuer deploys these contracts in her private blockchains. Examples for such issuers include Walmart, Comcast, Spotify, DHL, JPMorgan and MetLife.

[4]Communications in digital trading platforms are typically public. We discuss the implementation of private communications via smart contracts in the concluding section.

holds as long as the seller's contract space is rich enough to include all such contracts.

We note that our model assumes a certain amount of commitment power of the monopolist. Namely, if the contract in place determines an allocation, that allocation will be implemented and the seller cannot take further actions. The same assumption is maintained in the standard Coasian model: if the buyer is willing to buy at the posted price, trade will take place and the seller cannot reneg on the price.[5]

Our main result is that the monopolist's payoff is bounded away from the low valuation irrespective of the discount factor. We prove this result by constructing a simple and direct but suboptimal mechanism which never reveals any information to the seller. In the initial period, the buyer reports his valuation and the high type trades at a price less than his valuation with probability less than one. The low-type buyer does not trade in the first period. Even though the seller receives no signal from the mechanism, she updates her prior about the buyer's type whenever trade does not occur. In every subsequent period, the probability of trade is constant and does not depend on the buyer's type, so the seller's posterior remains the same unless there is sale. Furthermore, the price is the buyer's valuation. This means that the low-value buyer's payoff is zero and the high-value buyer earns rent only in the initial period. Finally, from the second period onwards, the discounted present value of the seller's payoff is larger than the low valuation.

If the seller abandons this contract, she loses its information content. The trading probabilities in this mechanism are specified so that the optimal full-commitment mechanism in all but the initial period is clearing the market at the low valuation. Since the seller's expected payoff is larger than the low valuation, the constraint guaranteeing that she does not abandon the mechanism is satisfied in each future period. The mechanism we construct may not be optimal: In the initial period, the seller might prefer to choose a different mechanism. But that would only imply that her equilibrium payoff is even larger than the one generated by the mechanism described above. Therefore, since the seller's expected profit is bounded away from the low valuation in our mechanism, her equilibrium payoff is also larger than the low valuation. That is, the Coase Conjecture fails.

Literature Review

The literature on dynamic contracting in the absence of commitment probably started with the papers by Laffont and Tirole (1988 and 1990). The authors offer two related yet distinct approaches to model such environments. The first one is to consider one-period contracts. In each period, the principal offers a contract which, if accepted by the agent, determines the allocation in that period as a function of contractible variables.[6] The second approach is to allow dynamic

[5]McAdams and Schwarz (2007) and Akbarpour and Li (2020) consider static mechanism design problems where the principal has even less committment power and she cannot credibly promise to follow the rules of her own mechanism.

[6]Examples for recent papers analysing dynamic screening problems with short-term contracts include Gerardi and Maestri (2020), Beccuti and Möller (2018), Acharya and Ortner (2017) and Tirole (2016).

contracts which can be voided and replaced if both parties wish to do so. That is, equilibrium contracts must be *renegotiation-proof*.[7]

The methodological contribution of our paper is to put forward another approach of modeling limited commitment which appears to be new. In order to explore the consequences of the absence of commitment in the context of a principal-agent relationship, it is desirable to consider a setting which differs from the full-commitment benchmark only in the assumption regarding the principal's commitment power. In the full-commitment benchmark, the principal has access to dynamic contracts and has full bargaining power. Therefore, our model combines the two approaches of Laffont and Tirole (1988 and 1990) in the following way. On the one hand, the set of mechanisms is not restricted to be one-period ones and the principal has access to infinite-horizon dynamic contracts. On the other hand, the principal can offer new contracts in each period and the agent's consent is not required to abandon the previous contract.

Doval and Skreta (2020a) also consider mechanism design problems with limited commitment. They generalize the approach in Laffont and Tirole (1988) and consider one-period contracts. Their mechanisms do not only determine allocations but can also reveal public information. The authors develop a Revelation Principle and show that the information revealed by a mechanism can be assumed to be the principal's posterior about the agent's type.[8] In their companion paper, Doval and Skreta (2020b) show that the Coase Conjecture still holds with such a contract space. Indeed, the authors demonstrate that in a Coasian environment, the seller optimally posts prices in each period.[9]

Our paper also contributes to the literature documenting failures of the Coase Conjecture in the 'gap case'. With multiple atomic buyers, Bagnoli, Salant and Swierzbinski (1989), von der Fehr and Kuhn (1995) and Montez (2013) show that the seller can maintain high posted prices until a trade occurs. Feinberg and Skrzypacz (2005) show that higher-order uncertainty can generate delay. Other papers demonstrate that the Coase Conjecture is not robust to the assumption that the seller's marginal cost of production is constant, see, for example, Kahn (1986), McAfee and Wiseman (2008), Karp (1993), and Ortner (2017).[10] Bulow (1982) argues that the monopolist benefits from renting the good rather than selling it.[11]

[7] Among others, Battaglini (2007) and Maestri (2017) generalize the results of Laffont and Tirole (1990) in various ways. Strulovici (2017) provides a foundation for renegotiation-proof contracts in a bargaining environment. Hart and Tirole (1988) and Breig (2019) compare the two modelling approaches in a dynamic buyer-seller relationship.

[8] Bester and Strausz (2001) also develop a Revelation Principle in finite-horizon environments and finite type spaces.

[9] Lomys and Yamashita (2021) introduce a mediator into the model of Doval and Skreta (2020a) who controls the communication between the contracting parties. The mediator cannot be replaced by the principal and can possess information which the principal does not have. The authors demonstrate that such a mediator expands the set of implementable allocations of Doval and Skreta (2020a).

[10] Also related is the literature on obsolescence or imperfect durability of the good, see Bulow (1986), Waldman (1993), and Fudenberg and Tirole (1998).

[11] Hart and Tirole (1988) point out that the aruments of Bulow (1982) relies on buyer-anonymity and show and that renting may make the seller worse-off.

Another approach to break the Coase Conjecture is to allow the seller to intratemporally screen, e.g. by producing a variety, see Wang (1998), Takeyama (2002), Hahn (2006), Inderst (2008), or Board and Pycia (2014). A notable contribution by Nava and Schiraldi (2019) demonstrates that all these results are consistent with the Coasian logic in the following sense. The seller's limit payoff is the maximal static monopoly profit subject to the market-clearing condition.

The Coase conjecture has been also proved to fail when market deterioration is prevented by the arrival of new buyers or stochastically changing values. Important contributions include Sobel (1991), Biehl (2001), Fuchs and Skrzypacz (2010), and Garrett (2016).

Our work is also related to the literature on smart contracts. The term 'smart contract' was first coined by Nick Szabo in the mid-90's, whose prototypical example of a vending machine highlights the ideas of automatic execution and immutability. Since then, with the advent of bitcoin and the popularization of blockchain technologies such as Ethereum, interest in smart contracts has heightened. For some recent papers on the blockchain, see Huberman, Leshno and Moallemi (2021) who provide an insightful analysis of the Bitcoin Payment System and Abadi and Brunnermeier (2018) who study the impossibility of any distributed ledger to satisfy certain desiderata.

Recent research on smart contracts has explored how these contracts can enlarge the space of implementable economic outcomes. Cong and He (2019) study the effects of smart contracts on industrial organization, while Tinn (2018) studies how financial contracting may be affected. Bakos and Halaburda (2020) delineate the effects of enhanced information generation of technologies dubbed the Internet-of-Things, and the automatic execution offered by smart contracts in a simple contracting game. Finally, Holden and Malani (2018) examine the use of smart contracts in the context of the hold-up problem. Two key properties of smart contracts underpin all of the above papers: (i) enhanced commitment power—for example, through lowering enforcement costs via automatic execution, or preventing renegotiation of terms altogether; and (ii) better information—for example, by reducing state-verification costs.

While we recognize that restoring some commitment power is possibly the main reason for the popularity of smart contracts, our paper intends to provide a different perspective. We take the view of Laffont and Tirole (1988) that the lack of intertemporal commitment is a form of contractual incompleteness. In other words, contracting parties may refrain from signing long-term, binding contracts due to potential unforeseen or non-contractible contingencies even if such contracts were feasible.[12] Although we do not model these contingencies, our assumption that the seller cannot commit not to switch off a deployed contract embodies the idea that she prefers a contract allowing for discretion in the future.[13] Our main result suggests that smart contracts

[12]Unexpected software security vulnerabilities, bugs, novel types of attack threats, the need for upgrades, and regulatory risk, are among the reasons one may willingly retain discretion over aspects of a smart contract, preventing its absolute immutability in practice.

[13]Such control can be exercised via 'admin keys' retained by the issuer. It is worth noting that 12 out of the 15 most popular Decentralized Finance protocols, governed by smart contracts, have such 'admin keys'

may turn out to be useful even in such environments because they can store information securely.

2 The Model

There is a seller of a durable, indivisible good and a buyer whose willingness-to-pay for the good is his private information. The buyer's valuation is either high, v_h, or low, v_l so that $v_h > v_l > 0$. The probability of high valuation, μ, is common knowledge. We assume that $\mu v_h > v_l$, so the static monopoly price is v_h. Time is discrete and indexed by 0, 1,... In the initial period, the seller offers a contract from the set $\mathcal{C}$ described below. This contract then determines the allocation, i.e., the probability of trade and the transfer, in every period unless it is replaced. In each subsequent period, the seller decides whether to proceed with the previous period's contract or to deploy a new one. If the seller deploys a new contract, then it will determine the allocation in that period as well in every future period until it is replaced. The game ends when the good is sold. We assume that both parties discount the future according to the common factor $\delta\,(\in (0,1))$. If the buyer's valuation is $v\,(\in \{v_l, v_h\})$, trade occurs in period T and the transfer is p_t at time t, then the payoffs of the buyer and seller are

$$\delta^T v - \sum_{t=0}^{\infty} \delta^t p_t \text{ and } \sum_{t=0}^{\infty} \delta^t p_t,$$

respectively. Moreover, both parties maximize their expected payoffs.

The Contract Space $\mathcal{C}$.— We describe a typical contract, c, from the seller's contract space $\mathcal{C}$. The contract specifies both the communication and the implemented allocation in each period when the contract is deployed. Formally, $c = \left(M_T^b, M_T^s, S_T^b, S_T^s, \mathbf{x}_T, \mathbf{p}_T, \rho_T\right)_{T=0}^{\infty}$, where M_T^b and M_T^s are the messages available to the buyer and the seller in a given period if the contract was already deployed T consecutive periods immediately preceding that period.[14] The sets S_T^b and S_T^s are the set of signals the buyer and the seller may receive privately. The functions $\mathbf{x}_T : \left(M_\gamma^b, M_\gamma^s\right)_{\gamma=0}^{T} \times \left(S_\gamma^b, S_\gamma^s\right)_{\gamma=0}^{T-1} \to [0,1]$ and $\mathbf{p}_T : \left(M_\gamma^b, M_\gamma^s\right)_{\gamma=0}^{T} \times \left(S_\gamma^b, S_\gamma^s\right)_{\gamma=0}^{T-1} \to \mathbb{R}$ specify the probability of trade and the transfer conditional on sale[15] as a function of histories of messages and signals. Finally, the function $\rho_T = \left(\rho_T^b, \rho_T^s\right) : \left(M_\gamma^b, M_\gamma^s\right)_{\gamma=0}^{T} \times \left(S_\gamma^b, S_\gamma^s\right)_{\gamma=0}^{T-1} \to \Delta\left(S_T^b, S_T^s\right)$ specifies the distributions of the signals revealed to the buyer and the seller as a function of the history of reports. To model the buyer's participation decision, we assume that, for each T, the buyer's message space, M_T^b, includes a special message, r, which triggers no trade. Sending this message is interpreted as *rejecting the contract.* If the buyer rejects the contract, $m_T^b = r$, then

(https://cointelegraph.com/news/how-many-defi-projects-still-have-god-mode-admin-keys-more-than-you-think).

[14]For an example, suppose that c is deployed at $t = 0, 2, 3$ but not at $t = 1$. Then, $T = 0$ at $t = 0, 2$ and $T = 1$ at $t = 3$.

[15]For notational simplicity, we assume that transfers are deterministic and paid only if there is trade. Allowing random transfers has no impact on our results.

$\mathbf{x}_T = \mathbf{p}_T = 0$.[16] We say that the contract c is *actively deployed* in a given period, if the seller deploys c and the buyer does not reject it in that period. The seller's contract space $\mathcal{C}$ is a set of contracts described above.

Note first that the signals revealed to the contracting parties are assumed to be private. However, when the signals are perfectly correlated, they are effectively public. In fact, contracts are defined to be general enough to also allow the mixture of private and public communication; signals may have both private and public components. Second, despite the seller having no private information to start with, it is important to allow a contract to condition on the seller's messages. The reason is that the seller learns over time and when she decides to deploy a contract, she may benefit from inputting her posterior and making the implemented allocations dependent on it.

Simple and Direct Contracts.— We are not making any additional assumption on the contract space except that it contains all those mechanisms which ask the buyer to report her valuation in the initial period of deployment but involve no additional meaningful communication. We call such contracts *simple and direct* and define them formally below. Again, let us describe a typical simple and direct contract, d. First, if d is deployed repeatedly then sending the message r only triggers a one period of delay. So, the easiest way to describe d is to index the message and signal spaces as well as the allocations defining these contracts by the number of those consecutive periods of deployment in which the contract d was not rejected. More precisely, at each history, let τ denote the number of previous periods in which d was actively deployed since a different contract was deployed.[17] Then, with a slight abuse of notation, the contract d is defined by the collection $(\mathbf{x}_\tau, \mathbf{p}_\tau)_{\tau=0}^{\infty}$, where $\mathbf{x}_\tau : \{v_l, v_h\} \to [0,1]$ and $\mathbf{p}_\tau : \{v_l, v_h\} \to \mathbb{R}$. In the initial period of deployment, and in every other period in which $\tau = 0$, the buyer is asked to report his valuation, so $M_0^b = \{v_l, v_h, r\}$. If the buyer reports $v \in \{v_l, v_h\}$ then trade occurs with probability $\mathbf{x}_0(v)$ at price $\mathbf{p}_0(v)$. If the buyer sends the message r and the seller deploys d in the next period, the buyer's message space is again $\{v_l, v_h, r\}$ and the allocation is determined by $(\mathbf{x}_0, \mathbf{p}_0)$. After the buyer does not reject d and reports a valuation, he can only accept or reject the contract, that is, $M_\tau^b = \{a, r\}$ for all $\tau > 0$. The seller is only informed whether or not the buyer rejected the contract, that is, $S_\tau^s = \{a, r\}$ for all τ and $\rho_\tau^s\left(m_\tau^b\right) = r$ if, and only if, $m_\tau^b = r$. The seller does not communicate to the contract and the buyer does not receive any information, so the seller's message spaces and the buyer's signal spaces are singletons. The set of such simple and direct contracts is denoted by $\mathcal{D}$ and we assume that $\mathcal{D} \subset \mathcal{C}$.

We point out that the set $\mathcal{D}$ is different from the set of contracts one may wish to call *direct* in our environment. In general, a contract should be defined to be direct if its message spaces in each period of its deployment are rich enough to allow the seller and the buyer to report their private information. Since such a contract may send signals to both parties and information may

[16] One may also find it natural to assume that the seller is informed about the buyer's rejection of the contract. Our main result holds irrespective of such an assumption.

[17] For example, if d was deployed at $t = 0, 1, 2, 3$ and was rejected only at $t = 1$, then $\tau = 2$ at $t = 3$.

also evolve in those periods when the contract is not deployed, a direct contract must allow the reporting if hierachies of beliefs. For example, the seller's type includes his posterior about the buyer's valuation, her belief about the buyer's belief about her posterior etc.

Equilibrium Concept and Existence.— We focus on Weak Perfect Bayesian Equilibria. That is, an equilibrium is defined as an assessment: a pair of system of beliefs and a (possibly mixed) strategy profile. The belief system specifies for each information set of the game a probability distribution over the set of nodes in that set, which is then interpreted as the belief of the contracting party who moves at that information set. An assessment is a Weak Perfect Bayesian Equilibrium if (i) the strategy profile is sequentially rational at each information set and (ii) beliefs are derived by Bayes' rule at those information sets which are reached with positive probability.

The concept of Weak Perfect Bayesian Equilibrium places little restrictions on the players' out-of-equilibrium beliefs. Since we provide a lower bound on the seller's equilibrium payoffs, one may suspect that this result is supported by constructing beliefs off the equilibrium path which may appear unreasonable. For example, if the seller believes that the buyer's willingness-to-pay is surely v_h whenever he rejects a contract, she would rationally offer a contract which specifies trade only at price v_h in subsequent periods. In fact, the seller may maintain this belief even after the buyer rejects contracts arbitrarily many times. This, in turn, may deter the buyer to reject an otherwise unattractive contract in the first place if it generates non-negative payoffs. We emphasize that our analysis does not rely on such arguments. In fact, in the spirit of the concept of Sequential Equilibrium, special care is taken to construct the seller's beliefs so that they are limit points of beliefs derived by Bayes' rule along a sequence of totally mixed strategy profiles converging to the equilibrium strategy profile.

It is not hard to show that equilibria exist in a discretized version of our model, i.e., the set of contracts, the message and signal spaces are all finite[18]. We also prove existence for the case when the seller only has access to simple and direct contracts, that is, $\mathcal{C} = \mathcal{D}$, see the Online Appendix[19]. In the rest of this paper, we assume that an equilibrium exists irrespective of the discount factor.

Revelation Principle.— The concept of the Revelation Principle in our environment is slightly more subtle than in the case of full-commitment. In the latter case, standard revelation principles state that the contract space can be restricted to be a canonical class without the loss of generality. That is, irrespective of the contract space, the seller's payoff can never exceed that generated by the optimal canonical contract. As mentioned in the Introduction, when commitment power is limited, the seller may benefit from reducing her contract space. Consequently, whether or not one may restrict attention to a certain canonical class of contracts depends on the original space. Nevertheless, it is not hard to show that if the contract space is rich enough, a Revelation Principle similar to the standard one still holds. More precisely, if the contract space includes all direct

[18]See Fudenberg and Levine (1983).

[19]Available at: tinyurl.com/36tzxh3k.

contracts, each equilibrium outcome can be reproduced by an incentive compatible direct contract which the seller deploys in each period and the buyer never rejects it.

By no means does the aforementioned Revelation Principle imply that the incentive compatible direct contract can also be assumed to be simple. That is, the equilibrium direct contract may involve communication with the contracting parties in addition to eliciting the buyer's willingness-to-pay in the initial period. One may naively suspect that such additional communication can be dispensed with because the seller offers the same direct contract in each period irrespective of the signals she received. So, it appears that unifying all the information nodes in each period at which the seller offers the same contract would make no difference. Similarly, revealing information to the buyer would only make it harder to satisfy his incentive compatibility constraints. Unfortunately, this reasoning is incomplete and, in general, the equilibrium contract cannot be assumed to be simple and direct. The reason is that the seller may benefit from receiving signals because the change in her beliefs due to these signals may deter her from deviating to other contracts. This is because knowing that the seller received information about the buyer's willingness-to-pay, the buyer would reject certain off-equilibrium contracts in the future which, in turn, can enable the seller to stick with the equilibrium contract. Since the direct contract guaranteed by the Revelation Principle may involve complex communication and reporting hierarchies of beliefs, this concept does not seem to be operational in characterizing equilibrium outcomes. Our forthcoming analysis will not rely on this concept and hence, we do not present formal proofs for the statements above.

3 Main Result

In order to state our main theorem, let $\pi(\mathcal{C}, \delta)$ denote the supremum of the seller's payoff across all equilibria if the contract space is $\mathcal{C}$ and the discount factor is δ.

Theorem 1 *There exist $\overline{\delta} \in (0,1)$ and $\underline{\pi} > v_l$, such that for all $\delta \in \left(\overline{\delta}, 1\right)$,*

$$\pi(\mathcal{C}, \delta) \geq \underline{\pi}.$$

We remark that this theorem implies the failure of the Coase Conjecture: no matter how close the discount factor is to one, the largest equilibrium payoff of the seller is bounded from below by a constant, $\underline{\pi}$, which is larger than the low valuation, v_l.

The key to the arguments leading to the statement of Theorem 1 is to analyze a particular set of simple and direct contracts, coined as *abiding contracts*. The identifying feature of these contracts is that if they are actively deployed forever then (i) the buyer's continuation payoff is weakly positive in each period irrespective of his type and (ii) the seller's expected continuation payoff is larger than her full-commitment profit in all but the initial periods. The proof of the theorem consists of two steps. We first show that the seller's largest equilibrium payoff cannot be smaller than her payoff generated by any of the abiding contract. The second step is to construct an abiding mechanism which generates a payoff to the seller which is larger than v_l and does not depend on the discount factor.

Next, we define incentive compatible and abiding contracts formally.

Incentive Compatible Simple and Direct Contracts.— Whether the buyer has incentive to report his willingness-to-pay truthfully after accepting a simple and direct contract depends on what contracts he expects to be deployed in the future. Moreover, it also depends on the discount factor, δ. In what follows, we define incentive compatibility conditional on the same contract being deployed forever. Before presenting the formal definition, recall that the allocation determined by a simple and direct contract, $(\mathbf{x}_\tau, \mathbf{p}_\tau)_0^\infty$, depends only on the initial report of the buyer. Observe that if a simple and direct contract is actively deployed in each period, the buyer's report, v, determines the unconditional probability of trade, $X_\tau(v)$, and the expected transfer, $P_\tau(v)$, in each period by

$$X_\tau(v) = \mathbf{x}_\tau(v) \Pi_{t=0}^{\tau-1} \left(1 - \mathbf{x}(v)_t\right) \text{ and } P_\tau(v) = \mathbf{p}_\tau(v) \mathbf{x}_\tau(v) \Pi_{t=0}^{\tau-1} \left(1 - \mathbf{x}(v)_t\right).$$

Vice versa, each simple and direct contract $d \in \mathcal{D}$ can be described by $(X_\tau, P_\tau)_{\tau=0}^\infty$, where $X_\tau : \{v_l, v_h\} \to [0,1]$ denotes the probability that trade occurs in period τ conditional d being actively deployed in each period and $P_\tau : \{v_l, v_h\} \to \mathbb{R}$ is the expected transfer in that period.

Note that if a simple and direct contract is deployed forever, the buyer may maximize his payoff by misreporting his type in the initial period of deployment and optimizing with respect to his rejection-acceptance strategy in the future. Let $U(v, \widehat{v}, d, \delta)$ denote the buyer's value if the contract d is deployed forever, the discount factor is δ, the buyer's valuation is v and he reported $\widehat{v}$ in the initial period. Recall that if the buyer rejects a simple and direct contract, he only induces a one-period delay. Therefore, he only rejects the contract if his continuation payoff is negative, in which case, he would reject it forever. Consequently,

$$U(v, \widehat{v}, d, \delta) = \sup_{T \geq 0} \sum_{t=0}^{T} \delta^t \left[X_t(\widehat{v}) v - P_t(\widehat{v})\right],$$

where T denotes the time period after which the buyer rejects the contract forever. We are now ready to define incentive compatibility.

Definition 1 *The contract $d = (X_\tau, P_\tau)_{\tau=0}^\infty \in \mathcal{D}$ is δ-incentive compatible if for $v \in \{v_l, v_h\}$*

$$v \in \arg \max_{\widehat{v} \in \{v_l, v_h\}} U(v, \widehat{v}, d, \delta).$$

Abiding Contracts.— As mentioned before, we intend to call a contract abiding if, conditional on the contract being actively deployed forever, the buyer's continuation payoff is non-negative and the seller's continuation payoff exceeds her full-commitment profit in each period. More precisely, we require that an abiding contract specifies trading probabilities with each type of the buyer so that, after the initial period, the static monopoly price becomes the low valuation. That is, conditional on not trading, the seller becomes so pessimistic regarding the buyer's willingness-to-pay that she would optimally clear the market at price v_l. Before providing the formal definition, let us introduce an additional piece of notation. If an incentive compatible, simple and direct contract is actively deployed in each period then $\mu_t(d)$ denotes the posterior probability that the buyer's willingness-to-pay is v_h in period t.

Definition 2 *The contract $d = (X_\tau, P_\tau)_{\tau=0}^{\infty} \in \mathcal{D}$ is δ-abiding if it is δ-incentive compatible and, in addition,*

(i) $\sum_{t=T}^{\infty} \delta^{t-T} [X_t(v) v - P_t(v)] \geq 0$ for all $v \in \{v_l, v_h\}$, $T \geq 0$,

(ii) $\mu_t(d) \leq v_l/v_h$ for all $t \geq 1$, and

(iii) $\mu_T(d) \sum_{t=T}^{\infty} \delta^{t-T} P_t(v_h) + (1 - \mu_T(d)) \sum_{t=T}^{\infty} \delta^{t-T} P_t(v_l) \geq v_l$ for all $T \geq 1$.

Condition (i) implies that if a δ-abiding contract is deployed forever, accepting the contract in each period is an optimal strategy of the buyer if his discount factor is δ. Conditions (ii) and (iii) require that the static monopoly price is v_l and the seller's continuation value is larger than v_l in all but the initial period if d is actively deployed forever.

Let $v(d, \delta)$ denote the seller's payoff if the incentive compatible, simple and direct contract $d = (X_\tau, P_\tau)_{\tau=0}^{\infty} \in \mathcal{D}$ is actively deployed forever, that is,

$$v(d, \delta) = \mu \sum_{t=0}^{\infty} \delta^t P_t(v_h) + (1 - \mu) \sum_{t=0}^{\infty} \delta^t P_t(v_l).$$

We are ready to state that the seller's value generated by any abiding contract is a lower bound on her largest equilibrium payoff.

Lemma 1 *Suppose that $d \in \mathcal{D}$ is a δ-abiding contract. Then $\pi(\mathcal{C}, \delta) \geq v(d, \delta)$.*

Let us explain the main arguments leading to this result. If the statement was false, the seller's payoff in each equilibrium would be strictly less than $v(d, \delta)$. Therefore, to prove the lemma, it is enough to argue that each such equilibrium can be modified so that, in the new equilibrium, the contract d is actively deployed forever. On the modified equilibrium path, the seller always deploys d and the buyer always accepts it. Off the equilibrium path the new equilibrium assessment is constructed based on the original equilibrium. In particular, the seller's payoff from offering a contract different from d in the initial period is the same as from offering that contract in the original equilibrium. Since the seller's payoff from offering any contract in the initial period is smaller than $v(d, \delta)$ in the original equilibrium, such deviations are not profitable. In subsequent periods, when d was already deployed a number of times, the seller's continuation payoff from offering it exceeds the full-commitment profit because d is abiding. Of course, if the seller abandons d and loses its information content, her continuation payoff cannot exceed the full-commitment profit. So even in later periods, the seller has no incentive to deviate from offering d. If the buyer ever rejects the contract, the seller's posterior belief remains the same.[20] Given this belief and that the buyer is expected to accept d, the seller rationally offers this contract even after many periods of rejection. In turn, knowing this, the buyer best-responds by accepting the seller's offer because d is abiding so it provides him with a non-negative continuation payoff irrespective of his willingness-to-pay.

Proof. We prove this lemma by contradiction. Suppose that the seller's payoff in each equilibrium is strictly smaller than $v(d, \delta)$. In what follows, we fix such an equilibrium and, by modifying it, we construct a new equilibrium so that the contract d is actively deployed forever and, consequently, the seller's payoff is $v(d, \delta)$, yielding a contradiction.

Let us first define the new equilibrium assessment at those information sets which are reached by paths along which no contract was offered but d. The seller always offers d and the buyer never rejects it. Moreover, in the initial period, the buyer reports his type truthfully. So, the equilibrium path, and hence payoffs, are determined by the repeated active deployment of d. If the seller moves at such an information set, her belief is defined to be $\mu_\tau(d)$ if d was actively deployed τ times before reaching that information set, irrespective of the number of times the buyer rejected the contract. In other words, when the buyer rejects d along a path where no other contract was offered, the seller does not update her belief.

Next, we define the assessment at each information set which is reached by a path along which a contract $c \neq d$ is offered. Observe that if d is actively deployed τ times before the seller deviates for the first time, her posterior is $\mu_\tau(d)$. Next, we show that even in the original equilibrium assessment there are information sets at which the seller's posterior is exactly $\mu_\tau(d)$. We accomplish this by demonstrating the existence of a simple and direct contract $c(d, \tau) = (X_\tau, P_\tau)_{\tau=0}^{\infty}$, with the following properties. In each equilibrium,

(i) the buyer accepts $c(d, \tau)$ in the initial period,

(ii) the buyer truthfully reports his type in the initial period if $c(d, \tau)$ is deployed and

(iii) the seller's posterior belief is $\mu_\tau(d)$ after the initial period if there is no trade.

To this end, let $P_0(v_l) = -v_h$, $X_0(v_l) = 1/2$, $P_0(v_h) = -v_h + q(v_l + \varepsilon)$ and $X_0(v_h) = 1/2 + q$, so that

$$\mu_\tau(d) = \frac{\mu\left(\frac{1}{2} - q\right)}{\mu\left(\frac{1}{2} - q\right) + (1 - \mu)\frac{1}{2}}$$

and $\varepsilon\,(> 0)$ is small enough so that $v_h - (v_l + \varepsilon) > \delta(v_h - v_l)$. Moreover, let $P_\tau(v) = X_\tau(v) = 0$ for all $\tau > 0$ and $v \in \{v_l, v_h\}$. One interpretation of this contract is that each type trades with probability half at price $-2v_h$ in the initial period. If the buyer reports v_h, he trades with an additional probability of q at a price just above v_l. After the initial period, the contract prescribes autarky. Note that accepting this contract generates an instantaneous payoff of at least v_h to the buyer. The sequential rationality of the seller implies that the expected continuation payoff of

[20] This belief is the limit of beliefs derived by Bayes' rule along a sequence of mixed strategy of the buyer over rejecting and accepting the contract, along which the probability of rejection goes to zero and does not depend on the buyer's valuation.

the buyer cannot exceed v_h, so the buyer accepts this contract in every equilibrium, yielding (i). To see part (ii), first recall that reporting v_h triggers trade with an additional probability of q at $v_l + \varepsilon$. Observe that the sequential rationality of the seller implies that the object is never sold at a price lower than v_l so the high-value buyer is better off trading at $v_l + \varepsilon$ with an additional probability of q whereas the low-value buyer is not. Hence, the buyer reports his value truthfully. To obtain part (iii), observe that q is defined so that the seller's posterior is computed by Bayes' rule and is $\mu_\tau(d)$.

We are now ready to specify the new assessment at those information sets which are reached by paths along which a contract $c \neq d$ is offered. To this end, consider an information set at which c is offered and along the paths reaching this set the contract d was actively deployed τ times and no other contract was ever offered. Therefore, the seller's posterior belief when offering c is $\mu_\tau(d)$. Of course, the continuation game starting at this information set is isomorphic to the continuation game in which the seller offers $c(d,\tau)$ in the initial period, the buyer accepts it and reports his type truthfully, trade does not occur and the seller offers c in the next period. Indeed, the seller's posterior is also $\mu_\tau(d)$ by the definition of $c(d,\tau)$. Therefore, we define the equilibrium assessment in the continuation game starting from offering c to be the same as the original equilibrium assessment in the continuation game in which the seller offers $c(d,\tau)$ and c in the first and second periods, respectively.

It remains to prove that the new assessment defined above is indeed an equilibrium assessment. We first argue that players are sequentially rational at each information set. In the initial period, the seller's payoff from offering $c(\neq d)$ is at most as large as her payoff in the original equilibrium. Since, $v(d,\delta)$ is larger than that, the seller rationally offers d. At those information sets which are reached by paths along which only d was offered, the seller's continuation payoff is larger than her full-commitment payoff given her posterior. So, even at those information sets, the seller rationally offers d. At any other information set, the seller's strategy is sequentially rational because it is defined by the sequentially rational original equilibrium assessment in the corresponding isomorphic continuation game. The buyer's strategy is also sequentially rational at those information sets which are reached by those paths along which no contract other than d was offered. The reason is that d provides the buyer with a non-negative payoff and rejecting d would only delay those payoffs given that the seller offers it again after any number of rejections. Since d is incentive compatible, the buyer rationally reports his type truthfully in the first period. At any other information set, the buyer's strategy is sequentially rational because it is defined by the original equilibrium assessment in the corresponding isomorphic continuation game. Also note that the seller's belief is defined by Bayes' rule at each information set which is reached with positive probability. Indeed, the seller's belief after the contract d was actively deployed τ times is $\mu_\tau(d)$. ■

Having established Lemma 1, in order to prove Theorem 1, we need to demonstrate the existence of a δ-abiding contract for each large enough δ which generates a payoff to the seller which is larger than v_l and does not depend on the discount factor. The next lemma states that such contracts exist.

Lemma 2 *There exists a $\overline{\delta} \in (0,1)$ such that for all $\delta \in \left(\overline{\delta}, 1\right)$ there is a δ-abiding contract $d_\delta \in \mathcal{D}$ so that $v\left(d_\delta, \delta\right) = \underline{\pi} > v_l$.*

In what follows, we construct a contract satisfying this lemma's statement. This contract will have the following properties. In its initial period of deployment, if the buyer reports v_h, the contract specifies a positive probability of trade, α, at price $p \in [v_l, v_h]$ and the buyer does not trade if he reports v_l. In any subsequent period, trade occurs with probability β with both types at a price equal to the buyer's report. In other words, this simple and direct contract depends on three parameters $(\alpha, \beta, p) \in [0,1]^2 \times [v_l, v_h]$ and can be formally defined as follows. In the initial period, $x_0(v_h) = \alpha$, $x_0(v_l) = 0$ and $p_0(v_h) = p_0(v_l) = p$.[21] For each $\tau > 0$, $x_\tau(v) = \beta$ and $p_\tau\left(v\right) = v$ for $v \in \{v_l, v_h\}$. In what follows, we express the constraints guaranteeing that the contract is not only incentive compatible but also abiding in terms of these parameters.

Let us first discuss incentive compatibility. First, note that the buyer with type v_l weakly prefers to report his willingness-to-pay irrespective of the parameter values. The reason is that if he does so, he always trades at price v_l and hence, his expected payoff is zero. On the other hand, if he reports v_h, the price is always weakly larger than v_l and hence, his expected payoff is non-positive. Consider now the buyer whose valuation is v_h. If he reports his true valuation, his payoff is $\alpha\left(v_h - p\right)$ because he trades with probability α at price p in the initial period and, any time in the future, the price is v_h. If he reports v_l, he does not trade in the initial period and, conditional on not trading before, he trades with probability β at price v_l in every period in the future. Therefore, if the high-type buyer misreports his type, the expected discounted present value of his payoff is

$$\delta \sum_{t=0}^{\infty} \delta^t (1-\beta)^t \beta (v_h - v_l) = \frac{\beta\delta}{1-\delta+\beta\delta}(v_h - v_l).$$

So, the incentive constraint of the buyer with type v_h is satisfied if

$$\alpha(v_h - p) \geq \frac{\beta\delta}{1-\delta+\beta\delta}(v_h - v_l). \tag{1}$$

Next, we investigate the set of those parameters for which the contract is abiding. First, note that whenever the buyer trades, the price is weakly smaller than his willingness-to pay. Therefore, if such a contract is actively deployed forever, the continuation value of each type is weakly larger than zero, so the contract satisfies part (i) of Definition 2 for any parameters. Let us now describe the constraint corresponding to part (ii) of Definition 2. That is, we describe conditions under which the seller's posterior in each future period is such that the full-commitment monopoly price is v_l and her continuation payoff exceeds v_l if the contract is deployed forever. To this end,

[21] Since $x_0\left(v_l\right) = 0$, $p_0\left(v_l\right)$ can be defined arbitrarily.

observe that, conditional on no trade, the seller's posterior remains the same after the initial period because probability of trade in future periods, β, does not depend on the buyer's report. This posterior depends only on the probability of trade in the initial period, α, and we denote it by $\widetilde{\mu}(\alpha)$. It can be computed by Bayes' Rule,

$$\widetilde{\mu}(\alpha) = \frac{(1-\alpha)\mu}{1-\mu+(1-\alpha)\mu}. \tag{2}$$

So, condition (ii) of Definition 2 holds and the static monopoly price is v_l if, and only if,

$$v_l \geq \widetilde{\mu}(\alpha)\, v_h. \tag{3}$$

We now compute the seller's continuation payoff in each period after the first deployment of the contract. Since neither the posterior distribution of types nor the probability of trade depends on time, this continuation payoff is also independent of time and can be expressed as

$$\sum_{t=T}^{\infty} \delta^{t-T}(1-\beta)^{t-T}\beta[\widetilde{\mu}(\alpha)\, v_h + (1-\widetilde{\mu}(\alpha))v_l] = \frac{\beta}{1-\delta+\beta\delta}[\widetilde{\mu}(\alpha)\, v_h + (1-\widetilde{\mu}(\alpha))v_l].$$

So, part (iii) of Definition 2 holds if this payoff is larger than v_l. In fact, we will construct parameter values so that that the seller's continuation value also exceeds his payoff from deploying this contract for one more period and selling the good at v_l immediately if the contract does not recommend trade, that is,

$$\frac{\beta}{1-\delta+\beta\delta}[\widetilde{\mu}(\alpha)\, v_h + (1-\mu(\alpha))v_l] \geq \beta[\widetilde{\mu}(\alpha)\, v_h + (1-\widetilde{\mu}(\alpha))v_l] + (1-\beta)\, v_l. \tag{4}$$

In order to prove Lemma 2, for each large enough δ, it is enough to show the existence of a triple, $(\alpha^*, \beta^*, p^*) \in [0,1]^2 \times [v_l, v_h]$, such that the constraints (1), (3) and (4) are satisfied. Furthermore, we need to demonstrate that the seller's payoff is larger than v_l and does not depend on δ.

Proof of Lemma 2. Let us explain how we construct the aforementioned triple of parameters. First, for each α we define $\widetilde{\beta}(\alpha)$ so that the abiding constraint (4) evaluated at $\beta = \widetilde{\beta}(\alpha)$ binds and ignore the constraint that $\widetilde{\beta}(\alpha)$ must be a probability. Second, we define $\widetilde{p}(\alpha)$ so that the incentive constraint (1) evaluated at $(\beta, p) = \left(\widetilde{\beta}(\alpha), \widetilde{p}(\alpha)\right)$ binds and ignore the constraint that $\widetilde{p}(\alpha) \in [v_l, v_h]$. Then, we consider the functional form of the seller's payoff at $\left(\alpha, \widetilde{\beta}(\alpha), \widetilde{p}(\alpha)\right)$, maximize it with respect to α subject to the constraint (3) and define α^* to be the maximizer. Finally, we show that, if δ is large enough, the parameters $\widetilde{\beta}(\alpha^*)$, $\widetilde{p}(\alpha^*)$ are feasible, that is, $\left(\widetilde{\beta}(\alpha^*), \widetilde{p}(\alpha^*)\right) \in [0,1] \times [v_l, v_h]$. Moreover, the seller's payoff generated by the contract corresponding to $\left(\alpha^*, \widetilde{\beta}(\alpha^*), \widetilde{p}(\alpha^*)\right)$ is strictly larger than v_l and does not depend on δ.

For each $\alpha \in [0,1]$, let $\widetilde{\beta}(\alpha)$ be defined so that the constraint (4) binds, that is,

$$\widetilde{\beta}(\alpha) = \beta = \frac{1-\delta}{\delta} \cdot \frac{v_l}{\widetilde{\mu}(\alpha)(v_h - v_l)}. \tag{5}$$

In addition, let us define $\widetilde{p}(\alpha)$ for each $\alpha \in [0,1]$ so that the high-type buyer's incentive constraint, (1) binds, that is

$$\alpha(v_h - \widetilde{p}(\alpha)) = \frac{\delta}{1-\delta+\widetilde{\beta}(\alpha)\,\delta}(v_h - v_l). \tag{6}$$

We now turn our attention to the seller's payoff generated by the contract corresponding to the triple $\left(\alpha, \widetilde{\beta}(\alpha), \widetilde{p}(\alpha)\right)$. We first compute the seller's continuation payoff in each period $t > 0$. As mentioned above, this continuation payoff does not depend on t. Since the abiding constraint (4) binds at $\beta = \widetilde{\beta}(\alpha)$, this payoff can be computed by plugging $\widetilde{\beta}(\alpha)$ into the right-hand side of this constraint,

$$\begin{aligned} & \frac{1-\delta}{\delta} \cdot \frac{v_l}{\widetilde{\mu}(\alpha)(v_h - v_l)}[\widetilde{\mu}(\alpha)\,v_h + (1-\widetilde{\mu}(\alpha))v_l] + \left(1 - \frac{1-\delta}{\delta} \cdot \frac{v_l}{\widetilde{\mu}(\alpha)(v_h - v_l)}\right) v_l \\ = \; & \frac{1-\delta}{\delta} \cdot \frac{v_l v_h}{(v_h - v_l)} - \frac{1-\delta}{\delta} \cdot \frac{v_l^2}{(v_h - v_l)} + v_l = \frac{v_l}{\delta}. \end{aligned}$$

We are now ready to compute the seller's payoff generated by the contract defined by $\left(\alpha, \widetilde{\beta}(\alpha), \widetilde{p}(\alpha)\right)$. Let $\nu(\alpha)$ denote this payoff. Observe that, in the initial period, the seller receives $\widetilde{p}(\alpha)$ with probability $\mu\alpha$ and, in the next period, her continuation payoff is v_l/δ. Therefore,

$$\nu(\alpha) = \mu\alpha\widetilde{p}(\alpha) + \delta(1-\mu\alpha)\frac{v_l}{\delta} = \mu\alpha\widetilde{p}(\alpha) + (1-\mu\alpha)\,v_l. \tag{7}$$

Substituting $\widetilde{p}(\alpha)$ from equation (6) and using equation (2) yield

$$\nu(\alpha) = v_l + (v_h - v_l)\left(1 - \frac{1-\mu}{1-\widetilde{\mu}(\alpha)} - \frac{\mu v_l}{\widetilde{\mu}(\alpha)\,v_h + (1-\widetilde{\mu}(\alpha))v_l}\right). \tag{8}$$

Finally, we define α^* to maximize $\nu(\alpha)$, subject to the constraint (3). That is, α^* solves

$$\max\left\{\nu(\alpha) : \alpha \in [0,1],\, \widetilde{\mu}(\alpha) \leq v_l/v_h\right\}. \tag{9}$$

We now show that α^* is uniquely determined. To this end, note that ν depends on α only through $\widetilde{\mu}(\alpha)$. Also note that, by (2), the function $\widetilde{\mu}$ is continuous, strictly decreasing in α and, in addition, $\widetilde{\mu}(0) = \mu$ and $\widetilde{\mu}(1) = 0$. Let $\Pi(\hat{\mu})$ denote $\nu\left(\widetilde{\mu}^{-1}(\hat{\mu})\right)$. In what follows, we characterize the unique solution, $\hat{\mu}^*$, of the following maximization problem

$$\max_{\hat{\mu} \in [0, v_l/v_h]} \Pi(\hat{\mu}).$$

Then, it follows that $\alpha^* = \widetilde{\mu}^{-1}(\hat{\mu}^*)$ is the unique solution of the problem (9). Note that

$$\Pi'(\hat{\mu}) = -(v_h - v_l)\left(\frac{1-\mu}{(1-\hat{\mu})^2} - \frac{\mu v_l (v_h - v_l)}{(\hat{\mu} v_h + (1-\hat{\mu})v_l)^2}\right),$$

so $\Pi'(\hat{\mu}) \geq 0$ if, and only if, $\hat{\mu} \leq \left[\sqrt{\frac{\mu}{1-\mu}} - \sqrt{\frac{v_l}{v_h - v_l}}\right] / \left[\sqrt{\frac{\mu}{1-\mu}} + \sqrt{\frac{v_h - v_l}{v_l}}\right]$. Therefore,

$$\hat{\mu}^* = \min\left\{\frac{\sqrt{\frac{\mu}{1-\mu}} - \sqrt{\frac{v_l}{v_h - v_l}}}{\sqrt{\frac{\mu}{1-\mu}} + \sqrt{\frac{v_h - v_l}{v_l}}}, \frac{v_l}{v_h}\right\} \tag{10}$$

and note that $\hat{\mu}^* \in (0, v_l/v_h]$ because $\mu \in (v_l/v_h, 1)$.

Let us now return to examine whether $\left(\widetilde{\beta}(\alpha^*), \widetilde{p}(\alpha^*)\right) \in [0,1] \times [v_l, v_h]$ if $\alpha^* = \widetilde{\mu}^{-1}(\hat{\mu}^*)$. Observe that, by equation (5), $\widetilde{\beta}(\alpha^*) \in [0,1]$ if, and only if,

$$\hat{\mu}^* \left(= \widetilde{\mu}(\alpha^*)\right) \geq \frac{1-\delta}{\delta} \cdot \frac{v_l}{v_h - v_l}.$$

Since the right-hand side is decreasing in δ and converges to zero as δ goes to one, there exists $\overline{\delta}$ such that $\widetilde{\beta}(\alpha^*) \in [0,1]$ whenever $\delta \in \left(\overline{\delta}, 1\right)$. Let us turn our attention to the first period's transfer, $\widetilde{p}(\alpha^*)$. By the definition of the function $\widetilde{p}$, it follows that $\widetilde{p}(\alpha^*) \leq v_h$ for all $\alpha \in [0,1]$. Furthermore, equation (7) implies that the seller's payoff, $\nu(\alpha)$, can be expressed as a convex combination of $\widetilde{p}(\alpha^*)$ and v_l. Therefore, in order to establish that $\widetilde{p}(\alpha^*) \in [v_l, v_h]$ we only need to show that $\nu(\alpha^*) > v_l$, what we will do next.

Before proceeding, we note that the construction of the parameters depends on the prior distribution of types, $\mu \in (\nu_l/v_h, 1)$. We now make this dependency explicit and express the seller's payoff induced by the contact constructed above as a function of μ. To this end, let us write the seller's posterior defined by (10) as a function of μ, $\hat{\mu}^*(\mu)$. Now, observe that, by equation (8), the seller's payoff can be written as

$$V(\mu) = v_l + (v_h - v_l)\left(1 - \frac{1-\mu}{1-\hat{\mu}^*(\mu)} - \frac{\mu v_l}{\hat{\mu}^*(\mu) v_h + (1-\hat{\mu}^*(\mu)) v_l}\right). \tag{11}$$

In order to prove that $\nu(\alpha^*) > v_l$, it is enough to show that V is strictly increasing on $(\nu_l/v_h, 1)$ and $\lim_{\mu \to v_l/v_h} V(\mu) = v_l$. To this end, note that V is continuous on $(\nu_l/v_h, 1)$. From equation (10), it follows that there is a cutoff value of μ, $\bar{\mu} \in (v_l/v_h, 1)$[22], such that $\hat{\mu}^*(\mu) = v_l/v_h$ whenever $\mu \in (\bar{\mu}, 1)$. On this domain, $V(\mu) = \mu v_h^2/(2v_h - v_l)$, which is indeed strictly increasing. Since $\hat{\mu}^*$ was chosen to maximize the seller's payoff, on the domain $(v_l/v_h, \bar{\mu})$, the Envelope Theorem implies that

$$\begin{aligned} V'(\mu) &= (v_h - v_l)\left[\frac{1}{1-\hat{\mu}^*(\mu)} - \frac{v_l}{\hat{\mu}^*(\mu) v_h + (1-\hat{\mu}^*(\mu)) v_l}\right] \\ &= (v_h - v_l)\left[1 - 2\frac{v_l}{v_h} + \sqrt{\frac{v_l}{v_h}\left(1 - \frac{v_l}{v_h}\right)}\left(\sqrt{\frac{\mu}{1-\mu}} - \sqrt{\frac{1-\mu}{\mu}}\right)\right]. \end{aligned}$$

It is clear from inspecting the expression in the second line that V' is strictly increasing on $(v_l/v_h, \bar{\mu})$. Furthermore, since $\lim_{\mu \to v_l/v_h} \hat{\mu}^*(\mu) = 0$ by (10), the first line of the previous equality chain implies that $\lim_{\mu \to v_l/v_h} V'(\mu) = 0$. Therefore, V' is strictly positive on $(v_l/v_h, \bar{\mu})$. Recall that V' is also strictly positive on $(\bar{\mu}, 1)$ and continuous on $(v_l/v_h, 1)$. Then, by noting that $\lim_{\mu \to v_l/v_h} V(\mu) = v_l$, we conclude that $V > v_l$ on $(v_l/v_h, 1)$.

[22] It can be shown that

$$\bar{\mu} = \frac{v_l(2v_h - v_l)^2}{v_l(2v_h - v_l)^2 + (v_h - v_l)^3}.$$

To summarize, we have constructed a triple of parameters, $(\alpha^*, \beta^*, p^*) = \left(\alpha^*, \widetilde{\beta}(\alpha^*), \widetilde{p}(\alpha^*)\right)$. We have demonstrated the existence of $\overline{\delta}$ such that $\left(\widetilde{\beta}(\alpha^*), \widetilde{p}(\alpha^*)\right) \in [0,1] \times [v_l, v_h]$, so these parameters indeed define a contract. By equations (5) and (7), this contract is incentive compatible and abiding. Finally, we have proved that the seller's value from deploying this contract forever is strictly larger than v_l. To conclude the lemma's statement, all is left to do is to argue that, provided that $\delta > \overline{\delta}$, the seller's value does not depend on δ. This, however, is evident from equations (11) and (10). ■

We are ready to argue that the statement of Theorem 1 follows from Lemmas 1 and 2.

Proof of Theorem 1. Recall that Lemma 2 guarantees the existence of a threshold value of the discount factor, $\overline{\delta} \in (0,1)$, as well as the existence of a δ-abiding contract $d_\delta \in \mathcal{D}$ for each $\delta \geq \overline{\delta}$ such that seller's value generated by d_δ is larger than v_l and does not depend on δ, that is, $v(d_\delta, \delta) = \underline{\pi} > v_l$. Then Lemma 1 implies that, whenever $\delta \geq \overline{\delta}$, the seller's largest equilibrium payoff exceeds $\underline{\pi}$, that is, $\pi(\mathcal{C}, \delta) \geq \underline{\pi}$. ■

We conclude this section by remarking that the statement of Theorem 1 holds for all discount factors and the requirement that the contracting parties are patient enough, $\delta \geq \overline{\delta}$, can be disposed of. To see this, for each δ, consider the contract which specifies trade with the high-type buyer in the initial period at a price $v_h - \delta(v_h - v_l)$. Note that this price makes the high-type buyer indifferent between buying the good and trading at v_l a period later. Furthermore, the contract specifies trade with the low-type buyer in the next period at price v_l. This contract is obviously δ-incentive compatible and δ-abiding. Let π_δ denote the seller's value generated by this contract and note that

$$\pi_\delta = \mu[v_h - \delta(v_h - v_l)] + (1-\mu)\delta v_l = \mu v_h (1-\delta) + \delta v_l > v_l(1-\delta) + \delta v_l = v_l,$$

where the inequality follows from $\mu > v_l/v_h$. Moreover, note that π_δ decreases in δ. Therefore, Lemma 1 implies that for all $\delta \in (0, \overline{\delta}]$, $\pi(\mathcal{C}, \delta) \geq \pi_{\overline{\delta}}$. This observation and Theorem 1 imply that, for each $\delta \in (0,1)$, the seller largest equilibrium payoff, $\pi(\mathcal{C}, \delta)$, exceeds $\min\left\{\pi_{\overline{\delta}}, \underline{\pi}\right\} > v_l$.

4 Discussion

Optimal Contracts.— Theorem 1 states that the seller's largest equilibrium payoff is bounded away from v_l but it provides no further information about this payoff. In fact, we do not know what the seller's optimal contract is generating her largest equilibrium profit. However, when we prove equilibrium existence for the case of $\mathcal{C} = \mathcal{D}$, we construct an equilibrium contract which induces a payoff to the seller which is significantly larger than the bound provided by Lemma 2 (see the Online Appendix[23]). For each δ, this equilibrium contract specifies trade before a certain date, $T(\delta)$, with probability one. In the initial period, only the buyer with valuation v_h trades with a positive probability at a price in (v_l, v_h). Ever after, the price is always the reported valuation,

[23]Available at: tinyurl.com/36tzxh3k.

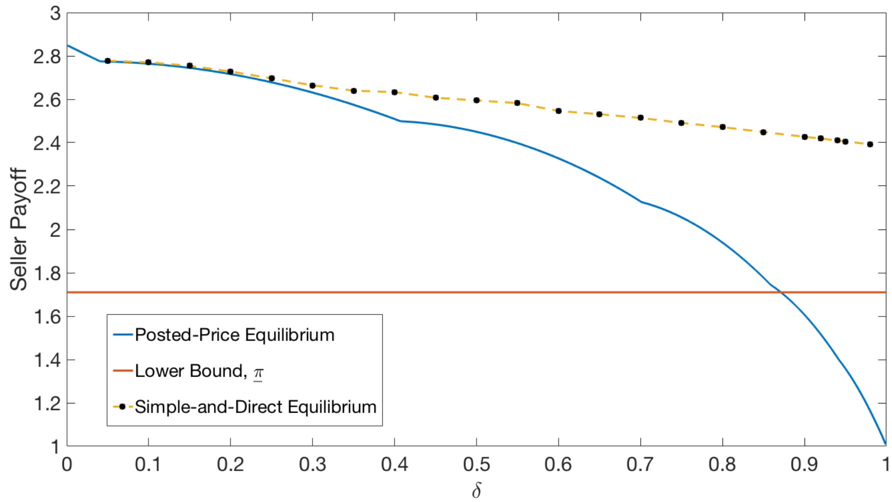

Figure 1: Comparison of the seller's profits when the discount factor varies.

just like in the case of the contract described in the proof of Lemma 2. In early time periods, only the high-type buyer trades and the seller is becoming more and more pessimistic. The low-type buyer only trades in the last and in the penultimate periods. Of course, as δ goes to one, $T(\delta)$ converges to infinity. Figure 1 plots the seller's payoff generated by this contract as a function of δ for the example where $v_l = 1$, $v_h = 3$ and $\mu = .95$.

What happens if the seller's contract space is larger than $\mathcal{D}$? As mentioned above, it is not hard to show that in different principal-agent models, the principal benefits from having access to contracts which reveal information to both contracting parties. It is also possible to construct examples where the principal is worse off if her contract space includes such contracts. Unfortunately, we were unable to establish whether the seller benefits or is hurt by enlarging the set $\mathcal{D}$.

One-period Contracts.— As mentioned in the Introduction, a common approach to model the lack of intertemporal commitment is to restrict the contract space to be the set of one-period contracts. Doval and Skreta (2020b) pursue this approach in the context of Coasian bargaining. In their setup, a contract of a given period determines the probability of trade and transfer in that period and reveals a public signal which can be assumed to be the seller's posterior. The authors show that the largest equilibrium payoff the seller can get, can be generated by a sequence of posted prices. Moreover, the seller's equilibrium profit converges to v_l as the discount factor goes to one, that is, the Coase Conjecture holds. Our main theorem highlights that the Coase Conjecture in Doval and Skreta (2020b) is not only the consequence of the seller's limited commitment power, but also of the restricted contract space.

Figure 1 also plots the seller's largest equilibrium profit in the model of Doval and Skreta (2020b) as a function of the discount factor. Note that when the discount factor is small, this profit level is larger than the one generated by the stationary contract described in Lemma 2. However, the profit induced by the aforementioned equilibrium contract is larger than the maximum profit in Doval and Skreta (2020b) irrespective of δ.

Side-Contracts.— If the seller decides not to proceed with the previous period's contract, the information content of the contract is lost. This feature enables the seller to redeploy contracts which implement allocations which are dominated from the viewpoints of both contracting parties. For example, in the context of the simple and direct contract of Lemma 2, the buyer trades with probability $\beta\,(< 0)$ at the price of the reported valuation in all but the initial periods. Of course, both the seller and the buyer would be (weakly) better off if this probability was larger. However, if the seller wants to replace the contract with another one with larger trading probabilities, she would need to pay information rent to the buyer again which makes such a deviation non-profitable. A possible way to circumvent such ex-post inefficiencies associated to a contract would be to consider the possibility of writing side-contracts. That is, the seller continues to redeploy the previous period's contract but can offer a side-contract which conditions on the outcome of the redeployed contract. Despite the fact that the seller cannot offer such contracts in our model, we next argue that our main result is robust to the introduction of side-contracts.

To this end, consider again the contract of Lemma 2 and note that the only information this contract ever reveals is whether there is trade. Of course, if the contract specifies trade, the game ends. So, a side-contract must condition on the event of no trade. Recall that the trading probability in the initial period is specified so that the static monopoly price is v_l in any subsequent period. Therefore, the seller's optimal side-contract would implement trade at price v_l if the outcome of the original contract is no trade. Note that the seller's payoff from offering such a side-contract is the right-hand side of equation (4). In other words, the abiding constraint (4) guarantees that the seller's continuation value from deploying the original contract exceeds her payoff from such a side-contract.

Implementation.— As explained in the Introduction, the contracts considered in our model have features resembling those of smart contracts used in digital markets. However, we emphasize that communications on blockchain-based software platforms, such as Ethereum, are typically public. For our main result to hold, it is essential that the buyer communicates his willingness-to-pay to the contract privately. Implementing such private communication is not hard using the cryptographic technology already employed in those markets. One way to do this is to encrypt the buyer's messages and let the buyer retain the decryption key. Then, in each period, the buyer can input the decryption key and an allocation is determined. Inputting an incorrect key would simply be treated as rejecting the contract.

References

[1] Abadi, Joseph, and Markus K. Brunnermeier. *Blockchain Economics.* Working Paper, 2019.

[2] Acharya, Avidit, and Ortner, Juan. *Progressive Learning.* Econometrica 85 (6), 1965-1990, 2017.

[3] Akbarpour, Mohammad, and Li, Shengwu. *Credible Auctions: A Trilemma.* Econometrica 88 (2), 425–467, 2020.

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

[5] Bagnoli, Mark, Salant, Stephen, and Swierzbinski, Joseph E. *Durable-Goods Monopoly with Discrete Demand.* Journal of Political Economy 97(6), 1459–1478, 1989.

[6] Bakos, Yannis and Halaburda, Hanna. *Smart Contracts, IoT Sensors and Efficiency: Automated Execution vs. Better Information.* Working Paper, Available at SSRN: https://ssrn.com/abstract=3394546, 2020.

[7] Battaglini, Marco. *Optimality and Renegotiation in Dynamic Contracting.* Games and Economic Behavior 60(2), 213–246, 2007.

[8] Battigalli, Pierpaolo, and Martin Dufwenberg *Dynamic Psychological Games.* Journal of Economic Theory 144 (1), 1-35, 2009.

[9] Beccuti, Juan, and Möller, Marc. *Dynamic adverse selection with a patient seller.* Journal of Economic Theory 173 (C), 95–117, 2018.

[10] Bester, Helmut, and Strausz, Roland. *Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case.* Econometrica: 69(4), 1077-1098, 2004.

[11] Biehl, Andrew R.. *Durable-Goods Monopoly with Stochastic Values.* The RAND Journal of Economics: 32 (3), 565–577, 2001.

[12] Blackwell, David. *Discounted Dynamic Programming.* Annals of Mathematics and Statistics 36 (1), 226-235, 1965.

[13] Board, Simon, and Pycia, Marek. *Outside Options and the Failure of the Coase Conjecture.* American Economic Review: 104(2), 656–671, 2014.

[14] Breig, Zachary. *Endogenous and Exogenous Commitment.* Economics Letters: 183 (C), 1–1, 2019.

[15] Bulow, Jeremy I. *Durable-Goods Monopolists.* Journal of Political Economy 90 (2): 314–332, 1982.

[16] Bulow, Jeremy I. *An Economic Theory of Planned Obsolescence*. The Quarterly Journal of Economics 101 (4): 729—749, 1986.

[17] Coase, Ronald H. *Durability and monopoly.* Journal of Law and Economics 15: 143–149, 1972.

[18] Cong, Lin William, and He, Zhiguo. *Blockchain Disruption and Smart Contracts.* The Review of Financial Studies 32 (5), 1754–1797, 2019.

[19] Doval, Laura, and Skreta, Vasiliki. *Mechanism Design with Limited Commitment.* Working Paper, Available at SSRN: https://ssrn.com/abstract=3281132, 2020a.

[20] Doval, Laura and Skreta, Vasiliki. *Optimal Mechanism for the Sale of a Durable Good.* Working Paper, Available at Arxiv: https://arxiv.org/abs/1904.07456, 2020b.

[21] Feinberg,Yossi, and Skrzypacz, Andrzej. *Uncertainty about Uncertainty and Delay in Bargaining.* Econometrica 73 (1): 69–91, 2005.

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

[23] Fudenberg, Drew and Levine, David. *Subgame-Perfect Equilibria of Finite- and Infinite-Horizon Games.* Journal of Economic Theory 31 (2): 251-268, 1993.

[24] Fudenberg, Drew and Tirole, Jean. *Upgrades, Tradeins, and Buybacks.* The RAND Journal of Economics 29 (2): 235–258, 1998.

[25] Fudenberg, Drew, Levine, David, and Tirole, Jean. *Infinite Horizon Models of Bargaining with One-sided Incomplete Information*, in “Bargaining with Incomplete Information” (Alvin Roth, Ed.), London/New York, Cambridge Univ. Press: 73–98, 1985.

[26] Garrett, Daniel F. *Intertemporal Price Discrimination: Dynamic Arrivals and Changing Values.* American Economic Review: 106 (11), 3275-3299, 2016.

[27] Gerardi, Dino, and Maestri, Lucas. *Dynamic Contracting with Limited Commitment and the Ratchet Effect.* Theoretical Economics 15 (2): 583-623, 2020.

[28] Gul, Faruk, and Sonnenschein, Hugo, and Wilson, Robert. *Foundations of Dynamic Monopoly and the Coase Conjecture.* Journal of Economic Theory 39: 155–190, 1986.

[29] Hahn, Jong-Hee. *Damaged Durable Goods.* The RAND Journal of Economics 37 (1): 121–133, 2006.

[30] Hart, Oliver D., and Jean Tirole. *Contract Renegotiation and Coasian Dynamics.* The Review of Economic Studies 55 (4): 509-540, 1988.

[31] Holden, Richard and Malani, Anup. *Can Blockchain Solve the Holdup Problem in Contracts?*. Working Paper, Available at SSRN: https://ssrn.com/abstract=3093879, 2017.

[32] Huberman, Gur, Leshno, Jacob D., and Moallemi, Ciamac. *Monopoly without a Monopolist: An Economic Analysis of the Bitcoin Payment System.* The Review of Economic Studies, 0 1-30, 2021.

[33] Inderst, Roman. *Durable Goods with Quality Differentiation.* Economics Letters 100 (2): 173–177, 2008.

[34] Kahn, Charles. *The Durable Goods Monopolist and Consistency with Increasing Costs.* Econometrica: 275–294, 1986.

[35] Karp, Larry S. *Monopoly extraction of a durable non-renewable resource: failure of the Coase conjecture.* Economica: 1–11, 1993.

[36] Laffont, Jean-Jacques, and Tirole, Jean. *The Dynamics of Incentive Contracts.* Econometrica 56 (5): 1153–1175, 1988.

[37] Laffont, Jean-Jacques, and Jean Tirole. *Adverse Selection and Renegotiation in Procurement.* The Review of Economic Studies 57 (4): 597–625, 1990.

[38] Lomys, Niccolò, and Yamashita, Takuro. *A Mediator Approach to Mechanism Design with Limited Commitment*, Working Paper, 2021.

[39] Maestri, Lucas. *Dynamic Contracting Under Adverse Selection and Renegotiation.* Journal of Economic Theory 171 (C): 136–173, 2017.

[40] McAdams, David, and Michael Schwarz. *Credible Sales Mechanisms and Intermediaries.* American Economic Review 97 (1): 260–276, 2007.

[41] Mcafee, R. Preston, and Wiseman, Thomas. *Capacity Choice Counters the Coase Conjecture.* The Review of Economic Studies 75 (1): 317—331, 2008.

[42] Montez, João. *Inefficient Sales Delays by a Durable-good Monopoly Facing a Finite Number of Buyers.* The RAND Journal of Economics 44 (3): 425–437, 2013.

[43] Nava, Francesco, and Schiraldi, Pasquale. *Differentiated Durable Goods Monopoly: A Robust Coase Conjecture.* American Economic Review: 109 (5), 1930–1968, 2019.

[44] Ortner, Juan. *Durable Goods Monopoly with Stochastic Costs.* Theoretical Economics 12 (2): 817–861, 2017.

[45] Perea, Andrés. *A Note on the One-Deviation Property in Extensive Form Games.* Games and Economic Behavior 40 (2), 322-338, 2002.

[46] Simon, Leo K. and Zame, William R. *Discontinuous Games and Endogenous Sharing Rules.* Econometrica 58 (4), 861-872, 1990.

[47] Sobel, Joel. *Durable Goods Monopoly with Entry of New Consumers.* Econometrica 59 (5): 1455-1485, 1991.

[48] Stokey, Nancy L. *Rational Expectations and Durable Goods Pricing.* Bell Journal of Economics 12: 112–128, 1982.

[49] Strulovici, Bruno *Contract Negotiation and the Coase Conjecture: A Strategic Foundation for Renegotiation-Proof Contracts.* Econometrica 85 (2): 585–616, 2017.

[50] Takeyama, Lisa N. *Strategic Vertical Differentiation and Durable Goods Monopoly.* The Journal of Industrial Economics 50 (1): 43–56, 2002.

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

[52] Tinn, Katrin. *'Smart' Contracts and External Financing.* Working Paper, Available at SSRN: https://ssrn.com/abstract=3072854, 2019.

[53] von der Fehr, Nils-Henrik, and Kuhn, Kai-Uwe. *Coase versus Pacman: Who Eats Whom in the Durable-Goods Monopoly?*. Journal of Political Economy, 103(4): 785–812, 1995.

[54] Waldman, Michael. *A New Perspective on Planned Obsolescence.* The Quarterly Journal of Economics 108 (1): 273–283, 1993.

[55] Wang, Gyu Ho. *Bargaining Over a Menu of Wage Contracts.* The Review of Economic Studies 65 (2): 295–305, 1998.