

Strulovici, Bruno

Working Paper

Contract negotiation and the coase conjecture: A strategic foundation for renegotiation-proof contracts

CSIO Working Paper, No. 0131

Provided in Cooperation with:

Department of Economics - Center for the Study of Industrial Organization (CSIO), Northwestern University

Suggested Citation: Strulovici, Bruno (2015) : Contract negotiation and the coase conjecture: A strategic foundation for renegotiation-proof contracts, CSIO Working Paper, No. 0131, Northwestern University, Center for the Study of Industrial Organization (CSIO), Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/142005>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Contract Negotiation and the Coase Conjecture: A Strategic Foundation for Renegotiation-Proof Contracts*

Bruno Strulovici
Northwestern University

July 11, 2015

Abstract

What does contract negotiation look like when some parties hold private information and negotiation frictions are negligible? This paper analyzes the above question and provides a foundation for renegotiation-proof contracts in a related environment. The model extends the framework of the Coase conjecture to situations in which the quantity or quality of the good is endogenously determined and to more general environments in which the traded goods are complements or substitutes. All equilibria converge to a unique outcome as frictions become negligible, which is separating, efficient, and straightforward to characterize.

1 Introduction

Real negotiations contain a puzzle: On the one hand, parties often try to withhold private information until a deal is reached. On the other hand, the very fact of agreeing to a deal, and the timing of this agreement, reveals to some extent the parties' real stakes in the negotiations. How is this new information incorporated in the final agreement, when players can freely renegotiate their contract?

This puzzle does not arise in the best-known models of negotiation, because of their specific structure. In Coase's model of a durable-good monopolist (Coase (1972)), for instance, buyers

*I am grateful for comments from Mehmet Ekmecki, Drew Fudenberg, Lucas Maestri, Jim Malcomson, Paul Milgrom, Alessandro Pavan, Marcin Peski, Ariel Rubinstein, Larry Samuelson, and Asher Wolinsky, Alex Wolitzky's discussion at the 2013 Cowles Conference in Economic Theory and Takuo Sugaya's discussion at the 2014 Winter Meeting of the Econometric Society, as well as for Mikhail Safronov's excellent research assistance. Various versions of this project were presented at Princeton University, Caltech, UCSD, Oxford University, the École Polytechnique Fédérale de Lausanne, the Séminaire Roy in Paris, the University of Illinois at Urbana-Champaign, the London School of Economics, the University of Essex, the University of Pennsylvania, Duke University, the 2013 Cowles Conference in Economic Theory, Kyoto University (KIER), the Stanford Institute for Theoretical Economics (SITE 2013), the Asian Meeting of the Econometric Society (2013), Columbia University, Western University, the University of Toronto, the 2014 Winter Meetings of the Econometric Society, the Toulouse School of Economics, and NYU Stern. I am grateful to participants for numerous reactions. This research is supported by an NSF Grant (No.1151410) and a fellowship from the Alfred P. Sloan Foundation.

only need one unit of the good. Any sale is therefore efficient, and further negotiation is pointless regardless of the information revealed about the buyer by the timing of his purchase. Similarly, when players with privately known patience bargain over splitting a pie (Rubinstein (1985)), any split of the pie is ex-post efficient and further negotiation is pointless regardless of what players learn about each other's patience.

In richer contractual environments, however, an initial agreement may be inefficient. For example, buyers may value multiple units or various qualities of the good being sold, in which case a buyer's early decision to buy the good may reveal to the seller the benefit of proposing more units or higher qualities of the good, spurring further negotiations. Furthermore, these considerations generally affect the buyer's and the seller's incentives from the very beginning of negotiations.

This paper's objective is to characterize the possible outcomes of negotiation in these richer contractual environments, when the ability to refine contracts is (almost) unrestricted. It provides a dynamic resolution of the above puzzle by describing the gradual succession of agreements leading to the final outcome. This resolution helps answer the following questions: How gradual are the agreements? How fast do they incorporate private information? How much efficiency is lost in the process? How do outcomes depend on the buyer's type and on the seller's initial belief about this type? And how do they depend on the initial contractual relationship between the two parties?

The analysis is based on an explicit protocol of negotiation in which one party – the agent – possesses private information while the other party – the principal – makes the offers. The framework is broader than the buyer-seller model with quasilinear preferences and linear cost. In particular, the parties' preferences may treat the traded goods as complements or substitutes.

The main result of this paper is that all equilibria of the negotiation game converge to a unique outcome, which is fully separating and efficient, as negotiation frictions become negligible. The outcome of negotiation is renegotiation-proof in the sense that no surplus may be gained from renegotiating them any further. Seen in this light, the paper provides a strategic foundation for renegotiation-proof contracts. The outcome depends only on the initial relation – or absence thereof – between the principal and the agent, not on the principal's initial belief about the agent's type.

In the negotiation protocol considered here, the principal can propose new contracts or changes to a previously accepted contract at any time, including after the agent has accepted or rejected a previous proposal. The principal is thus unable to commit *not* to renegotiate the contract. Time is divided into negotiation rounds; at each round, the principal can propose various changes to the current agreement. The agent then accepts a proposal or rejects all changes. This exchange captures in a stylized fashion the idea of a gradual agreement formation in real negotiations; here, any agreement is interpreted as a (possibly oral) bilateral contract, binding unless both parties

agree to replace it by another contract.¹

For the final outcome to be well defined, negotiations must end somehow, and the protocol relies on a particular concept of negotiation friction: at the end of each round, negotiations are exogenously interrupted with a fixed probability, η , in which case the current agreement is implemented. This interruption may be interpreted in various ways. For example, suppose that parties are negotiating a risk-sharing contract, each dimension of which concerns a state of the world. Then, the realization of the state of the world (or its public announcement) makes any further negotiation moot. In a sales contract application, interruption may be coming from a third party, supplier or customer, demanding a commitment or service which requires the immediate implementation of the contract.² The interruption probability captures the negotiation friction. When η is equal to 1, the protocol reduces to full commitment since the first proposal is also the last one. In this case, it is well-known that the principal typically distorts the allocation of some types of the agent, causing ex post inefficiency. This paper's interest lies in the opposite case, in which negotiation frictions are negligible (η goes to 0). This case should be interpreted as parties having arbitrarily frequent opportunities to negotiate: the time interval between consecutive rounds is so small that parties become unlikely to be exogenously interrupted in any such interval. In this interpretation, the time interval between rounds is of order η . Therefore, even though it takes an increasing number rounds for negotiations to stop, as η goes to zero, the expected stopping time of negotiation is independent of η . In particular, it remains bounded as η goes to zero.³

While the ability to freely modify past agreements seems necessary to guarantee an ex post efficient outcome, establishing that it is sufficient involves complex issues. To appreciate the difficulty, consider again the standard durable-good monopolist. Efficiency in that context means that the good is sold without any delay, and was established by Gul, Sonnenschein, and Wilson (1986) as the discount rate between consecutive periods goes to zero.⁴ The proof is sophisticated even in

¹In practice, reneging on an agreement is costly even if they was made orally or informally. Many jurisdiction recognize oral contracts as legally binding. Reneging on an informal agreement is also costly by damaging the reputation of the reneging party and putting a strain on further negotiations.

²Parties may also be forced to implement the current form of the contract due to one of the parties becoming unable to pursue negotiations further or due to some unforeseen contingency.

³Precisely, suppose that the time interval between two rounds is equal to $a \times \eta$ for some $a > 0$. Then the interruption time has mean a regardless of η and becomes approximately exponentially distributed with parameter $1/a$ as η goes to zero.

⁴The result is shown for the “gap” case and the “no gap” case under some Lipschitz condition on the distribution of types, for weak Markov equilibria (see also Sobel and Takahashi (1983) and Fudenberg, Levine, and Tirole (1985)). Ausubel and Deneckere (1989) show that the conjecture can fail when more general equilibria are allowed. The analysis of the Coase conjecture has been extended to various environments: interdependent values (Deneckere and Liang (2006)), incoming flow of new buyers (Fuchs and Skrzypacz (2010)), and outside options for the buyer (Board and Pycia (2013)). Skreta (2006) takes a mechanism design approach and shows the optimality price posting. All

this considerably simpler contractual environment, where each contract amounts to a single posted price. The key question is to determine whether the seller can benefit from distorting the allocation of the low-valuation buyer by inefficiently delaying the sale, in order to extract some rent from the high-valuation buyer. In richer environments, the question is more complex because i) the signature of any contract may be followed by further negotiations (e.g., contractual covenants, increases in quantities or qualities), ii) the principal may benefit from proposing multiple new contracts at each round instead of single one,⁵ iii) each type of the agent can randomize over all such contracts, and iv) in many contracting problems, the utility of the agent need not be linear or even separable in the contract components.

Due to the complexity of the analysis, the model focuses on a binary information structure: the agent can be of two types, and the corresponding utility functions satisfy a standard single-crossing condition. As a result, there is common knowledge of gains from renegotiation: as long as the types of the agent have not been fully separated, there is a strictly positive surplus to be extracted. The main result is that, as negotiation frictions become negligible zero, *all* equilibria (PBEs) of the negotiation game converge to a unique outcome. The outcome allocation is separating and efficient, and thus renegotiation-proof. Put differently, flexible renegotiation provides a dynamic implementation, without commitment, of efficient allocations.

The type-specific contracts to which all PBE outcomes converge are straightforward to characterize and determine graphically. Unlike the full commitment case (but like the Coase conjecture, which it generalizes), these contracts are *independent* of the initial (non degenerate) belief that the principal holds about the type of the agent. They do depend on the initial contract (or absence thereof, which is formally equivalent to a ‘null’ contract), which may lie in three possible regions of the contract space. In the “No-Rent” region, the principal extracts all surplus of renegotiation regardless of the agent’s type. In the other two regions, there is one region-specific type (“*L*”, say) who gains nothing from negotiation while the other type (“*H*”) gets a positive rent.

The fact that almost efficient contracts are proposed immediately implies that renegotiation plays a relatively minor role *in equilibrium*, even though the *possibility* of renegotiation has a major impact on the the outcome. This suggests that, empirically, one should not infer that renegotiation is impossible or difficult in practice just because the observed renegotiation activity seems negligible. Instead, it may well be that negotiation is feasible and cheap, but finds its expression in the very first contracts that are proposed.

Another contribution of this paper is to establish the existence of a PBE for a negotiation

these models focus on the case in which the buyer can only buy one unit of the good, a single quality of the good is available, and utility functions are quasilinear.

⁵For example, the principal may propose one contract for each type of the agent, or propose multiple almost identical contracts as a communication device to emulate cheap talk.

game with a (relatively) rich contract space. In the present setting, backward induction techniques cannot be applied. Instead, the proof takes a two-step approach: first, prove the existence of an equilibrium in an auxiliary game of perfect information between the principal and the high type of the agent, based on Harris (1985). Second, use that equilibrium to construct an equilibrium of the negotiation game with private information.

Relation to Existing Concepts of Contract Renegotiation

Contract renegotiation with private information has traditionally been studied from two different angles. The first approach is axiomatic, and focuses on “renegotiation-proof” contracts.⁶ It essentially *assumes* that renegotiation leads to an efficient contract, despite the potential difficulties arising when one party holds private information. The second approach focuses on simple renegotiation protocols, in which the principal has a single shot at renegotiating the contract, by making a take-it-or-leave-it offer. This approach typically results in *inefficient* contracts.⁷ The second approach seems incomplete: what, in reality, should prevent the principal from proposing a new contract after learning the inefficiency of the current contract? Such restriction amounts to a strong form of commitment power for the principal, and can even result in full commitment outcomes.⁸ The paper shows that, by dropping the restriction on the number of negotiation rounds, one can reconcile these two approaches.

Relation to Coasian Dynamics for Multi-Period Contracts

While the model’s main interpretation concerns the implementation of a single contract arising from a sequence of temporary agreements, one may alternatively interpret it as follows. In each

⁶See Dewatripont (1989), Maskin and Tirole (1992), Battaglini (2007), Maestri (2015), and Strulovici (2011, 2013). A similar approach has been used to study renegotiation in repeated games with complete information. See, e.g., Bernheim and Ray (1989) and Farrell and Maskin (1989).

⁷See Hart and Tirole (1988) and Fudenberg and Tirole (1990). Wang (1998) considers a more flexible protocol, in which the principal proposes contracts until an agreement is reached. Such protocol leaves a high commitment power to the principal, since he cannot renegotiate any agreement. Indeed, Wang’s main result is that, with this protocol, the principal achieves the full commitment allocation, which is also ex post inefficient.

⁸For example, imposing any *finite* number k of negotiation opportunities results in the full commitment outcome: the principal simply passes the first $k - 1$ opportunities to negotiate the contract, and then proposes the full commitment allocation in the last round. Similarly, Wang (1998) has shown that the principal can implement the full commitment allocation if renegotiation stops as soon as the agent accepts an offer. In Wang’s model, the principal can repeatedly propose menus over quality–wage contracts but renegotiation stops as soon as the buyer has accepted a contract. Beaudry and Poitevin (1993) obtain a similar result if renegotiations break down as soon as a new proposal is *rejected*. In Beaudry and Poitevin, the informed party has the bargaining power and proposes a single new contract at each round. Renegotiation stops in the first round at which the other party refuses the new contract. In a different setting with moral hazard, Matthews (1995) considers one-shot renegotiation by the informed party and obtains ex post efficiency.

round, the current contract is implemented and parties receive the payoff corresponding to this contract. The interruption probability is reinterpreted as the discount rate between periods. In each period, the principal has a single opportunity to renegotiate the contract for future periods. The model becomes formally equivalent to an infinite-horizon version of Hart and Tirole (1988) with divisible goods as in Laffont and Tirole (1990), arbitrary utility and cost functions, and in which the contract is constant until renegotiated. Thus interpreted, this paper shows that *all* equilibria (PBEs) converge arbitrarily fast relative to the discount rate to an efficient, separating outcome, as the discount rate goes to zero.

Maestri (2015) analyzes renegotiation-proof equilibria in a divisible-good, binary-type version of the Hart-Tirole framework. The paper proposes a renegotiation-proofness refinement of PBE, formulated recursively: in each period, the contracts offered by the principal maximize his revenue among all renegotiation-proof continuations, given the current contract and beliefs. The principal can propose at most two contracts in each period and the probability of facing the high type or decreases monotonically as long as this type is not fully revealed. As the discount rate goes to zero, any renegotiation-proof equilibrium becomes efficient arbitrarily fast relative to the discount factor. By contrast, the present paper implies that all PBEs become arbitrarily fast, allowing for arbitrarily large menus of contracts, general utility and cost functions, and an infinite horizon, and focusing on contracts which are constant until renegotiated. Conceptually, the papers' objectives are complementary: while Maestri studies renegotiation-proof equilibria when trade happens over time, the present paper shows – under its main interpretation – that if parties have frequent opportunities to bargain before a one-time trade decision, then all equilibria will be approximately efficient, thus providing a non-cooperative foundation for renegotiation-proof contracts.⁹

Relation to Screening and Reputation Models

The paper is also related to the literature on reputation, in which some players are trying to determine whether other players have a “commitment” type.¹⁰ Both this literature and the present paper concern a dynamic screening problem and, in both cases, much of the analysis is devoted to establishing properties of beliefs (such as bounds and convergence rates) along candidate equilibria. Compared to this literature, the present model differs in several ways: i) the “actions” of the players (the types) are endogenous, because the principal chooses which contracts the agent chooses from in each round, ii) the state space is larger, because it includes the last accepted contract, in addition

⁹The technical contributions are also different. For example, an important part of Maestri's analysis lies in its conceptual work to define renegotiation-proof contracts. By contrast, this paper analyzes all equilibria, which creates important challenges described in later sections.

¹⁰See in particular Fudenberg and Levine (1989), Schmidt (1993), Abreu and Gul (2000), Cripps et al. (2005), and Atakan and Ekmekci (2012).




Figure 1: Setting (trade interpretation)

to the principal's belief, and iii) all types of the agent are strategic.¹¹

2 Setting and Overview of the Results

There are two players, a principal (P) and an agent (A) who negotiate a contract lying in some compact and convex subset \mathcal{C} of \mathbb{R}^2 whose components are denoted x_1 and x_2 .

A has a utility function $u_\theta : \mathcal{C} \rightarrow \mathbb{R}$, where $\theta \in \{L, H\}$ denotes his type, and P has a cost function $Q : \mathcal{C} \rightarrow \mathbb{R}$. The functions u_L , u_H , and Q are twice continuously differentiable and have strictly positive derivatives with respect to x_1 and x_2 ; u_L and u_H are concave and Q is convex. Although it is convenient to think of the agent's type as being "high" (H) or "low" (L) – especially to connect the paper's results to the existing literature – the role played by each type is in fact determined by the initial contract, as explained below.

A contract $C = (x_1, x_2) \in \mathcal{C}$ is θ -efficient if it is the cheapest contract in \mathcal{C} providing θ with some given utility level. For each θ , let \mathcal{E}_θ denote the set of θ -efficient contracts in the interior of \mathcal{C} ; θ 's iso-utility curve and P's iso-cost curve are tangent at any such contract. To rule out pathological

¹¹The richer state space combined with the nonlinear geometry of the problem (due to the agent's and principal's nonlinear utility functions) make the problem particularly difficult to analyze. Some of the listed differences can formally be incorporated into the standard reputation framework. For example, the action space of the agent may be assumed to be fixed by setting a default value if the agent chooses anything outside of the principal's proposed set. This formal similarity does not however affect the substantive differences between the settings.

cases, it is assumed that the efficiency curves $\mathcal{E}_L, \mathcal{E}_H$ are smooth and upward sloping. The setting is represented on Figure 1 in the context of a trade application (other applications are described later in this section). It will also be assumed that, given any contract C on \mathcal{E}_θ , P 's isocost curve going through C and θ 's isoutility curve going through C do not both have a zero curvature at C . This assumption is satisfied when the agent has a quasi-linear utility function and his valuation for the good is strictly concave.¹²

The functions u_L and u_H are required to satisfy a standard single-crossing condition: iso-utility curves of L are steeper than those of H at their intersection point. This implies that the efficiency curve \mathcal{E}_L lies to the lower right of \mathcal{E}_H . \mathcal{C} can therefore be partitioned into three regions separated by \mathcal{E}_L and \mathcal{E}_H . Contracts in the inner region are said to be in the ‘No Rent’ configuration, while contracts below \mathcal{E}_L (above \mathcal{E}_H) are in the ‘ H -Rent’ (‘ L -Rent’) configuration. The set of contracts in the H -Rent configuration will be denoted by \mathcal{H} . \mathcal{C} represents an Edgeworth box, delimited by the sum of endowments of the agent and the principal. Each contract C represents the a final allocation for the agent, the efficiency curve \mathcal{E}_θ is the ‘contract curve’ corresponding to type θ , and the status quo R_0 represents the endowment of the agent before any trade.

The Negotiation Game

The game unfolds as follows. First, the agent privately observes his type θ ; P has a prior characterized by the probability $\beta_0 = Pr(\theta = H)$. The game starts with a default “contract” $R_0 \in \mathcal{C}$, which represents the current engagement between the principal and the agent. In many settings this initial “contract” would simply represent the absence of any past engagement, as in most contracting models.¹³ In the durable-good monopolist application, the initial contract is the “no sale” outcome; in the trade application described below, the initial contract corresponds to the initial endowment of the agent before any trade. Specifying R_0 explicitly is useful for two reasons: firstly, it will allow us to treat the initial round like any later round in which an agreement has already been made, which significantly simplifies the exposition. Secondly, there are environments such as the trade application in which the initial contract (initial endowment of the agent) plays an important role on the outcome of negotiations.

There are countably many potential rounds, indexed by $n \in \mathbb{N}$. At each round n , P can propose

¹²Alternatively, the agent’s utility could be in the linear in the good if P ’s cost function is strictly convex. The assumption does rule out settings in which both A ’s and P ’s indifference curves are linear, but even this case can be approximated arbitrarily closely. It guarantees that the two curves take off cleanly from each other as one moves away from C , a property used to compute a lower bound for the inefficiency of some contracts (see Lemma 12).

¹³Indeed, assuming no pre-existing contract is equivalent to assuming a contract that entails no transaction. Seen this way, any contracting model may be seen as implicitly assuming an initial contract, perhaps trivial, between parties.

a menu M_n of contracts in \mathcal{C} . In terms of interpretation, proposing contracts or *changes* to the current contract are formally equivalent; the former formulation is used here. The number of contracts in M_n is bounded by some constant $G \geq 2$ that is arbitrary but fixed throughout the game.¹⁴ The agent chooses a contract in M_n or holds on to the last accepted contract, R_n . Any mixed strategy over the choice set $M_n \cup \{R_n\}$ is allowed; indeed, as in the standard Coase conjecture and in reputation models, mixing plays a key role in the analysis. At the end of each round, negotiations are frozen with probability $\eta \in (0, 1]$, in which case the last accepted contract, R_{n+1} , is implemented. Otherwise, negotiations move on to the next round. The event of a negotiation freeze will be hereafter referred to as a “breakdown.” This term means that *future* negotiations are terminated. However, agreement formed before the breakdown (or the initial contract, if no agreement was formed) is still valid. Although it is not necessary for the analysis, one should interpret η as being proportional to the time interval between two rounds. With this interpretation, even though it takes an increasing number rounds for negotiations to stop, as η goes to zero, the expected stopping time of negotiation is independent of η . In particular, it remains bounded as η goes to zero.¹⁵

Letting $\{R_n\}$ denote the stochastic process of contracts entering each round n , the agent’s expected utility is equal to

$$\mathcal{V}_\theta = E \left[\sum_{n \geq 0} (1 - \eta)^n \eta u_\theta(R_{n+1}) \right]$$

while P’s expected cost is

$$\mathcal{Q} = E \left[\sum_{n \geq 0} (1 - \eta)^n \eta Q(R_{n+1}) \right].$$

The parameter η represents the *negotiation friction* of the game.¹⁶ The objective of this paper

¹⁴There is no guarantee that proposing only two contracts at each round is without loss of generality. As Bester and Strausz (2001) have shown, the set of implementable outcomes can require more “messages” (or contracts) than the number of types of the agent, even in a two-period setting. While their modified revelation principle without commitment implies that *incentive efficient* contracts may require only 2 messages, here one must consider all possible continuation equilibria, including incentive inefficient ones (indeed, an inefficient continuation equilibrium may provide incentives at earlier stages of the game). Since \mathcal{C} is a continuum, one could also imagine that P exploits a higher G by proposing several almost identical contracts as a ‘cheap’ way of communicating with the agent. The assumption that G is finite is used in Lemma 6, and only this lemma, to bound the probability that the principal’s posterior jumps up as a result of an unlikely contract choice by the agent.

¹⁵Precisely, suppose that the time interval between two rounds is equal to $a \times \eta$ for some $a > 0$. Then the expected interruption time is equal to a for all η , and the interruption time becomes approximately exponentially distributed with parameter $1/a$, as η goes to zero.

¹⁶There is another interpretation of the setting where η is the discount rate and the parties receive payoffs at each period of the on going relationship. This interpretation is discussed at the end of Section 5.




Figure 2: Renegotiation outcomes

is to characterize the PBEs of the game as the friction η goes to zero. Existence of a PBE is guaranteed by Theorem 1, whose proof is in Appendix A.

THEOREM 1 *For each $\eta \in (0, 1]$, there exists a PBE of the negotiation game.*

For any contract $R \in \mathcal{C}$, let $E_H(R)$ and $E_L(R)$ denote the cheapest pair of H - and L -efficient contracts such that each type $\theta \neq \theta'$ weakly prefers $E_\theta(R)$ to $E_{\theta'}(R)$ and to R . That pair is well defined for each possible configuration of R .¹⁷ Figure 2 represents these concepts for the case of CRRA utility functions and a linear cost function, and where \mathcal{C} is the Cartesian product $[0, \bar{x}_1] \times [0, \bar{x}_2]$.

The proof of Theorem 2, below, exploits the inefficiency stemming from both types of the agent getting the same contracts at the time of a breakdown. To guarantee that such an inefficiency does exist, one must rule out situations in which contracts lying at the boundary at the contract space

¹⁷If R is in the No-Rent configuration, $E_\theta(R)$ is simply the θ -efficient contract that gives θ the same utility as R . If R is in the H -Rent configuration, then $E_L(R)$ is similarly defined, while $E_H(R)$ is the H -efficient contract that gives H the same utility as $E_L(R)$. Because that contract gives a strictly higher utility to H than the initial contract R , H must be getting a positive rent in any equilibrium, hence the name of that configuration. A symmetric construction obtains if R is instead in the L -Rent configuration.

arise in equilibrium, since such contracts may be efficient for both types.¹⁸

This is achieved by focusing on the following contracts: a contract R_0 is **regular** if either it is in the No-Rent configuration or if it satisfies the following condition, stated when R_0 lies in the H -Rent configuration (an analogous condition is imposed when R_0 is in the L -Rent configuration): for any $R' \in \mathcal{H}$,

$$u_H(E_L(R')) \geq u_H(R_0) \Rightarrow E_L(R') \neq E_H(R') \quad (1)$$

When $R_0 \in \mathcal{H}$, one may show that any contract $R' \in \mathcal{H}$ arising in equilibrium satisfies $u_H(E_L(R')) \geq u_H(R_0)$ (Proposition 1, part iv). Therefore, the premise of the condition is always satisfied on path. If R_0 is regular, the condition then implies that any breakdown with a non-degenerate belief yields some inefficiency. The condition is straightforward to check for any R_0 and is satisfied as long as R_0 is so bad for the agent that negotiated contracts can hit the lower boundary of the contract space.¹⁹ In Figure 2, all contracts are regular except for the origin.

THEOREM 2 *Consider any regular contract R_0 , belief $\beta_0 \in (0, 1)$, and $\varepsilon > 0$. There exists $\bar{\eta} > 0$ such that the following statements hold for any $\eta \leq \bar{\eta}$ and corresponding PBE:*

A: The expected utility of each type θ is bounded below by $u_\theta(E_\theta(R_0)) - \varepsilon$.

B: The probability that each type θ gets a contract within a distance²⁰ ε of $E_\theta(R_0)$ when renegotiation breaks down is greater than $1 - \varepsilon$.

Statement B implies that the outcomes of renegotiation must get arbitrarily close to ex-post efficiency as the renegotiation friction η goes to zero, since each contract $E_\theta(R_0)$ is θ -efficient. That statement is a corollary of Statement A, to which the quasi-totality of the proof is devoted.

Theorem 2 implies that P always gets some surplus from negotiation: When the contract is in the No-Rent configuration, P extracts, in fact, all the surplus regardless of the agent's type. When R_0 is in the H -Rent configuration, P extracts all the surplus from negotiating with L , and extracts some additional surplus in case he is facing H (the surplus obtained by moving from $E_L(R_0)$ to $E_H(R_0)$).

¹⁸At the boundary, even the strict single crossing property is ineffectual in separating efficient contracts: even if the isoutility curves for the low and high types have different slopes for a contract lying at the boundary, that contract may be efficient for both types.

¹⁹Even when regularity does not hold, it is easy to slightly perturb the agent's utility function so as to satisfy it, as illustrated by Section 3 to recover the standard Coase conjecture. Furthermore, if the utility and cost functions are defined on some upper orthant $\mathcal{O} = [\ell_1, +\infty) \times [\ell_2, +\infty)$ containing the contract space \mathcal{C} , one can always expand \mathcal{C} to the smallest rectangle \mathcal{C}' containing \mathcal{C} and such that \mathcal{E}_L (resp. \mathcal{E}_H) hits the boundary of \mathcal{C}' on its right (resp. upper) edge. The regularity condition then only rules out initial contracts that gives such a low utility to the agent that negotiated contracts may hit a lower boundary of \mathcal{C} .

²⁰The statement holds for any norm on \mathbb{R}^2 .

Applications

1. **Durable Good Monopolist.** A is a buyer with quasi-linear utility $u_\theta(C) = \theta \bar{u}(x_2) + x_1$, where x_2 is the quantity of the good sold by P, x_1 is A's wealth, and u is A's concave utility function.²¹ The initial contract, R_0 , is equal to $(\bar{x}_1, 0)$ where \bar{x}_1 is A's initial wealth. P's cost is $Q(x_1, x_2) = cx_2 + x_1$, where $c > 0$ is the marginal cost for producing the good and x_1 captures how much wealth P "leaves" to A.²²

2. **Labor Contract.** P is a potential employer and A is a worker. $-x_2$ represents A's effort and x_1 is his wage. A gets a utility $u_\theta(C) = \theta \psi(-x_2) + x_1$ from contract C , where ψ is a factor entering A's cost of effort, increasing in its argument, and θ is a worker-specific factor entering this cost. The status quo $R_0 = (0, 0)$ represents unemployment, while P's profit is $\Pi(x_1, x_2) = -Q(x_1, x_2) = -x_2p - x_1$, where $p > 0$ is the unit price of the good.

3. **Consumption Smoothing and Insurance.** There are two periods and a single good. The dimensions of \mathcal{C} represent A's consumption in each period. P is a social planner or a bank who can help the agent smooth his consumption. The type θ may be a privately known patience/discount factor, or a distribution parameter that describes how likely the agent is to value the good in the second period. For example, $u(x_1, x_2)$ may be equal to $v(x_1) + \theta v(x_2)$ or to $v(x_1) + E[w(x_2, \tilde{\rho})|\theta]$ where $\tilde{\rho}$ is a taste shock whose distribution is increasing in θ in the sense of first-order stochastic dominance and where w is supermodular, so that $E[\partial w / \partial x(x_2, \tilde{\rho})|\theta]$ is increasing in θ .²³ R_0 is A's autarkic income stream. $Q(x_1, x_2) = p_1x_1 + p_2x_2$, where p_t is the market price for the good in period t .

4. **Trade.** More generally, the model describes a trade environment in which the dimensions of \mathcal{C} represent distinct goods, with x_i denoting the quantity of good i consumed by A. Type L cares more about the first good than the second, relative to H . P (like A) has convex preferences and Q is the negative of a utility function representing P's preferences. R_0 denotes A's initial holdings of the goods.

5. **Risk Sharing** Suppose that parties are negotiating a risk-sharing contract. Each dimension corresponds to a state of the world. The quantity x_i denotes the transfer of a good from P to A in state i . The relevant transfer is performed when the state of the world is realized. According to one interpretation, types differ in their subjective beliefs about the state of the world, with L assigning a higher probability to the first state of the world than H does. Another interpretation

²¹The iso-level curves of u_θ have positive curvature as long as the second derivative \bar{u} is strictly negative, as is easily checked. A similar observation applies to ψ in the labor contract application.

²²P's profit is $\Pi(t, x_2) = t - cx_2$, where t is how much the agent pays P. Letting $t = \bar{x}_1 - x_1$, we obtain the formulation in terms of the cost function Q .

²³This application is explored in detailed by Strulovici (2013).

is simply that L values the good more in state 1 than in state 2, relatively to H .

3 Relation to the Coase conjecture

In the standard Coase conjecture, buyers value only one unit of the good. As a result, the set of efficient outcomes is the same for all buyer types: it consists of all outcomes for which the buyer gets the good (regardless of the sale price).²⁴ However, when goods are divisible or available in multiple qualities, it becomes rather restrictive to assume that the same contracts which are efficient for one type of the agent are also the ones which are efficient for other agent types. In fact, when efficient contracts lie in the interior of the contract space, the strict single-crossing property implies that whatever contract is efficient for one agent type is inefficient for the other.

This distinction explains why, in this paper, the principal can always extract some surplus from the “high” type, in contrast to the Coase conjecture. It also explains why, when the initial contract lies between efficiency curves of both types – something that is obviously impossible in the setting of standard Coase conjecture, since the curves coincide – the principal can extract *all* the surplus from negotiations. In this sense, the Coase conjecture appears to be non-generic because it relies on the assumption that agent types share the same efficient outcomes.

However, because the model does not impose any lower bound on “how far” the efficiency curves of both types have to be, it is easy to recover the Coase conjecture as a limit of Theorem 2. To see this, suppose that the first contractual dimension represents the agent’s wealth, while the second dimension represents the probability that the agent gets the good (alternatively, it could represent the quantity, between 0 and 1, of a divisible good). The initial contract is $(W, 0)$, where W is the agent’s initial wealth. The principal incurs a cost c for the good (or marginal cost c , in the divisible-good interpretation), resulting in linear isocost curves. To bridge the two settings, we fix an arbitrarily small constant δ and define the agent’s utility as follows. The agent’s utility is quasilinear²⁵ for $x_2 < 1 - 2\delta$, given by $u(x_1, x_2; \theta) = \theta x_2 + x_1$ with $\theta_H > \theta_L > c$; the iso-utility curves are then kinked for $x_2 > 1 - 2\delta$ as depicted on Figure 3, so as to guarantee that all efficient contracts for H (L) involve a fixed probability $x_2 = 1 - \delta(1 - 2\delta)$ of getting the good. The red (blue) curves represent the iso-utility curve of the high (low) type. The boundary of regular contracts is

²⁴This statement concerns the “positive gap” case (lowest buyer valuation exceeds seller cost), which is the relevant comparison here.

²⁵The paper assumes a nonzero curvature condition for iso-utility curves. However, the curvature can be arbitrarily close to zero, so that the quasilinear case can be approximated arbitrarily well, and this can be done simultaneously as δ is taken to zero. Moreover, the Coase conjecture really concerns the value of the agent for zero and one unit of the good. It doesn’t have to be linear for intermediary quantities of the good, though this restriction is natural if one thinks of the “quantity” as being the probability of that the agent gets the good.




Figure 3: Recovering the standard Coase conjecture

shown on the left of the figure. All contracts of \mathcal{H} lying to the right of that boundary, including R_0 , are regular.

This situation approximates, for δ arbitrarily small, the setting of the Coase conjecture, in which it is always efficient to sell the good to the agent. For any $\delta > 0$, however, it also satisfies the assumptions of Theorem 2. In particular, the efficiency curves of the types are separated.

The Coase conjecture is then recovered as follows: if P were sure to face H , he would move to the contract C_H appearing on Figure 3. With uncertainty about the buyer's type, however, Theorem 2 implies that the outcome is given by the contracts $E_H(R_0), E_L(R_0)$, which converge to the same contract as δ goes to zero. Both types of the buyer obtain essentially the same outcome, which is the (almost) sure sale of the good at the same price. The high type gets a rent corresponding to the distance between $E_H(R_0)$ and C_H , while L gets no rent.

As the previous discussion illustrates, the Coase conjecture is rather special in its assumption of a common efficiency curves. Seen in this light, Theorem 2 provides a way to check which part of the Coase conjecture is robust to a more general setting. In particular, the efficiency part continues to hold, but the seller's inability to extract any surplus disappears.

Finally, the richer contract space considered here gets rid of a stark discontinuity arising in the Coase conjecture between the gap and no gap cases. With two types, in particular, H 's rent increases as L 's valuation v_L (and hence the equilibrium price) becomes lower. However, when L 's valuation reaches P's marginal cost c , turning into the “no gap” case, H 's rent suddenly drops to zero and P's profit leaps up from zero to $\beta_0(v_H - v_L)$. Consider a similar exercise in the setting

of Theorem 2, where L' efficiency curve is lowered until it goes through R_0 . The surplus that P extracts from H then varies continuously, until \mathcal{E}_L goes exactly through R_0 and P extracts all surplus from H .²⁶

4 Explaining Theorem 2

This section provides some intuition for Theorem 2 and explains the structure of its proof.

Intuition


An important aspect of Theorem 2, among its most surprising ones, is the fact that *all* equilibria converge to a unique and efficient outcome. The coarsest explanation for this fact is that the situation resembles the setting of the Coase conjecture with two types. In this environment we know (Gul, Sonnenschein, and Wilson (1986)) that there is a unique equilibrium, which may be determined by backward induction on the probability of facing the high type, and which becomes efficient (no delay) as the discount rate converges to zero (discounting and negotiation breakdowns are equivalent in this case). The present setting is more complex, as there is no ‘final’ round to which backward induction may be applied. The key is to show that all equilibria have the same essential features: Firstly, a *quasi*-backward induction argument can be provided to replace exact backward induction. Secondly, all equilibria contain at least one on-equilibrium-path sequence that behaves qualitatively like the unique equilibrium of the Coase conjecture with a positive gap.

Theorem 2 stands in contrast to Ausubel and Deneckere’s (1989) analysis of the no-gap case for the durable good monopolist, in which many equilibria exist. As explained above, the intuition for this discrepancy is that the present setting concerns an “endogenous but positive gap” which functions quite differently from the no gap setting which they analyze.


This intuition may be further developed by considering several restrictions on the principal’s type and number of contacts that he may propose. The strongest of these restrictions turns the setting into one which is analytically equivalent to the standard Coase conjecture (Level ‘1’ below). It is then relaxed in several steps (Levels ‘2’ and ‘3’) to reach the most general setting (Level ‘4’). This section will explain why one should expect the conclusions of Theorem 2 to continue to hold after each relaxation of the restrictions.

Preliminaries Theorem 2 encompasses in one statement two very different configurations: the No-Rent and H -Rent ones (the L -Rent configuration is symmetric). The No-Rent configuration is the easiest one to analyze because the full-commitment allocation is ex post efficient in this case: no


²⁶The rent that P extracts from H need not be monotonic throughout this transformation of \mathcal{E}_L , but it evolves continuously and reaches its maximum as \mathcal{E}_L goes through R_0 .




(1) Like Coase: Contracts for H until P gives up



(2) Reference R_n becomes gradually more efficient



(3) “Free” contracts



(4) General Menu (Non monotonic beliefs)

Red dots are contracts proposed at round n

Figure 4: From the Coase conjecture to the General Case

type benefits from mimicking the other type if P proposes the efficient contracts extracting all the surplus from each type (Figure 2). While it is straightforward to show that P immediately proposing these contracts is an equilibrium, it is more challenging to show that this equilibrium is unique. The core of the proof is a “cost undercutting” argument similar to the one applied to utility levels in Rubinstein’s (1982) bargaining model with complete information. Some challenges arise due to the presence of private information, and the precise definition of PBE (as defined by Fudenberg and Tirole (1991)) is used to show uniqueness. Furthermore, the geometry of the problem requires some specific use of the single-crossing crossing condition to apply the undercutting argument (see Proposition 1).

The analysis is significantly harder when R_0 is in the H -Rent configuration. To understand the arguments underpinning it, it is helpful to decompose it into four layers of complexity, corresponding to decreasingly stringent restrictions on P 's strategies, and represented on Figure 4.

Level 1 considers a restriction in which P can propose, in each round n , either an H -efficient contract C_n or the pair $\{E_H(R_0), E_L(R_0)\}$. In this case, L rejects all contracts until being offered $E_L(R_0)$, just like the low-valuation buyer in the Coase conjecture rejects all offers until the price drops to his valuation. In equilibrium, the contracts C_n 's are chosen so as to make H indifferent between accepting them or waiting for future rounds. H randomizes in each round between taking C_n and waiting. This situation is equivalent to the standard Coase conjecture: as long as H rejects the H -efficient contracts C_n 's, any breakdown causes the parties to be stuck with the initial contract R_0 . Since R_0 is inefficient for L , it is optimal for P to jump to $E_L(R_0)$ if the probability β_n of facing H is low enough. Backward induction in beliefs can be applied from there and used to uniquely pin down the equilibrium.

Now suppose (Level 2) that P is allowed to propose *two* contracts at each round, C_n and R_{n+1} . The contract C_n is H -efficient, as in Level 1, while the contract R_{n+1} gives L the same utility as R_0 , but is more efficient as it gets gradually closer to $E_L(R_0)$. This makes the analysis more challenging because even when β_n is small, the L -inefficiency of R_n may be small enough that P doesn't mind continuing the negotiation. As a result, there is no belief threshold below which P gives up trying to extract any rent, and backward induction can no longer be used. To address this problem, the strategy of the proof is to replace proper backward induction by an infinite-horizon equivalent described by Proposition 3. Intuitively, this proposition states that when β_n is small, the current contract R_n must be extremely close E_{R_0} (as it would if P just gave up extracting rent) so that the backward-induction argument is almost valid.

In general, the contracts R_{n+1} proposed by P need not give L the same utility as R_0 .²⁷ This new difficulty is represented on the third part of Figure 4 (Level 3). The proof addresses it by providing bounds on how R_{n+1} can "wobble" away from L 's initial isoutility curve. This step involves some elementary differential geometry combined with incentives bounds.

Finally, because we cannot presume the efficiency of continuation equilibria following any given round – this is, after all, what we are trying to establish, and only in the limit! – one cannot a priori rule out equilibria in which P proposes more than two contracts in each round.²⁸ This takes complexity to a much higher level because there can now exist on path choice sequences along which the belief β_n is non monotonic: depending on which of the proposed contracts is chosen

²⁷Proposing an H -efficient contract C_n which only H may accept in equilibrium *is* without loss of generality, as explained in later sections.

²⁸Bester and Strausz's (2001) revelation principle result only applies to constrained efficient continuations.

(the red dots of Level 4) in Figure 4), the belief could a priori jump back up arbitrarily close to 1. Here, the proof consists in showing the existence of *one* sequence of equilibrium choices by the agent along which both β_n and the contracts R_{n+1} behave as in the simpler cases. This part is the most challenging and requires several conceptual and technical steps. It should also prove useful for future work, to extend the analysis to three or more types.

Overview of the Proof

While the increasingly complex levels described above are useful to convey the intuition for Theorem 2, the actual proof of this result directly tackles the general case and is organized with a distinct structure. This part provides a road map of the key steps and ideas used to establish Theorem 2. Most arguments and concepts are purposely simplified; all proofs are in the appendix.

As a first pass to understanding the arguments, the reader may wish to assume that P can only propose two contracts in each round (as in Level 3, above), so that β_n is decreasing along the (unique) choice sequence consisting of contracts in \mathcal{H} . The proof is significantly simpler and Part II (below) becomes unnecessary in this case.

Preliminaries

When P assigns probability 1 to either type of the agent, or when the last accepted contract is in the No-Rent configuration, it is comparatively easy to show that there is a unique continuation PBE: P immediately extracts all the rent from negotiation and efficiency obtains (Proposition 1). The most challenging part of the analysis is to prove the theorem when R_0 is in the H -Rent or the L -Rent configuration; this section focuses on the former case, without loss of generality.

With R_0 in the H -Rent configuration, one may show that, along any PBE, L accepts only contracts in the H -Rent configuration (Lemma 3). Without loss, a PBE looks like this: P proposes at each round an H -efficient contract, C_n , and some contracts in the H -Rent configuration. H mixes over all contracts, while L mixes over all contracts *except for* C_n .

“Jump” deviation: It will be shown that P can, at any round n , propose the efficient contracts $E_H(R_n)$ and $E_L(R_n)$ and have them accepted by their respective types (Lemma 2). Jumping to these efficient contracts is therefore a feasible deviation for P, in any equilibrium. It provides a key upper bound on P’s continuation cost and forms the basis of the incentive compatibility equations used in the proof.

Main question and strategy: The crux of the proof is to show that P cannot do any better than the jump deviation, as η goes to zero. To achieve this, it is sufficient to show that P must leave to H a utility arbitrarily close to $u_H(E_H(R_0))$; the other claims of Theorem 2 follow from that statement.

Contradiction hypothesis and block construction

The proof proceeds by contradiction. Fixing some $\varepsilon > 0$, suppose one may find some arbitrarily small η and an associated PBE for which P extracts a rent of at least ε from H , compared to $u_H(E_H(R_0))$. To fix ideas, let $u_H(0)$ denote H 's expected utility at the beginning of round 0, and suppose that $u_H(0) = u_H(E_H(R_0)) - \varepsilon$.²⁹

The contradiction argument starts by constructing *blocks* of rounds delimited by continuation-utility thresholds for H , denoted $\hat{u}_0 = u_H(0) < \hat{u}_1 < \hat{u}_2 < \dots$, and an on-equilibrium-path sequence of contract offers and choices along which the probability of facing H is guaranteed to drop by at least some factor $q < 1$ between consecutive blocks. The argument and all variables below pertain to this particular choice sequence: they correspond to one specific realization among all possible equilibrium paths.

The thresholds are defined inductively: for block $k + 1$ (starting with block 1, corresponding to $k = 0$) \hat{u}_{k+1} is defined by the equation

$$\frac{\hat{u}_{k+1} - \hat{u}_k}{\hat{e}_k - \hat{u}_k} = \frac{1}{t} < 1 \quad (2)$$

where $\hat{e}_0 = u_H(E_H(R_0))$ and, for $k \geq 1$, we have³⁰

- $\hat{e}_k = u_H(E_H(R_{n_k}))$ where n_k is the last³¹ round of block k , and
- $t > 1$ is a factor that will be defined shortly.

These thresholds were chosen so as to generate a geometric sequence for $\hat{e}_k - \hat{u}_k$, which is then used to compute an upper bound on $\hat{e}_0 - \hat{u}_0$.³² This upper bound is of clear interest because $\hat{e}_0 - \hat{u}_0 = u_H(E_H(R_0)) - u_H(0)$ is exactly the quantity which must be shown to be less than ε in the contradiction hypothesis above.

It is easily shown that H 's continuation utility at the beginning of round n , $u_H(n)$, increases by steps of order at most η between consecutive rounds. In fact, H 's Bellman equation implies that

$$u_H(n) = \eta u_H(R_{n+1}) + (1 - \eta)u_H(n + 1) \quad (3)$$

for any contract R_{n+1} chosen by H with positive probability, which implies the η bound on $u_H(n + 1) - u_H(n)$. Therefore, each block k contains at least in the order of $(\hat{u}_k - \hat{u}_{k-1})/\eta$ rounds. Since each

²⁹In general, $u_H(0)$ is strictly less than $u_H(E_H(R_0)) - \varepsilon$. However, this distinction can be easily addressed by an initialization phase. See Lemma 4.

³⁰The actual definition of \hat{e}_k is more complicated and involves the entire history of play.

³¹Lemma 14 shows that H 's continuation utility must eventually exceed \hat{u}_k in finite time, so that the final round n_k is well-defined.

³²Precisely, one can show that $\hat{e}_0 - \hat{u}_0 \leq [\frac{t}{t-1}]^k (\hat{e}_k - \hat{u}_k)$ for any relevant k . See the argument preceding (22).

round is followed by a breakdown with probability η , this puts a lower bound, of order $\hat{u}_k - \hat{u}_{k-1}$ on the probability of a breakdown within block k . Moreover, each breakdown creates an inefficiency, conditional on facing H , that is bounded below by some constant $D > 0$.³³ Given this inefficiency, the only reason for P to proceed with the equilibrium instead of directly jumping to the efficient contracts $\{E_H(R_n), E_L(R_n)\}$ is to extract some rent from H . However, this rent is bounded above by $\hat{e}_k - \hat{u}_k$ because \hat{e}_k is what P would leave to H if he gave up on rent extraction, while \hat{u}_k is H 's continuation utility if P goes on with the equilibrium. This implies an upper bound of the form $a(\hat{e}_k - \hat{u}_k)$, for some $a > 0$, on the cost reduction that P can hope to get through rent extraction.

Taken together, these observations imply an upper bound on the probability $\hat{\mu}_k$ that H rejects all H -efficient contracts (the C_n 's meant only for him) throughout block k :

$$\hat{\mu}_k \leq \frac{a}{a + D} \frac{\hat{e}_k - \hat{u}_k}{\hat{u}_{k+1} - \hat{u}_k} = t^{-1} < 1 \quad (4)$$

where the equality comes from (2) and from choosing a particular value for the factor t introduced earlier: $t = \sqrt{\frac{a+D}{a}}$. Since H 's probability of accepting only contracts in \mathcal{H} is bounded above by (4) and L accepts contracts in \mathcal{H} with probability 1, P 's posterior probability of facing H conditional of facing these contracts must go down by *at least* some factor related the bound provided by (4), *at least for some choice sequence* of contracts in \mathcal{H} . Letting $\hat{\beta}_k$ denote P 's assigned probability of facing H at the end of block k , this implies that

$$\hat{\beta}_k \leq g\hat{\beta}_{k-1}$$

for some factor $g < 1$ independent of η . As previously mentioned, this inequality is shown to hold only along a specific choice sequence on the equilibrium path, but this will suffice for making the contradiction argument.

Adapting backward induction to an infinite horizon

The above construction is interrupted at the first block K for which the difference $\hat{e}_K - \hat{u}_K$ is less than $\bar{W}\eta$, where \bar{W} is chosen high enough to guarantee that $\hat{e}_K - \hat{u}_K$ lies *above* the threshold $\underline{W}\eta$ for some constant $\underline{W} > 0$. This lower bound $\underline{W}\eta$ is guaranteed to exist and plays an important role in establishing the contradiction, as explained below.

The remainder of the analysis hinges on the value of the posterior $\hat{\beta}_K$ at the end of block K . To see this, consider the standard Coase conjecture with 2 types. In that simpler setting, there exists a belief threshold $\hat{\beta}$ below which the seller immediately sets the price at the low-buyer valuation, completely giving up on rent extraction. If such a threshold existed in our problem,

³³The existence of such a constant hinges on the geometric assumption of regularity, as explained before stating Theorem 2. See also Lemma 16.

then because $\hat{\beta}_k \leq g^k \beta_0$, this would put a *fixed* bound on the number of blocks until $\hat{\beta}$ is reached. Proceeding backward on the utility thresholds, this would then imply that the initial difference $u_H(E_H(R_0)) - u_H(0) = \hat{e}_0 - \hat{u}_0$ is also of order η . This, for η small enough, would yield the desired contradiction.

Unfortunately, such a belief threshold does not exist here. Unlike the positive gap case of the Coase conjecture, where the “no sale” outcome generates a first-order inefficiency even for the lowest type, here the default contract R_n varies over time and becomes asymptotically L -efficient. Instead of jumping to an L -efficient contract when β_n becomes small enough, the principal makes active proposals yielding ever smaller improvements to the contract until negotiations exogenously break down. As a result, the relatively simple backward induction argument in the standard Coase conjecture must be modified. The key is to establish the following result.

Proposition (Belief Bound) *For any $d > 0$, there exists an upper threshold $\bar{\eta} > 0$ such that $\hat{\beta}_K > \eta^d$ for all PBEs corresponding to any $\eta < \bar{\eta}$.*

To apply this result, let $\rho > 0$ by (uniquely) defined by $g^{-\rho} = \frac{t}{t-1}$ and set $d = \frac{1}{2} \min\{\frac{1}{\rho}, 1\} \in (0, 1/2]$. Proceeding by backward induction on the constructed blocks and applying the above proposition to this value of d , one may easily show that $\hat{e}_0 - \hat{u}_0$ is of order $\sqrt{\eta}$ (see (22) in the Appendix), which yields the desired contradiction. Proving this proposition is challenging, however, and takes up most of the proof.

The intuition for the above proposition is – deceptively – simple: if $\hat{\beta}_K$ were smaller than η^d , for η arbitrarily small, then P would be willing to sacrifice almost all rent extraction from type H , who is extremely unlikely anyway, in order to avoid any inefficiency on type L . This would imply that the current contract must be approximately L -efficient. Formally, this would imply that the rent extraction $\hat{e}_K - \hat{u}_K$ is of order η^{1+d} and contradict, for η small enough, the lower bound of $\underline{W}\eta$ which was mentioned above.

Deriving the belief bound for $\hat{\beta}_K$

The idea for proving the proposition is to transform (beyond recognition!) H ’s Bellman equation (3) into the following dynamic belief equation:

$$\frac{\beta_{n+2}}{\beta_{n+1}} \geq \frac{\beta_{n+1}}{\beta_n} - c\sqrt{\beta_{n+1}}. \quad (5)$$

This transformation requires a number of conceptual and technical steps, outlined below, and involves the observation that, for n large and η small, one can set without loss of generality

$$\frac{aw_n}{\eta D} = \frac{\beta_{n+1}}{\beta_n} \quad (6)$$

where $w_n = u_H(E_H(R_n)) - u_H(n)$ is the *rent extraction index*. This index w_n must converge to zero as n goes to infinity (Proposition 2): Intuitively, the probability β_n of facing H converges to

zero and R_n converges to an L -efficient contract as n goes to infinity. This implies that $E_H(R_n)$ gives roughly the same utility to H as R_n does for n large (see Figure 2) or, in other words, that w_n converges to zero. It turns out however (Lemma 9) that if $\hat{\beta}_K = \beta_{n_K}$ is too small (in particular, below the belief bound that we are trying to establish), then no solution to the dynamic equation (5) is such that β_{n+1}/β_n converges to zero. From (6), this implies that w_n cannot converge to zero as n goes to infinity, yielding the desired contradiction.

Transforming H 's Bellman equation into the dynamic belief equation (5)

The transformation of H 's Bellman equation is based on P 's IC constraint. One may consider, at each round n , P 's incentive to jump to $\{E_H(R_n), E_L(R_n)\}$ instead of pursuing active negotiations (this “jump deviation” is always available to P , as mentioned earlier). For this deviation to be suboptimal, we must have

$$w_n a \beta_n \geq \sum_{R_{n+1} \in (M_n \cup \{R_n\}) \cap \mathcal{H}} [\beta_n \mu_n^H(R_{n+1}) \eta D + (1 - \beta_n) \mu_n^L(R_{n+1}) \eta (Q(R_{n+1}) - Q(E_L(R_n)))] \quad (7)$$

where $\mu_n^\theta(R_{n+1})$ is type θ 's probability of accepting any given contract R_{n+1} of \mathcal{H} that lies in the menu $M_n \cup R_n$ of available contracts. The left-hand side is an upper bound on P 's gain, similar to the one used in the block construction described earlier. The right-hand side consists of the loss on the high and low types, respectively, for each possible chosen contract R_{n+1} . Rewriting (7) as

$$w_n a \beta_n \geq \sum_{R_{n+1} \in (M_n \cup \{R_n\}) \cap \mathcal{H}} \mu_n^L(R_{n+1}) \left[\beta_n \frac{\mu_n^H(R_{n+1})}{\mu_n^L(R_{n+1})} \eta D + (1 - \beta_n) \eta (Q(R_{n+1}) - Q(E_L(R_n))) \right], \quad (8)$$

the right-hand side may now be seen as a convex combination of these losses, weighted by L 's probability $\mu_n^L(R_{n+1})$ of choosing each contract R_{n+1} . Therefore, there must exist a contract R_{n+1} , chosen with positive probability, for which³⁴

$$w_n a \beta_n \geq \mu_n \beta_n \eta D + (1 - \beta_n) \eta (Q(R_{n+1}) - Q(E_L(R_n))), \quad (9)$$

where $\mu_n = \frac{\mu_n^H(R_{n+1})}{\mu_n^L(R_{n+1})}$. That condition is then broken up into two weaker conditions³⁵

$$\frac{w_n a}{\eta D} \geq \mu_n, \quad (IC_n^{LL})$$

and

$$\beta_n w_n a \geq (1 - \beta_n) \eta (Q(R_{n+1}) - Q(E_L(R_n))) \quad (IC_n^{LH})$$

³⁴Some “loss” terms in (8) may be negative, but very slightly so. The argument can be adapted to account for this, as shown in the appendix.

³⁵Again, some of the terms in the right-hand side of (9) may be negative, but only very slightly so. This complication is addressed in the appendix.

This step establishes the existence of a sequence $\{R_n\}$ of on-path contracts such that the coupled inequalities IC_n^{LL} and IC_n^{LH} are satisfied for all n large enough and for which the sequence $\{\beta_n\}$ of posterior beliefs satisfies in good approximation $\beta_{n+1} = \beta_n \mu_n$. To exploit this system of conditions and derive the desired contradiction, it is helpful to make it more tractable by replacing all inequalities IC_n^{LL} by equalities. This is achieved as follows: starting from some judiciously chosen round N , increase μ_N so as to satisfy IC_N^{LL} as an equality. We also increase β_{N+1} to preserve the relation $\beta_{N+1} \sim \beta_N \mu_N$. After these modifications of μ_N and β_{N+1} , we get $\frac{\beta_{N+1}}{\beta_N} \sim \mu_N = \frac{w_N a}{\eta D}$. Crucially, this change does not violate any inequality for indices $n \geq N$. One can then repeat increase μ_{N+1} and β_{N+2} , etc, while satisfying all the relevant inequalities. Moreover, this operation on beliefs and rejection probabilities has no impact on the *contracts* involved in the inequalities. Therefore, if one can show that these contracts violate some equilibrium property (in particular, asymptotic convergence to efficiency), this will establish a contradiction, irrespective of the fact that the beliefs were modified to establish that violation.

Subtracting $u_H(E_H(R_n))$ from both sides of H 's Bellman equation (3) and rearranging (also recalling that, by definition, $w_n = u_H(E_H(R_n)) - u_H(n)$) yields

$$w_{n+1} = w_n - \eta(u_H(E_H(R_n)) - u_H(R_{n+1})) + \eta w_{n+1} + (1 - \eta)(u_H(E_H(R_{n+1})) - u_H(E_H(R_n))). \quad (10)$$

The last two terms are shown in the appendix to be negligible, resulting in the simpler equation

$$w_{n+1} = w_n - \eta(u_H(E_H(R_n)) - u_H(R_{n+1})). \quad (11)$$

Finally, one can further show by exploiting IC_n^{LH} and some geometric inequalities (Lemma 12) that H 's utility difference in (11) is of order $\sqrt{\beta_{n+1}}$. Since w_n is proportional to $\eta \frac{\beta_{n+1}}{\beta_n}$, as explained above, it easily follows that (11) can then be turned into (5).

Bridging blocks and the dynamic belief equation

The previous paragraphs have naturally made abstraction from a number of complications arising in the actual proof. One intermediary step must be mentioned here, however, as it takes up almost a third of the actual proof (Part II in the Appendix). The dynamic belief equation and the non-convergence to zero that it implies only work as long as β_n remains small during all the rounds following block K . Unfortunately, however, one cannot guarantee that β_n remains small: even though, *on average*, L is more likely than H to choose contracts in \mathcal{H} (since H can also choose the H -efficient contract C_n), there may be contracts in the menu $M_n \cap \mathcal{H}$ that are chosen with much higher probability by H than L , resulting in a spike up of the posterior β_n . And, at least in principle, one cannot rule out the possibility that those contracts creating the spike are precisely the contracts chosen in P 's IC constraint (8) to get the simpler condition (9).

Fortunately, this issue does not arise if w_n is small enough (precisely, below $\frac{\eta D}{2a}$), because in that case μ_n can be shown to be smaller than 1, implying that β_n *must* be decreasing. But the last block K ended up with the condition $w_n \leq \eta \bar{W}$, and the constant \bar{W} could a priori be well above $\frac{D}{2a}$. To address this difficulty, the idea is to insert a number of intermediary blocks, following block K , and build a hybrid argument drawing on the arguments used along the blocks constructed earlier and the arguments used to derive the dynamic belief equation. Along those new blocks, w_n is guaranteed to decrease until it drops below $\frac{\eta D}{2a}$, but not *too* fast, as it must remain exactly of order η : remember that, in the above contradiction argument, we used the fact $\hat{w}_K = \hat{e}_K - \hat{u}_K$ was *greater* than $\eta \bar{W}$ for some fixed $\bar{W} > 0$. Moreover, these intermediary blocks are built so that β_n remains of order η^d . The dynamic belief equation argument can then be applied at the end of those intermediary blocks.

5 Discussion

Whether it concerns trade, production, employment, finance, or other economic activities, contract negotiation is central to economic analysis and has been widely studied from both pure and applied perspectives. However, much remains to understand when i) contracts are more complex than binary sales or divisions of a given “pie”, ii) some parties hold private information, and iii) negotiation is not limited to a single take-it-or-leave-it-offer. In these common situations, private information is endogenously revealed through negotiation even in the absence of any new exogenous information, and this affects ongoing negotiations, providing a dynamic interaction between beliefs and contractual agreements.

I have taken the view that parties should be able to react to this endogenous flow of information, particularly when it reveals some inefficiency of the current agreement, instead of being stuck with this agreement. This view captures the idea of unconstrained negotiation, and seems essential to generate a foundation for renegotiation-proof contracts.³⁶ While the analysis has focused on a specific negotiation protocol (just like foundations of the Nash bargaining solution and of the Coase conjecture), it would be useful to explore, in future research, more general protocols of negotiation. One may conjecture that contract negotiation should lead to efficient outcomes under these more general protocols as long as it is unconstrained in the above sense and does not entail additional frictions, such as explicit renegotiation costs.³⁷

In the present setting, where the efficiency gap converges to zero (for type L),³⁸ it is perhaps

³⁶For instance, Wang (1998) shows that when the first signed contract is always implemented, outcomes can be ex post efficient. In fact, those outcomes are, in his model, identical to the full commitment outcome.

³⁷See, e.g., Brennan and Watson (2013).

³⁸The gap for type H remains bounded away from zero but the probability of facing type H converges to zero.

surprising that efficiency obtains for all equilibria, including non-stationary ones, as frictions vanish. Future work should consider extensions of this result to more than two types or contract dimensions.³⁹ While those extensions do not seem straightforward, the techniques developed here should prove helpful to analyze them. With more than two types, one may conjecture that the protocol of this paper will lead to ex post efficiency, and that the limit contracts will also be easily determined by ex post efficiency.⁴⁰

Allowing shared bargaining power between the principal and the agent would also be natural.⁴¹ However, this extension would not provide a foundation for the renegotiation-proof contracts, common in the literature,⁴² which are *proposed by the principal*. Negotiation protocols in which both parties have some bargaining power should typically *not* yield the principal-optimal renegotiation-proof contracts analyzed in this literature.

Finally, the analysis has focused on a single “delivery” time at which the contract is implemented (an alternative interpretation, in which payoffs occur in each round, was discussed in the introduction). To provide a foundation for multi-period renegotiation-proof contracts, one should consider a more general model with multiple “physical” events (payments or efforts are made, new information arrives) and with a renegotiation protocol like the one analyzed here inserted between each consecutive events. Each renegotiation stage would pertain to continuation contracts pertaining to the remaining events. In such a model, each event arrival would trigger the end of the current negotiation. Alternatively, events could occur at integer dates and renegotiation rounds between event dates $t - 1$ and t would occur at times $\tau_n^t = t - \frac{1}{2^n}$, for $n \geq 1$. This double time scale is natural in many contexts where physical deliveries have a particular calendar structure (e.g., monthly wage, weekly delivery, quarterly report, etc.), but parties’ ability to negotiate has no reason to be thus constrained.⁴³

³⁹Another extension concerns the case of interdependent values. Deneckere and Liang (2006) provide an in-depth analysis of this case for the case of binary sales.

⁴⁰With three types (H, M, L) for instance, the contracts $\{E_\theta(R_0)\}_{\theta \in \{H, M, L\}}$ would be the cheapest θ -efficient contracts that are incentive compatible. The conjecture extends trivially to finitely many types and can also be extended to a differential equation characterizing $\{E_\theta(R_0)\}_{\theta \in \Theta}$ for a continuum of types, although proving the conjecture in that case seems particularly challenging.

⁴¹Ausubel and Deneckere (1992) analyze a model with shared bargaining power. In that model, the informed party (buyer) is getting all the surplus, even when the seller is making all the offers. When given the opportunity to make offers, therefore, the buyer does not lose anything by remaining silent. In the present setting however, the agent does not get all the surplus and it is clear that, even in the absence of private information, sharing bargaining power would affect the equilibrium allocation of surplus. The key question is whether shared bargaining power would create some efficiency loss.

⁴²See, e.g., Dewatripont (1989), Maskin and Tirole (1992), and Battaglini (2007), in addition to the work already cited.

⁴³The time structure described here allows unrestricted renegotiation between physical events, as there are infinitely

Appendices

A Proof of Theorem 1 (Existence of a PBE)

It suffices to prove the result for R_0 in the H -Rent configuration and $\beta_0 \in (0, 1)$: the degenerate prior and No-Rent cases obtain as direct consequences of Proposition 1 below, which proves (independently of this section) the existence and uniqueness of an equilibrium in those cases, while the L -Rent case obtains by symmetry of the H -Rent case. The proof proceeds in two steps:

- Step 1 - Prove the existence of an equilibrium in an auxiliary game played between P and H .
- Step 2 - Construct a strategy profile of the original game based on the equilibrium established in Step 1, and verify that it defines a PBE of the original game.

Step 1: Auxiliary game

The game starts with a contract $R_0 \in \mathcal{H}$ in the H -Rent configuration and a parameter $\beta \in (0, 1)$. For this auxiliary game, β is just a parameter affecting the payoff functions and is devoid of its interpretation as a belief.

The auxiliary game is a dynamic game with infinitely many rounds. At each round n , starting in state R_n , P proposes new contracts $R_{n+1} \in \mathcal{H}$ and $C_n \in \mathcal{E}_H$ subject to the constraints

$$u_L(R_{n+1}) \geq u_L(R_n) \quad (12)$$

$$u_L(R_{n+1}) \geq u_L(C_n) \quad (13)$$

$$u_H(C_n) \geq u_H(R_n). \quad (14)$$

H then chooses a number $\mu_n \in [0, 1]$. The interpretation of this choice is that H accepts R_{n+1} with probability μ_n and C_n with probability $(1 - \mu_n)$. For this auxiliary game, however, μ_n is simply an action deterministically affecting payoffs.

The principal's cost, for strategies $\{R_{n+1}, C_n\}$ and $\{\mu_n\}$, is given by

$$\begin{aligned} \mathcal{Q}(\{R_{n+1}, C_n\}, \{\mu_n\}) = & \sum_{n \geq 0} Q(C_n) \beta (1 - \eta)^n (1 - \mu_n) \prod_{k=0}^{n-1} \mu_k \\ & + \sum_{n \geq 0} Q(R_{n+1}) \left(\beta (1 - \eta)^n \eta \prod_{k=0}^n \mu_k + (1 - \beta) (1 - \eta)^n \eta \right), \end{aligned} \quad (15)$$

while H 's payoff is

$$\mathcal{V}(\{R_{n+1}, C_n\}, \{\mu_n\}) = \sum_{n \geq 0} u_H(C_n) (1 - \eta)^n (1 - \mu_n) \prod_{k=0}^{n-1} \mu_k + \sum_{n \geq 0} u_H(R_{n+1}) (1 - \eta)^n \eta \prod_{k=0}^n \mu_k. \quad (16)$$

These payoffs correspond to the expected cost and utility that P and H would obtain in an equilibrium of the *original* game in which P proposes two contracts at each round, the breakdown probability is η , $\{\mu_n\}$ is the mixing strategy of H , L always accepts R_{n+1} , and the initial probability of facing H is equal to β .

many renegotiation rounds, and could in principle also be captured by a continuous-time model, in which the principal can propose new contracts at any instant between physical events occurring on a discrete time set.

LEMMA 1 *For any initial R_0 and $\beta \in (0, 1)$, there exists a perfect equilibrium of the auxiliary game*

Proof. The result is direct consequence of Theorem 1 in Harris (1985). We check Assumptions 1–5 of that theorem. The payoff function of the principal is simply the negative of his cost, Q . P’s (unconstrained) action set in round n is $S_{P_n} = \mathcal{H} \times \mathcal{E}_H$, while H ’s action space is $S_{H_n} = [0, 1]$ which are both compact and Hausdorff spaces. Hence, Assumptions 1 and 2 are satisfied. P’s feasible set at each round n , as defined by the constraints (12) and (14), is closed and depends continuously on the current state. Therefore, the set \mathcal{S}_f of feasible sequences is closed in $\mathcal{S} = \times_n (S_{P_n} \times S_{H_n})$ endowed with the product topology, and the set of feasible actions in round n depends continuously on past play. Thus, Assumptions 3 and 4 are satisfied. Finally, the payoffs $-Q$ and V are clearly continuous on their domain \mathcal{S}_f , so Assumption 5 is satisfied as well. The result follows. ■

REMARK 1 *We can similarly define an auxiliary game and equilibrium when instead R_0 is in the L-Rent configuration. This equilibrium yields an expected utility for H , as a passive player of the auxiliary game, given, by*

$$V_H(\beta) = \sum_{n \geq 0} (1 - \eta)^n \eta u_H(R_{n+1}). \quad (17)$$

This equilibrium and payoff is used to define H ’s strategy, off the equilibrium path, in the PBE construction for the original game.

Step 2: Equilibrium of the original game

Starting from $R_0 \in \mathcal{H}$ and a belief $\beta_0 \in (0, 1)$, the equilibrium strategies are defined as follows:

At each round n :

- P proposes the sequence of contracts $\{C_n, R_{n+1}\}$ corresponding to the auxiliary game started at (R_0, β_0)
- L accepts R_{n+1} with probability 1, while H accepts R_{n+1} with probability μ_n and C_n with probability $(1 - \mu_n)$, where μ_n is H ’s equilibrium choice in the auxiliary game. If $R_{n+1} \neq R_n$ and the agent accepts R_n , P assigns probability 1 to H , so the continuation play is trivially defined in that case, by Proposition 1 (whose proof is independent of Theorem 1).
- If P proposes, at some round n , a menu M_n that does *not* correspond to the pair of contracts defined by the auxiliary game, let \bar{R}_{n+1} denote the contract of $M_n \cup \{R_n\}$ that maximizes L ’s utility and \bar{C}_n denote the contract of $M_n \cup \{R_n\}$ that maximizes H ’s utility.⁴⁴ By construction, \bar{R}_{n+1} and \bar{C}_n satisfy (12)–(14). Let \hat{R}_{n+1} denote the L -efficient contract that gives L the same utility as \bar{R}_{n+1} and \hat{C}_n denote the H -efficient contract that gives H the same utility as \bar{C}_n . There are three cases to consider: a) $u_H(\hat{C}_n) \geq u_H(\hat{R}_{n+1})$ and $u_L(\hat{R}_{n+1}) \geq u_L(\hat{C}_n)$, b) $u_H(\hat{C}_n) < u_H(\hat{R}_{n+1})$ and $u_L(\hat{R}_{n+1}) \geq u_L(\hat{C}_n)$, and c) $u_H(\hat{C}_n) \geq u_H(\hat{R}_{n+1})$ and $u_L(\hat{R}_{n+1}) < u_L(\hat{C}_n)$. Because of the single-crossing property, the fourth and last logical case cannot occur, as is easily checked.

Continuation play is then defined as follows, according to each case.

⁴⁴If there are several maximizers, the equilibrium selects any of them. There must be at least one maximizer, because the menu is finite.

- a) L chooses \bar{R}_{n+1} with probability 1, H chooses \bar{C}_n with probability 1. If the agent chooses any other contract R in $M_n \cup \{R_n\}$, then the principal assigns probability 1 to a type θ of the agent such that the other type $\theta' \neq \theta$ cannot benefit from choosing that contract if the principal put probability 1 on θ .⁴⁵ There always exists at least one such type, as is easily checked.
- b) Case b) can occur only if \bar{R}_{n+1} is in the H -Rent configuration. Continuation play is defined by the continuation equilibrium of the auxiliary game in which, following R_n , P proposes \bar{R}_{n+1} and \hat{C}_n , but replacing \hat{C}_n by \bar{C}_n .⁴⁶ In particular, L accepts \bar{R}_{n+1} with probability 1, and H randomizes between the contracts \bar{C}_n and \bar{R}_{n+1} according to the probability μ_n of coming from the auxiliary equilibrium if \bar{C}_n is replaced by \hat{C}_n . If the agent picks any contract other than \bar{R}_{n+1} and \bar{C}_n , P assigns probability 1 to one type according to the same rule as in Case a). Continuation for rounds $m \geq n+1$ is also determined by the equilibrium of the auxiliary game.
- c) Case c) is symmetric to Case b), and can only occur if \bar{C}_n is in the L -Rent configuration. The continuation equilibrium is defined by the continuation equilibrium, from period 1 onwards (see Remark 1) of the auxiliary game starting in period 0 with belief $\tilde{\beta}_0 = 1 - \beta_n$ (since L now plays the role of H and vice versa) and at some fictitious contract \tilde{R} in the L -Rent configuration such that \hat{R}_{n+1} and \bar{C}_n satisfy equations (14) and (12), respectively (the inequalities are reversed, because the equilibrium is in the L -Rent configuration).

This construction defines the continuation strategies after any possible history. We now verify that the strategy profile forms an equilibrium, using the one-shot deviation principle, which applies since the breakdown probability η has the effect of discounting the utility of future rounds at a geometric rate.

Consider first L 's strategy, assuming that P follows the prescribed sequence of contracts. From (13), L cannot benefit from picking C_n : indeed, doing so causes β_n jumps to 1, and L to be stuck with utility $u_L(C_n)$, which is (weakly) lower than $u_L(R_{n+1})$ and hence lower than his continuation utility if he chooses R_{n+1} .⁴⁷ Similarly, if $R_{n+1} \neq R_n$ and L chooses R_n , then β_n jumps to 1, and L 's continuation utility is bounded above⁴⁸ by $u_L(R_n)$, which is weakly dominated by accepting R_{n+1} by (12) (guaranteeing that L 's continuation utility is bounded below by $u_L(R_{n+1})$).

Let us now consider the optimality of H 's strategy. From (14), $u_H(C_n) \geq u_H(R_n)$. Therefore, if H holds on to R_n , his continuation utility is equal to $u_H(R_n)$, which is weakly dominated by taking C_n . Moreover, given that H randomizes between C_n and R_{n+1} , his expected payoff is given by (16), and by perfection of the auxiliary equilibrium, the strategy $\{\mu_n\}$ is a best response to the sequence of contracts.

Consider now the agent's strategy after a deviation by P . In Case a), if L chooses \bar{C}_n , his utility is bounded above by $\max\{u_L(\bar{C}_n), u_L(\hat{C}_n)\}$, which is less than $u_L(\bar{R}_{n+1})$, by definition of Case a). Similarly, if L picks any other contract R , then either P puts probability 1 on L , in which case L gets utility $u_L(R)$, which is less than $u_L(\bar{R}_{n+1})$, by definition of \bar{R}_{n+1} , or P puts probability 1 on H , but in this case L cannot

⁴⁵That is, θ' prefers the contract that he is supposed to take with probability 1 in equilibrium (e.g., \bar{C}_n if $\theta' = H$) to the θ -efficient contract that gives θ the same utility as R .

⁴⁶By construction, $\bar{R}_{n+1} \in \mathcal{H}$, \bar{C}_n is H -efficient, and the contracts satisfy conditions (12), (13), and (14), so the contract pair is feasible for the auxiliary games.

⁴⁷That utility is always weakly higher than $u_L(R_{n+1})$, since L can always hold on to R_{n+1} .

⁴⁸Indeed, P then proposes the H -efficient contract R that gives H utility $u_H(R_n)$, and $u_L(R) \leq u_L(R_n)$ by the single-crossing property and the fact that R_n is in the H -Rent configuration.

benefit from this erroneous belief. The same reasoning applies to H : it is optimal for that type to choose \bar{C}_n .

In Case b), L prefers \bar{R}_{n+1} over any other contract in $M_n \cup \{R_n\}$, by an argument similar to Case a). Now consider H 's response to P 's deviation. First, H cannot benefit from choosing a contract R other \bar{C}_n and \bar{R}_{n+1} , for the reason explained in Case a). Moreover, given the continuation play, which is defined by the auxiliary equilibrium, randomizing according to the probability μ_n coming from the auxiliary equilibrium in which \bar{R}_{n+1} and \hat{C}_n are proposed is optimal.⁴⁹

Case c) is similar to Case b).

There remains to show that P 's strategy is optimal. By construction of the auxiliary equilibrium, P 's strategy is optimal among all strategies that propose contracts (R_{n+1}, C_n) satisfying (12), (13), and (14). As shown by Lemma 3 (whose proof is independent of Theorem 1), P can never benefit from any deviation in which L accepts a contract that is not in the H -Rent configuration. Moreover, any contract R accepted by H with positive probability and that is *not* in the H -Rent configuration immediately results, at the next round, in an H -efficient contract that gives H the same utility as R and is less costly to P than R . We can therefore, without loss of generality, consider deviations in which P proposes one H -efficient contract, \bar{C}_n , and a number of contracts in the H -Rent configuration, among which \bar{R}_{n+1} maximizes L 's utility, and such that $u_L(\bar{C}_n) \leq u_L(R_{n+1})$. Given the agent's strategy, the menu is equivalent to just proposing \bar{C}_n and \bar{R}_{n+1} , which is a feasible strategy in the auxiliary equilibrium and thus has to be weakly dominated by the equilibrium menu, by subgame perfection of that menu in the auxiliary game.

B Results holding for all friction levels

B.1 Statements

PROPOSITION 1 *The following holds for any PBE and η :*

- i) *If the prior β puts probability 1 on some type θ , P immediately proposes the θ -efficient contract that leaves θ 's utility unchanged and θ accepts it.*
- ii) *If R_0 is θ -efficient, P immediately proposes $E_{\theta'}(R_0)$ ($\theta' \neq \theta$), and θ' accepts it.*
- iii) *If R_0 is in the No-Rent configuration, P immediately proposes $E_L(R_0)$ and $E_H(R_0)$, and each type θ accepts $E_\theta(R_0)$.*
- iv) *If R_0 is in the H -Rent (L -Rent) configuration, H 's (L 's) expected utility is bounded above by $u_H(E_H(R_0))$ ($u_L(E_L(R_0))$).*

The next result is crucial for the analysis: for any PBE and round n , P can always propose the contracts $E_H(R_n)$ and $E_L(R_n)$ and have them accepted by types H and L , respectively. This deviation puts an upper bound on P 's continuation cost as a function of the current contract R_n . The deviation will henceforth simply be referred to as the “jump.”

Let β_n denote the probability, at the beginning of round n , that P assigns to type H .

⁴⁹Because $u_H(\hat{C}_n) = u_H(\bar{C}_n)$, H gets exactly the same utility as in the auxiliary equilibrium, even though the contract \bar{C}_n is not in the H -Rent configuration.

LEMMA 2 (JUMP) *If R_n is in the H -Rent configuration and P proposes the contracts $E_H(R_n)$ and $E_L(R_n)$, with $E_H(R_n)$ augmented by an arbitrarily small amount $\varepsilon > 0$, then H accepts $E_H(R_n)$ with probability 1 and L accepts $E_L(R_n)$ with probability 1. Therefore, P 's continuation cost is bounded above by $\bar{Q}_n = \beta_n Q(E_H(R_n)) + (1 - \beta_n) Q(E_L(R_n))$*

Proof. The result follows from Part iv) of Proposition 1: $E_H(R_n)$ plus any small amount gives a strictly higher utility to H than what he can get under any continuation utility, and also gives him strictly more utility than $E_L(R_n)$. Therefore, H accepts the contract with probability 1. Because L strictly prefers $E_L(R_n)$ to $E_H(R_n)$ in the H -Rent configuration, and because the agent's type is revealed in round n unless L takes the strictly suboptimal contract $E_H(R_n)$, it is optimal for L to accept $E_L(R_n)$. ■

LEMMA 3 *If R_0 is in the H -Rent configuration, then in any PBE, L accepts only contracts that are in the H -Rent configuration.*

Given any PBE, any contract sequence $\{R_n\}$ that is accepted by L with positive probability (until the exogenous negotiation breakdown) will be called a **choice sequence**. When R_0 is in the H -Rent configuration, choice sequences will play a particular role: we will see that, without loss, any accepted contract sequence is a choice sequence until H accepts an H -efficient contract. Moreover, choice sequences have several important properties. First, as indicated by Lemma 3, any choice sequence consists of contracts that are in the H -Rent configuration. Other properties are described by the following proposition.

PROPOSITION 2 *Suppose that R_0 is in the H -Rent configuration. Along any choice sequence $\{R_n\}$, i) β_n converges to zero, ii) R_n converges to an L -efficient contract, denoted \bar{C}_L , and iii) the quantity $w_n = u_H(E_H(R_n)) - u_H(n)$ (where we recall that $u_H(n)$ is H 's continuation utility at the beginning of round n) converges to zero as n goes to infinity.*

B.2 Proofs

PROOF OF PROPOSITION 1

Part i) Let \bar{u} denote the agent's supremum over his expected utility, given his type θ , over all possible continuation PBEs starting from R_0 at which P puts probability 1 on type θ , and let $u = u_\theta(R_0)$. Suppose by contradiction that $\bar{u} > u$. By time homogeneity, \bar{u} will be the same in the next round if the agent rejects new offers from P in round 0 and P continues to assign probability 1 on facing type θ . In such case, the agent's continuation payoff is bounded above by $\tilde{u} = \eta u + (1 - \eta)\bar{u} < \bar{u}$. Consider any PBE that gives θ an expected utility $u_0 \in (\tilde{u}, \bar{u})$ (such PBE must exist, by definition of \bar{u}). Suppose that the principal deviates by proposing the θ -efficient contract C that give θ a utility level u' in (\tilde{u}, u_0) . By definition of a PBE,⁵⁰ P continues to assign probability 1 to type θ after his *own* deviation. If the agent accepts C with probability 1, the deviation is strictly profitable to P since C is the cheapest way of providing utility $u' < u_0$ to the agent. If the agent rejects the offer with positive probability, then by Bayes rule, P must continue to assign probability 1 to type θ , which implies that his continuation utility is bounded above by \bar{u} . Therefore, the agent's rejection is strictly suboptimal, implying that the agent must accept C with probability 1 and

⁵⁰See Fudenberg and Tirole (1991), part iii) of the definition.

the deviation is profitable.⁵¹ Let \underline{Q} denote the cost of the θ -efficient contract, \underline{C} , that provides utility u to θ . Clearly, any PBE must cost exactly \underline{Q} , otherwise P has a profitable deviation which is to propose the θ -efficient contract that gives θ slightly more than u and costs less than following the PBE. Moreover, the only way of achieving \underline{Q} is to propose \underline{C} in the first round and have it accepted with probability one.

Part ii) Suppose without loss that $\theta = L$ (the opposite case is treated identically). Let $u_L = u_L(R_0)$ and $u_H = u_H(R_0)$. Also let $\bar{u}_H(\beta)$ denote the supremum utility that H can achieve over any continuation PBE starting from R_0 when P assigns probability β to H , and let $\bar{u}_H = \sup_{\beta \in [0,1]} \bar{u}_H(\beta)$. Suppose by contradiction that $\bar{u}_H > u_H$. Then, for any small $\varepsilon > 0$, there exists $\bar{\beta}$ and an associated PBE for which H 's continuation utility is above $\bar{u}_H - \varepsilon > u_H$. For that PBE, because L gets at least u_L and C is L -efficient, $\bar{Q}_L \geq Q$, where $Q = Q(R_0)$, and \bar{Q}_L is P 's expected cost in that PBE conditional on facing θ_L . Since not proposing any new contract is always feasible for P , and costs Q , the continuation cost \bar{Q}_H conditional on facing H must satisfy $\bar{Q}_H \leq Q$ to offset the weakly higher cost conditional on facing L . Suppose that P deviates from that PBE by proposing the H -efficient contract that gives θ_H utility $\bar{u}_H - \varepsilon - \epsilon$, for arbitrarily small ϵ . Because, for small enough ε and ϵ , $\bar{u}_H - \varepsilon - \epsilon > \eta u_H + (1 - \eta)\bar{u}_H$, H accepts this proposal with probability 1. For any strategy that θ_L chooses and continuation equilibrium, this proposal strictly reduces P 's expected cost (since $\bar{Q}_H \leq Q$), yielding a contradiction. This shows that $\bar{u}_H(\beta) = u_H$ for all β .⁵² To conclude, suppose that P proposes the H -efficient contract that gives H utility $u_H + \tilde{\epsilon}$, for $\tilde{\epsilon}$ arbitrarily small. From the previous observation, H must accept that contract regardless of L 's strategy. This shows that P can and, hence, does achieve the full-commitment optimal cost under any PBE. This proves Part ii).

Part iii) Suppose without loss that $Q_L \geq Q_H$, where $Q_\theta = Q(E_\theta(R_0))$ (the opposite case is proved symmetrically). Let \bar{Q} denote the maximal expected cost incurred by P over all PBEs and beliefs $\beta \in [0, 1]$, starting from R_0 .

We start by showing that $\bar{Q} \leq Q_L$. Suppose by contradiction that $\bar{Q} > Q_L$ and consider any PBE that achieves \bar{Q} .⁵³ Now suppose that P deviates by proposing the pair \tilde{C}_L, \tilde{C}_H of contracts such that \tilde{C}_θ is efficient for θ and costs $\bar{Q} - \varepsilon$ for some ε arbitrarily small compared to η . Those contracts maximize each type's utility subject to costing P at most $\bar{Q} - \varepsilon$. Because these contracts are efficient and incentive compatible, Part ii) guarantees that no type ever chooses the contract meant for the other type. Moreover, no matter what belief and continuation PBE follows rejection of these contracts, P 's continuation cost must be less than \bar{Q} , by definition of \bar{Q} . But this latter bound implies that there must be at least one type θ of the agent who is getting a lower payoff if he rejects \tilde{C}_θ than if he accepts it: conditional on rejection P has to be spending weakly less on at least one type of the agent than under \tilde{C}_θ (up to ε , which is negligible compared to η). Moreover, the contract \tilde{C}_θ maximizes that type's utility subject to P spending less than \tilde{C}_θ . Since rejection leads to a renegotiation breakdown with a probability η , which gives that type a strictly lower utility than \tilde{C}_θ , accepting \tilde{C}_θ is strictly more profitable than rejection for that type, and thus he accepts \tilde{C}_θ with probability 1. As a result, a rejection fully reveals that the agent is of the other type. From Part i), that agent gets $u_\theta(C)$ after rejection, which is strictly less than the utility he gets from \tilde{C}_θ (since that

⁵¹The continuation play after P 's deviation must be a PBE of the corresponding continuation game. Therefore, if θ 's continuation strategy, after P 's deviation, is to reject the proposed deviation with positive probability, Bayes rule applies. I am grateful to Marcin Peski for proposing the current version of this argument.

⁵²If $\beta = 0$, P does not propose anything new, from i) and L -efficiency of R_0 , and the result trivially holds in that case too.

⁵³If the supremum \bar{Q} is not achieved, the argument below can easily be adapted by considering a PBE whose expected cost is arbitrarily close to \bar{Q} .

contract maximizes the agent's utility subject to a higher cost than what P incurs with C). Therefore, both types accept their contract, and this reduces the cost of the principal strictly below \bar{Q} , showing that this is a profitable deviation. Thus, necessarily, $\bar{Q} \leq Q_L$.

Since L cannot get utility less than $u_L(R_0)$, under any PBE, and Q_L is the cheapest cost of providing that utility, this means that in all PBEs starting with $\beta \in (0, 1)$, P must spend weakly less than \bar{Q} on the high type, in order to guarantee that $\bar{Q} \leq Q_L$. Let \bar{u}_H denote the supremum expected utility that H gets over all PBEs and beliefs $\beta > 0$. Since P spends less than Q_L on H , \bar{u}_H is bounded by the utility \hat{u}_H obtained from the H -efficient contract \hat{C}_H that costs Q_L . We will show that $\bar{u}_H = u_H(E_H(R_0))$. Suppose by contradiction that $\bar{u}_H > u_H(E_H(R_0))$, and consider a PBE that achieves \bar{u}_H .⁵⁴ The expected cost Q from that PBE must be above $\beta Q(\bar{C}_H) + (1 - \beta)Q_L$, where \bar{C}_H is the H -efficient contract that gives utility \bar{u}_H to H . Suppose that P deviates by proposing the contracts \tilde{C}_L, \tilde{C}_H such that \tilde{C}_L is L -efficient and gives utility $u_L(C) + \varepsilon^2$ to L and \tilde{C}_H is H -efficient and gives utility $\bar{u}_H - \varepsilon$ to H , for ε small compared to η . H accepts \tilde{C}_H , since rejection leads to a continuation utility bounded above by \bar{u}_H and to a strictly lower payoff in case of a breakdown. Given that, L also accepts, since rejection will reveal his type, and, by Part i), result in a utility of $u_L(E_L(R_0))$. The cost reduction on the high type is of order ε compared to $Q(\bar{C}_H)$, while the cost increase on the low type is of order ε^2 , compared to Q_L . Therefore, this deviation is strictly profitable for ε small enough. This shows that $\bar{u}_H = u_H(E_H(R_0))$. Proceeding as in the end of the proof of Part i), this shows that L 's maximal utility across all PBEs for $\beta \in (0, 1)$ is $u_L(E_L(R_0))$.

Part iv) The argument is similar to the proof of Part iii). Let \bar{Q} denote P's maximal expected cost over all PBEs and beliefs, starting from R_0 . We will start by showing that $\bar{Q} \leq Q(E_L)$, where $E_L = E_L(R_0)$. Suppose by contradiction that \bar{Q} is strictly greater than $Q(E_L)$ and achieved for some PBE and belief,⁵⁵ and consider the following deviation: P proposes the contracts \tilde{C}_θ that are efficient for each type and cost $\bar{Q} - \varepsilon$ for ε arbitrarily small. It is easily shown that these contracts are IC, and by a similar argument as in Part iii), rejecting those contracts is always a strictly dominated strategy for one of the two types, and hence for both types. This is a strictly profitable deviation for P, yielding a contradiction. Hence, $\bar{Q} \leq Q(E_L)$. Since L gets an expected utility of at least $u_L(R_0)$ in all PBEs, and providing that utility costs at least $Q_L = Q(E_L)$ to P, this means that P spends at most Q_L on H , in all PBEs, and for all initial beliefs $\beta > 0$. This implies that H 's expected utility is bounded above by the utility it achieves with the H -efficient contract that costs Q_L . We now show that H 's expected utility is bounded above by $u_H(E_L)$. Suppose not, and consider a PBE that gives H its highest utility, across PBEs and beliefs, denoted $\bar{u}_H > u_H(E_L)$. The expected cost Q from that PBE must be above $\beta Q(\bar{C}_H) + (1 - \beta)Q_L$, where \bar{C}_H is the H -efficient contract that gives utility \bar{u}_H to H . Suppose that P deviates by proposing the contracts \tilde{C}_L, \tilde{C}_H such that \tilde{C}_L is L -efficient and gives utility $u_L(C) + \varepsilon^2$ to L , and \tilde{C}_H is H -efficient and gives utility $\bar{u}_H - \varepsilon$ to H , for ε arbitrarily small. Because \tilde{C}_H gives strictly more to H than \bar{u}_H , H will accept \tilde{C}_H and, hence L will accept \tilde{C}_L . Repeating the proof of Part iii), one can show that this deviation is strictly profitable, establishing the desired contradiction. The only difference with that earlier proof lies in showing that the proposed contracts are incentive compatible. This is indeed true, for ε small enough, because $\bar{u}_H > u_H(E_L)$ so H does not want to mimic L .⁵⁶

PROOF OF LEMMA 3

⁵⁴Again, the proof is easily adapted if the supremum is not achieved, by considering a PBE that gets very close to providing \bar{u}_H .

⁵⁵As before, one can use a PBE that yields a cost arbitrarily close to \bar{Q} , in case it is not exactly achieved.

⁵⁶It is straightforward to show that L does not want to mimic H , since P spends less on H than on L , and L is already getting his maximal utility given the cost that P incurs conditional on facing L .

Consider any PBE starting with R_0 in the H -Rent configuration. Consider, by contradiction, the first round n such that i) R_n is the H -Rent configuration and ii) L accepts with positive probability a contract R_{n+1} that is in a different configuration. Suppose first that R_{n+1} is in the No-Rent configuration. Then $u_L(n) = u_L(R_{n+1})$, by Part iii) of Proposition 1. This immediately implies that $u_L(R_n) \leq u_L(R_{n+1})$: R_{n+1} is on a weakly higher iso-utility curve of u_L than R_n . Moreover, because H can always accept R_{n+1} , $u_H(n) \geq u_H(R_{n+1}) > u_H(E_H(R_n))$, where the strict inequality comes from the fact that u_H is increasing along the iso-utility curve of u_L in the direction of \mathcal{E}_H .⁵⁷ This implies that the continuation cost for P is strictly above $\beta_n Q(E_H(R_n)) + (1 - \beta_n)Q(E_L(R_n))$, which contradicts Lemma 2. Now suppose that R_{n+1} is in the L -Rent configuration. Part iv) of Proposition 1 applied to the L -Rent configuration implies that, by choosing R_{n+1} , L gets a continuation utility of at most $u_L(\tilde{E}_L(R_{n+1}))$ where $\tilde{E}_L(\tilde{R})$ is defined – when \tilde{R} is in the L -Rent configuration – similarly to $E_H(R)$ when R is in the H -Rent configuration. Therefore, $u_L(\tilde{E}_L(R_{n+1}))$ must be weakly greater than $u_L(R_n)$. However, notice that when \tilde{E}_L is constructed, we use L 's isoutility curve between the efficiency curves \mathcal{E}_H and \mathcal{E}_L , which is steeper than H 's isoutility curve at $\tilde{E}_L(R_{n+1})$, from the single-crossing property. As can be easily seen visually, this implies that $u_H(R_{n+1})$ must have been *strictly* greater than $u_H(E_H(R_n))$, contradicting Part iv) of Proposition 1 applied to H . ■

PROOF OF PROPOSITION 2

i) Observe, first, that negotiation cannot end endogenously at a finite round N in the sense that $\beta_n = \beta_N > 0$ and $R_n = R_N \in \mathcal{H}$ for all $n \geq N$. If that were the case, P could strictly reduce his cost at round N by proposing the H -efficient contract $E_H(R_N)$ and have it accepted by H with probability 1, by Part iv) of Proposition 1. Hence, consider the case in which P keeps proposing new contracts until renegotiation is exogenously interrupted, and suppose by contradiction that there is a choice sequence with an associated belief subsequence $\{\beta_{n(k)}\}_{k \in \mathbb{N}}$ that converges to $\beta^* > 0$. Let $u_H^* = \sup\{u_H(R_n)\}$ where the supremum is taken among all contracts in the choice sequence. For H to accept R_n with positive probability infinitely often, $u_H(R_n)$ must converge to u_H^* for any subsequence, including along the subsequence $\{n(k)\}$.⁵⁸ However, that implies that proposing the H -efficient contract C_H that gives u_H^* to H is a strictly profitable deviation as $\beta_{n(k)}$ gets arbitrarily close to β^* : it does not change P's cost conditional on facing L but it strictly reduces P's expected cost by an amount arbitrarily close to $\beta^*[Q(C_L) - Q(C_H)]$, where C_θ is the θ -efficient contract that provides H with utility u_H^* .⁵⁹

ii) Suppose that there exists $\varepsilon > 0$ and a subsequence of rounds, indexed by m , for which $Q(R_m) - Q(E_L(m)) \geq \varepsilon$. For m large enough, β_m converges to zero, from part i), and is thus bounded above by $\frac{\eta\varepsilon}{2\Delta_Q}$, where $\Delta_Q = \max_{C \in \mathcal{C}} Q(C) - \min_{C \in \mathcal{C}} Q(C)$. Therefore, P can deviate by proposing $E_L(m)$, $E_H(R_m)$, which are respectively accepted by L and H . This deviation yields an immediate gain of $\eta\varepsilon$ on L and a loss of at most $\frac{\eta\varepsilon}{2}$ on H , given the upper bound on β_m , and is thus strictly profitable. This shows that the limit points of $\{R_n\}$ are all L -efficient. Let $u_L^* = \sup\{u_L(R_n)\}$. There is a subsequence \tilde{m} for which $u_L(R_{\tilde{m}})$

⁵⁷More explicitly, we have $u_H(R_{n+1}) > u_H(E_L(R_{n+1})) \geq u_H(E_L(R_n)) = u_H(E_H(R_n))$.

⁵⁸Otherwise, there must exist a subsequence of rounds for which $u_H(R_{m+1})$ is bounded above away from u_H^* by some constant $\delta > 0$. However, H 's continuation utility, $u_H(m)$, is nondecreasing and becomes arbitrarily close to u_H^* . (Monotonicity comes from the fact that H can always hold on to the last accepted contract and is proved formally in Lemma 5.) When H 's continuation gets within $\varepsilon\eta$ of u_H^* for some ε arbitrarily small, this implies that accepting R_{m+1} causes a loss of order $\eta\delta$, due to the probability of an immediate breakdown, and contradicts the fact that $u_H(m)$ is within $\varepsilon\eta$ of u_H^* .

⁵⁹Since $u_H(R_n)$ gets arbitrarily close to u_H^* and R_n lies in \mathcal{H} , $Q(R_n)$ becomes arbitrarily close to (or above) $Q(C_L)$ as n gets large.

converges to u_L^* . Moreover, since L can always hold on to any contract R_n along the choice sequence, and thus in particular for contracts occurring along the subsequence $\{\tilde{m}\}$, $u_L(R_n)$ must converge to u_L^* for all subsequences. Combining these observations, $\{R_n\}$ must converge to the L -efficient contract \bar{C}_L such that $u_L(\bar{C}_L) = u_L^*$. ■

iii) Parts i) and ii) have shown that β_n converges to zero and R_n converges to an L -efficient contract as n goes to infinity. This implies that $E_H(R_n)$ gives asymptotically the same utility to H as R_n does for n large (see Figure 2) and, hence, that w_n converges to zero.⁶⁰

C Proof of Theorem 2

Without loss of generality, it suffices to prove the theorem when R_0 is in the H -Rent configuration: Proposition 1 already addresses the case in which R_0 is in the No-Rent configuration, and the L -Rent configuration can be proved by symmetry. Let us thus assume that $R_0 \in \mathcal{H}$. From Lemma 3, L accepts only contracts in \mathcal{H} . Moreover, any contract C_n that is accepted only by H in equilibrium can be replaced by an H -efficient contract \tilde{C}_n that gives H the same utility, without affecting anyone's incentive (assuming that H accepts \tilde{C}_n with the same probability as the one with which he was accepting C_n in the initial equilibrium) and reduces P 's cost. Therefore, we can without loss of generality focus on PBEs in which P only proposes, at each round, a number of contracts in \mathcal{H} , and the H -efficient contract that gives H his continuation utility. This assumption is maintained throughout the analysis.

ORGANIZATION OF THE PROOF

The proof of Theorem 2 (Statement A) proceeds by contradiction. We will suppose that there exists $\varepsilon > 0$, a decreasing sequence $\{\eta_m\}_{m \in \mathbb{N}}$ of breakdown probabilities converging to zero, and a PBE associated to each η_m for which H 's expected utility $u_H(0)$ at round 0 is below $u_H(E_H(R_0)) - \varepsilon$. Throughout, $u_\theta(n)$ will denote θ 's continuation utility at the beginning of round n .

In what follows, we focus entirely on that sequence of η 's and corresponding PBEs. The expression “as η goes to zero” will refer to the elements of that sequence and corresponding PBEs.⁶¹

The difference $w_0 = u_H(E_H(R_0)) - u_H(0)$ can be thought of as a *rent extraction index* for type H . It defines how much rent P is extracting from H , relative to the immediate jump: $u_H(0)$ is H 's continuation utility while $u_H(E_H(R_0))$ is the maximal utility that P can give H in any equilibrium, as shown by Proposition 1, part iv).

The proof consists of the following steps.

Step 1: For each PBE of the sequence, show that one can construct a choice sequence ending at some finite round \tilde{N} for which the *augmented* rent extraction index,

$$\bar{w}_{\tilde{N}} = \max_{m \leq \tilde{N}} \{u_H(E_H(\tilde{m}))\} - u_H(\tilde{N}),$$

is of order η , and there exists $d > 0$ such that either

- a) $\beta_{\tilde{N}} \geq \eta^d$ and $w_0 \leq \hat{w}\sqrt{\eta}$, where $\hat{w} > 0$ is exogenous, or

⁶⁰Put differently, the contract \bar{C}_H defined in Part iii) of the proof of Proposition 1 satisfies $u_H(\bar{C}_H) = u_H(\bar{C}_L)$.

⁶¹Without loss of generality, we focus on ε small enough so that the constant $D(2\varepsilon)$ defined by (57) is strictly positive.

b) $\beta_{\tilde{N}} < \eta^d$.

Proving that step is the object of Part I below. Of course, Case a) above implies that w_0 could not have been greater than ε , for η small enough. Therefore, it suffices to rule out Case b).

Step 2: Show that, in Case b), there must exist a round $N \geq \tilde{N}$ for which $\bar{w}_N \leq \frac{\eta D}{2a}$ but $\bar{w}_N \geq w\eta$ and $\beta_N \leq \eta^d$, for some exogenous constants a, D, w . This is done in Part II.

Step 3: Show that at round N , one must have $\bar{w}_N \leq \bar{w}\eta^{1+d}$ for some $\bar{w} > 0$. This contradicts, for η small enough, the inequality of Step 2 involving w , which rules out Case b). (Part III).

To avoid cluttering the exposition, the longest proofs for each of these parts are given in separate appendices (Appendix E for Parts I and II and Appendix F for Part III). Several of these proofs are based on geometric inequalities and other inequalities based on incentive constraints which are collected in Appendix D. Once Statement A of Theorem 2 has been proven, showing Statement B is relatively straightforward; the proof is provided in Appendix G.

PART I: BLOCK CONSTRUCTION

The strategy of the proof is to build a sequence of *blocks* (each consisting of finitely many rounds), and a choice sequence going through these blocks, with the following properties: i) within each block, for the PBE to be profitable to P compared to an immediate jump, H must accept H -efficient contracts with a high enough probability, which drives the posterior β closer to zero, by a controlled amount, ii) P's potential gain, conditional on facing type H , shrinks geometrically across blocks. This construction ends at some terminal block, K , such that P's maximal potential gain on H is of order η , and the posterior $\hat{\beta}_K$ is bounded above by $g^K \beta_0$ for some factor $g < 1$. When $\hat{\beta}_K > \eta^d$ for some power $d > 0$ that is judiciously chosen, this yields an upper bound on the number K of blocks which, by using the geometric series backwards, implies that the initial gain on H must have been small as well, for η small enough, contradicting the existence of a sequence $\{\eta_m\}$ and corresponding PBEs for which the initial rent index w_0 always exceeds the constant ε . The ulterior parts (Parts II and III) of the proof establish that $\hat{\beta}_K > \eta^d$ is the only possible case, provided that η is small enough.

For any round n and choice sequence up to round n , let $\bar{e}_n = \max\{u_H(E_H(R_m)) : m \leq n\}$ and $\bar{w}_n = \bar{e}_n - u_H(n)$.

Notation Throughout the analysis, variables with upper bars, such as \bar{w}_n , refer to specific *rounds*, while variables with hats, such as \hat{w}_k , refer to specific *blocks*.

We begin the proof by the following observation.

LEMMA 4 *If $u_H(0) < u_H(E_H(R_0)) - \varepsilon$, there exists a choice sequence and a round n_0 such that i) $\beta_{n_0} \leq \beta_0$ and ii) $\bar{w}_{n_0} \in [\varepsilon/2, \varepsilon]$.*

Block 1 starts after the choice sequence and round n_0 coming from Lemma 4. At round n_0 , one has by construction $\bar{w}_{n_0} = \bar{e}_{n_0} - u_H(n_0) \in (\varepsilon/2, \varepsilon)$. Let $\hat{u}_0 = u_H(n_0)$, $\hat{e}_0 = \bar{e}_{n_0}$, and $\hat{\beta}_0 = \beta_{n_0} \leq \beta_0$. The last round of Block 1 is determined as follows: let \hat{u}_1 be defined by the equation

$$\frac{\hat{e}_0 - \hat{u}_0}{\hat{u}_1 - \hat{u}_0} = t > 1$$

where $t > 1$ is a fixed threshold to be determined shortly and let $n_1 = \inf\{n : u_H(n) \geq \hat{u}_1\}$ denote the first round at which H 's continuation utility exceeds the threshold \hat{u}_1 . We set the last round of Block 1 equal to n_1 . Because $\hat{u}_1 < \hat{e}_0$, Lemma 14 (Appendix E) guarantees that n_1 is finite with probability 1.⁶²

To get a control on how much β must have dropped within Block 1, let $\hat{\mu}_0$ denote the probability, evaluated at round n_0 , that H accepts only contracts in \mathcal{H} until round n_1 (i.e., the probability that H does not fully reveal himself during Block 1). Lemma 15 (Appendix E) shows that there must exist a choice sequence such that, upon observing that sequence up until H 's utility passes \hat{u}_1 , the posterior probability $\hat{\beta}_1$ of facing H satisfies

$$\hat{\beta}_1 \leq \frac{\hat{\beta}_0 \hat{\mu}_0}{\hat{\beta}_0 \hat{\mu}_0 + (1 - \hat{\beta}_0)}. \quad (18)$$

This inequality need not hold for all on path choice sequences, but it must hold for at least one choice sequence arising on the equilibrium path, and we focus on such a sequence, called the **pushdown** sequence. The construction parses together choice sequences: until round n_0 , we use the choice sequence guaranteed to exist by Lemma 4, and between n_0 and n_1 , we use the continuation choice sequence guaranteed to exist by Lemma 15. After n_1 , we will similarly define choice sequences that can be appended to the choice sequence built until round n_1 , etc.

At round n_0 , P can always implement the jump to the contracts $(E_H(R_{n_0}), E_L(R_{n_0}))$, by Lemma 2. For that deviation to be suboptimal, the **net gain** from extracting some rent from H , compared to the jump, must outweigh the **net loss** resulting from a negotiation breakdown at an inefficient contract. In the argument below, it suffices to exploit the inefficiency loss on H . (The loss on L is exploited later in the proof of Theorem 2.)

We now compute an upper bound on this gain and a lower bound on the loss. Comparing these bounds will yield an upper bound on P 's posterior belief of facing H after the first block, following the pushdown choice sequence. The following lemma will be helpful to derive these bounds, as well as in later parts of the proof.

LEMMA 5 *Along any choice sequence, H 's continuation utility at the beginning of round n , denoted $u_H(n)$, is nondecreasing in n and satisfies*

$$u_H(n+1) - u_H(n) \leq \eta \Delta_H,$$

where $\Delta_H = \max_{C \in \mathcal{C}} u_H(C) - \min_{C \in \mathcal{C}} u_H(C)$.

Proof. Given the current contract R_n at round n , let R_{n+1} denote any contract chosen by H with positive probability among $R_n \cup \{M_n\}$. H 's utility satisfies the dynamic equation⁶³

$$u_H(n) = \eta u_H(R_{n+1}) + (1 - \eta) u_H(n+1). \quad (19)$$

Therefore, $u_H(n)$ is a convex combination of $u_H(R_{n+1})$ and $u_H(n+1)$. Because H can hold on to R_{n+1} in all rounds $m \geq n+1$, $u_H(n+1)$ is bounded below by $u_H(R_{n+1})$. Combining these observations yields $u_H(n) \leq u_H(n+1)$. Since $u_H(n+1) - u_H(n) = \eta(u_H(n+1) - u_H(R_{n+1}))$, the second claim of the lemma

⁶²Indeed, $u_H(R_n)$ must eventually exceed any utility level below $\max\{u_H(E_H(R_m)) : m \leq n_0\}$, along any choice sequence, as n gets large enough.

⁶³More generally, H 's utility satisfies the Bellman equation $u_H(n) = \max_{R \in \{R_n\} \cup M_n} \{\eta u_H(R) + (1 - \eta) u_H(n+1)\}$. Equation (19) then follows for all contracts that are optimal for H in round n .

follows. The intuition is simple: if the utility jump was higher between two rounds, H would prefer to wait until the next round rather than accept any contract today. ■

The net gain, between rounds n_0 and n_1 , is bounded above by $\hat{\beta}_0(1 - \hat{\mu}_0)a(\hat{e}_0 - \hat{u}_0)$ for some Lipschitz constant $a > 0$. Indeed, $\hat{\beta}_0(1 - \hat{\mu}_0)$ is the probability that the agent is of type H and that he accepts some H -efficient contract at some round of the first block. Because H accepts only H -efficient contracts that give him at least his continuation utility,⁶⁴ and because that continuation utility is nondecreasing, by Lemma 5, the smallest utility that H when choosing H -efficient contract within that first block is \hat{u}_0 . By contrast, $\hat{e}_0 \geq e_{n_0}$ is an upper bound on the utility that P provides to H if he chooses the immediate jump. Therefore, the maximum rent that P can extract from H is $\hat{e}_0 - \hat{u}_0$. The constant a is a Lipschitz constant that “translates” utility differences for H along \mathcal{E}_H into cost differences for P , and is derived in the “Inequalities” section of the appendix (Lemma 10).

Similarly, the expected net gain made after round n_1 , but seen from round n_0 , is bounded above by $\hat{\beta}_0\hat{\mu}_0a(\hat{e}_0 - \hat{u}_1)$, because $\hat{\beta}_0\hat{\mu}_0$ is the probability of facing H and reaching round n_1 , and \hat{u}_1 is the smallest utility that P must provide to H at any round following n_1 .

To get a lower bound on the net loss, the intuition is that, as long as H accepts contracts in \mathcal{H} , he is getting contracts that are inefficient, and hence costly to P relative to the immediate jump to $E_H(R_{n_0})$. Lemma 16 shows that this inefficiency cost must be greater than some constant $D > 0$ whenever the rent extraction index at the beginning of a block is less than 2ε , which holds without loss of generality (see Remark 3 in Appendix E). This cost is only incurred in case of a breakdown. To compute the probability of a breakdown between rounds n_0 and n_1 , the key is to observe that H ’s utility jumps, at each round, by at most $\eta\Delta_H$, by Lemma 5. Therefore, there must be at least $\underline{n}(1) = \lfloor (\hat{u}_1 - \hat{u}_0)/\eta\Delta_H \rfloor$ steps to get to \hat{u}_1 , for any choice sequence.

This implies that the breakdown probability is bounded below by⁶⁵

$$1 - (1 - \eta)^{\underline{n}(1)} = 1 - \exp(\underline{n}(1) \ln(1 - \eta)) \geq -\underline{n}(1) \ln(1 - \eta) - \frac{1}{2}\underline{n}(1)^2 (\ln(1 - \eta))^2. \quad (20)$$

Because the gain is of order ε , which is small, while the loss conditional on a breakdown is of order D , the probability of a breakdown must be (at most) of order ε , which means that $\underline{n}(1) \ln(1 - \eta)$ must also be small. The quadratic term of (20) is therefore negligible. Moreover, because we are focusing on the case where η is small, $\ln(1 - \eta)$ can be approximated by $-\eta$. Combining these bounds on gains and losses yields⁶⁶

$$\beta a[(\hat{e}_0 - \hat{u}_0)(1 - \hat{\mu}_0) + (\hat{e}_0 - \hat{u}_1)\hat{\mu}_0] \geq \beta \hat{\mu}_0 D \frac{\hat{u}_1 - \hat{u}_0}{\Delta_H}. \quad (21)$$

Recall that \hat{u}_1 was defined in terms of an arbitrary threshold t . We now define t precisely by the following

⁶⁴Indeed, by accepting such contract, H reveals his type, and his continuation utility is exactly the one provided by the last accepted contract, by Proposition 1, Part i).

⁶⁵The inequality comes from the standard inequality $1 - \exp(x) \geq -x - x^2/2$, valid for all $x \leq 0$, which may be shown as follows. The function $x \mapsto \exp(x) - 1 - x - \frac{x^2}{2}$ vanishes at 0, as do its first and second derivatives. Since its third derivative is positive (equal to $\exp(x)$), its first derivative is convex and, from the previous observations, must have a minimum at zero. This implies that the function itself is increasing and, since it vanishes at 0, that it is negative for $x \leq 0$.

⁶⁶For expositional simplicity, the “floor” operator is dropped. This change is negligible because $\underline{n}(1)$ is large, since $\hat{u}_1 - \hat{u}_0 = \frac{1}{t}(\hat{e}_0 - \hat{u}_0) \gg \eta\Delta_H$, for η small. That observation applies to each block k ; see Footnote 69.

equation⁶⁷

$$t^2 = \frac{a + D/\Delta_H}{a} > 1.$$

With this value of t , we have

$$\hat{\mu}_0 \leq \frac{a}{a + D/\Delta_H} \frac{\hat{e}_0 - \hat{u}_0}{\hat{u}_1 - \hat{u}_0} = t^{-1}.$$

Combining this inequality with (18) implies that, upon observing the constructed pushdown choice sequence until round n_1 , the posterior $\hat{\beta}_1$ satisfies

$$\hat{\beta}_1 \leq \frac{\hat{\mu}_0 \hat{\beta}_0}{\hat{\mu}_0 \hat{\beta}_0 + (1 - \hat{\beta}_0)} \leq \hat{\beta}_0 \frac{t^{-1}}{\hat{\beta}_0 t^{-1} + (1 - \hat{\beta}_0)} = g \hat{\beta}_0.$$

where $g = \frac{t^{-1}}{\hat{\beta}_0 t^{-1} + (1 - \hat{\beta}_0)}$ and we used the inequality $\hat{\beta}_0 \leq \beta_0$. Because $t^{-1} < 1$, g is strictly less than 1. This achieves the goal of guaranteeing that the posterior $\hat{\beta}_1$ drops by some fixed factor along the first block, for some choice sequence.

To initiate the second block, we use the value \hat{u}_1 that was defined as part of Block 1.⁶⁸ Note that the actual value of $u_H(n_1)$ may be slightly above \hat{u}_1 , but by no more than $\Delta_H \eta$, by Lemma 5. This observation is useful to bound below the number of rounds in each block. The level $\hat{e}_1 = \max_{m \leq n_1} \{u_H(E_H(R_m))\}$ is the maximum value that H gets if P jumps at any round $m \leq n_1$ along the particular choice sequence constructed so far. Having defined \hat{u}_1 and \hat{e}_1 , we define $\hat{u}_2 \in (\hat{u}_1, \hat{e}_1)$ by

$$\frac{\hat{e}_1 - \hat{u}_1}{\hat{u}_2 - \hat{u}_1} = t.$$

Let $\hat{\mu}_1$ denote the probability, seen from round n_1 and following the pushdown choice sequence used for Block 1, that H takes a contract in \mathcal{H} at all rounds $n \geq n_1$ until \hat{u}_2 is reached. Repeating the previous analysis, there exists a pushdown choice sequence for Block 2 such that, upon observing that sequence up to \hat{u}_2 , the probability $\hat{\beta}_2$ of facing H satisfies $\hat{\beta}_2 \leq \frac{\hat{\mu}_1 \hat{\beta}_1}{\hat{\mu}_1 \hat{\beta}_1 + (1 - \hat{\beta}_1)}$. Let n_2 denote the round at which \hat{u}_2 is first exceeded. Repeating the analysis used for the first block, we have

$$\hat{\beta}_2 \leq \frac{\hat{\mu}_1 \hat{\beta}_1}{\hat{\mu}_1 \hat{\beta}_1 + (1 - \hat{\beta}_1)} \leq \hat{\beta}_1 \frac{t^{-1}}{\hat{\beta}_1 t^{-1} + (1 - \hat{\beta}_1)} \leq g^2 \hat{\beta}_0.$$

The value of \hat{e}_2 is determined by the pushdown sequence of the second block, by $\hat{e}_2 = \max_{m \leq n_2} \{u_H(E_H(R_m))\}$.

By induction, this defines a sequence of blocks indexed by k . To each block k corresponds a terminal round, n_k , as well as values \hat{u}_k, \hat{e}_k and $\hat{\beta}_k = \beta_{n_k}$, which is P 's belief at the end of the k^{th} block following the successive pushdown sequences. The potential overshoot of $u_H(n_k)$ above \hat{u}_k , $\Delta_H \eta$ is negligible when computing the lower bound on number of blocks, because we stop the block construction when $\hat{u}_{k+1} - \hat{u}_k$ is still large relative to $\Delta_H \eta$, as explained in the next paragraph. Upon observing the pushdown choice sequence across blocks 1 to k , we get

$$\hat{\beta}_k \leq g^k \hat{\beta}_0 \leq g^k \beta_0.$$

⁶⁷ D is defined independently of t (and of this entire block construction), so there is no circularity in the definition.

⁶⁸The next block is defined only following the pushdown choice sequence that we constructed in Block 1: what matters to us is to understand what happens along a particular choice sequence constructed by piecing together pushdown sequences constructed for each block.

The construction stops at the first block, K , such that $\hat{e}_k - \hat{u}_k < \bar{W}\eta$ for some constant \bar{W} such that $\bar{W} > \max\{\frac{t-1}{t}(1 + \Delta_H), \frac{\bar{W} + \Delta_H}{t\Delta_H}\}$, where \bar{W} is an arbitrarily large constant.⁶⁹ Such a block must exist, because $\hat{w}_k = \hat{e}_k - \hat{u}_k$ converges to zero, as shown in Part ii) of Lemma 14. Let ρ be defined by $g^{-\rho} = \frac{t}{t-1}$. Since the ratio is greater than 1 and $g < 1$, ρ is strictly positive. Also let $d = \frac{1}{2} \min\{\frac{1}{\rho}, 1\} \in (0, 1/2]$.

As explained earlier, the key to showing Theorem 2 is the following proposition, whose proof is the object of Parts II and III.

PROPOSITION 3 *There exists $\tilde{\eta} > 0$ such that $\hat{\beta}_K > \eta^d$ for all $\eta < \tilde{\eta}$ and $\hat{\beta}_K$ resulting from piecing together pushdown sequences as constructed above, given any PBE with friction parameter η .*

It must be emphasized that the proposition holds for any $d > 0$. However, we only need it for the value of d defined above. Taking Proposition 3 as given for now, we compute an upper bound on the initial rent, by backward induction. For each block $k \leq K$, we have

$$\hat{e}_K - \hat{u}_k = (\hat{e}_K - \hat{u}_{k+1}) + (\hat{u}_{k+1} - \hat{u}_k) \leq (\hat{e}_K - \hat{u}_{k+1}) + \frac{1}{t-1}(\hat{e}_k - \hat{u}_{k+1}) \leq \frac{t}{t-1}(\hat{e}_K - \hat{u}_{k+1}).$$

By construction, moreover, $\hat{e}_K - \hat{u}_K \leq \bar{W}\eta$, which implies that

$$\hat{e}_K - \hat{u}_0 \leq \left(\frac{t}{t-1}\right)^K \bar{W}\eta. \quad (22)$$

Since $\hat{\beta}_K \geq \eta^d$ and $\hat{\beta}_K \leq g^K \hat{\beta}_0 < 1$, we must also have

$$\frac{1}{g^K} \eta^d \leq 1.$$

Combining these inequalities yields

$$\hat{e}_K - \hat{u}_0 \leq \left(\frac{t}{t-1}\right)^K \bar{W}\eta = \bar{W}\eta \left(\frac{1}{g}\right)^{\rho K} \leq \eta \bar{W} \eta^{-\rho d} \leq \bar{W} \eta^{1/2}.$$

Since $\hat{e}_K \geq \hat{e}_0$, this shows that $\hat{e}_0 - \hat{u}_0 = O(\eta^{1/2})$, which contradicts the existence of the sequence of $\{\eta_m\}$, converging to zero, and corresponding PBEs for which $\hat{e}_0 - \hat{u}_0 \in (\varepsilon/2, \varepsilon)$.

To avoid confusion, we remind the reader of the ‘ o ’ and ‘ O ’ notation: A function $f(\eta)$ is ‘ $o(\eta)$ ’ if $\lim_{\eta \downarrow 0} |f(\eta)|/\eta = 0$. It is ‘ $O(\eta)$ ’ if there exists a constant $\bar{f} > 0$ such that $|f(\eta)| \leq \bar{f}\eta$ for all η small enough.⁷⁰

PART II: BRIDGING ARGUMENT

The objective of Parts II and III is to prove Proposition 3. Suppose, by contradiction, that $\hat{\beta}_K < \eta^d$. By definition of K , the previous block $K-1$ must satisfy $\hat{e}_{K-1} - \hat{u}_{K-1} > \bar{W}\eta$. This implies that⁷¹

$$\hat{w}_K = \hat{e}_K - \hat{u}_K \geq \bar{W}\eta \quad (23)$$

⁶⁹The number of rounds in each block $k \leq K$ is bounded below by $\frac{\hat{u}_k - \hat{u}_{k-1} - \Delta_H \eta}{\Delta_H \eta} \geq \frac{1}{t\Delta_H \eta}(\hat{e}_{k-1} - \hat{u}_{k-1} - \Delta_H \eta) \geq \frac{\bar{W}\eta - \Delta_H \eta}{t\eta\Delta_H} > \bar{W}$, which can be made arbitrarily large by choosing \bar{W} appropriately. The reason for also requiring that $\bar{W} > \frac{t-1}{t}(1 + \Delta_H)$ is explained at the beginning of Part II.

⁷⁰Equivalently, $\limsup_{\eta \downarrow 0} |f(\eta)|/\eta$ is finite.

⁷¹This inequality comes from the fact that $\hat{e}_{K-1} - \hat{u}_{K-1} = t(\hat{u}_K - \hat{u}_{K-1})$, by construction of the blocks in Part I, and the fact that $\hat{e}_K \geq \hat{e}_{K-1}$.

where $\bar{W} = \frac{t-1}{t}\bar{W}$. Since we chose $\bar{W} > \frac{t}{t-1}(1 + \Delta_H)$, we have $\bar{W} > 1 + \Delta_H$. Combining this with Lemma 5, we obtain, for the augmented index evaluated at round n_K ,⁷²

$$\hat{w}_K = \hat{e}_K - u_H(n_K) \geq (\bar{W} - \Delta_H)\eta \geq \eta. \quad (24)$$

Part III (Proposition 5) establishes that if $\hat{w}_K \leq \frac{D\eta}{2a}$, then $\hat{w}_K \leq \hat{w}\eta^{1+d}$, which contradicts (24) for η small enough.

Unfortunately, nothing guarantees that \hat{w}_K lies below $\frac{D\eta}{2a}$. The purpose of Part II is to bridge Parts I and III when $\hat{w}_K \in \left(\frac{D\eta}{2a}; \bar{W}\eta\right)$. We will analyze the dynamics of β_n and \bar{w}_n along some appropriate choice sequence between the levels \hat{w}_K and $\frac{\eta D}{2a}$, and establish the following result.

PROPOSITION 4 (BRIDGE) *Suppose that $\hat{\beta}_K < \eta^d$. Then, letting $N \geq n_K$ denote the first round for which $\bar{w}_N \leq \frac{\eta D}{2a}$, there exists a choice sequence such that*

1. $\bar{w}_N \geq \frac{\eta D}{2a} - o(\eta)$
2. $\beta_N = O(\eta^d)$.

Thus, if $\hat{\beta}_K < \eta^d$, there must exist a choice sequence and a round N to which the contradiction argument of Part III can be applied.

To construct a choice sequence that yields the two conclusions of Proposition 4, we start by exploiting P's IC constraint, similarly to what was done in Part I. This time, however, there are no blocks: the equation is used at every single round n , and exploits the losses on *both* types. For each $R_{n+1} \in M_n \cup \{R_n\}$, let $\mu_n^\theta(R_{n+1})$ denote the probability that θ accepts R_{n+1} . Because P can always jump to $E_L(R_n), E_H(R_n)$ (cf. Lemma 2), P's IC constraint implies as explained below that

$$w_n a \beta_n \geq \sum_{R_{n+1} \in (M_n \cup \{R_n\}) \cap \mathcal{H}} \beta_n \mu_n^H(R_{n+1}) \eta D + (1 - \beta_n) \mu_n^L(R_{n+1}) \eta (Q(R_{n+1}) - Q(E_L(R_n))) \quad (25)$$

$$= \sum_{R_{n+1} \in (M_n \cup \{R_n\}) \cap \mathcal{H}} \mu_n^L(R_{n+1}) [\beta_n \mu_n(R_{n+1}) \eta D + (1 - \beta_n) \eta (Q(R_{n+1}) - Q(E_L(R_n)))], \quad (26)$$

where $\mu_n(R_{n+1}) = \mu_n^H(R_{n+1}) / \mu_n^L(R_{n+1})$ and D is the lower bound on the loss on H given in Lemma 11.⁷³

The left-hand side of (25) is an upper bound on the gain, relative to the immediate jump, made on the high type: given his continuation utility $u_H(n)$, the lowest achievable cost that provides this utility is the cost of the H -efficient contract that gives $u_H(n)$. From Lemma 10 of Appendix D, this gain is bounded above by $a(u_H(E_H(R_n)) - u_H(n)) = aw_n$ (that bound is computed using a ‘best-case scenario’ for P, in which H accepts with probability 1 the H -efficient contract C_n providing $u_H(n)$).⁷⁴ The first term of the right-hand side is the net loss on H if he accepts a contract in the H -Rent configuration and, hence, far from efficient,

⁷²The reason for using $u_H(n_K)$ instead of \hat{u}_K is that H 's continuation utility at round n_K need not be exactly equal to \hat{u}_K : it lies between \hat{u}_K and $\hat{u}_K + \Delta_H \eta$.

⁷³We can assume without loss of generality that $\mu_n^L(R_{n+1})$ is strictly positive for all $R_{n+1} \in (M_n \cup \{R_n\}) \cap \mathcal{H}$: first, if any contract in that set is not chosen with any probability, we can construct an equilibrium in which those contracts are removed. And if any contract R'_{n+1} in that set is chosen only by H with positive probability, then Proposition 1 implies that H gets the H -efficient contract C that gives him the same utility as R'_{n+1} , so that the equilibrium can be modified by having P propose C instead of R'_{n+1} . That change reduces P's cost and does not affect incentives.

⁷⁴This is an upper bound on the gain, since C_n is the cheapest way of providing H with his continuation utility.




Figure 5: Rent extraction concepts (contracts are in black, utilities are in purple)

in case of a breakdown at the end of round n . This loss is bounded below by D as long as $w_n \leq 2\varepsilon$, which will be true along the choice sequence that we consider. The last term is the net loss on L in case of such a breakdown.

To establish Proposition 4, we introduce the variable $y_n = u_H(E_H(R_n)) - u_H(R_{n+1})$, i.e., H 's utility gap, for any choice R_{n+1} , between the immediate jump and his utility in case of a negotiation breakdown at round n (the breakdown occurs *after* the agent has chosen the new contract, R_{n+1} , which explains the index). This quantity y_n is important for the analysis, because it provides a control on the decrements of w_n and makes sure that we do not overshoot the threshold $\frac{\eta D}{2a}$ by too much. Indeed, subtracting $u_H(E_H(R_n))$ from (19) and rearranging (and recalling that $w_n = u_H(E_H(R_n)) - u_H(n)$) leads, along any choice sequence, to

$$w_{n+1} = w_n - \eta y_n + \eta w_{n+1} + (1 - \eta)(u_H(E_H(R_{n+1})) - u_H(E_H(R_n))).$$

These concepts are represented on Figure 5

Proposition 4 is based on the following lemma, which is proved in Appendix E. Fix a positive integer \bar{N} , positive constants $\bar{\beta}$ and \bar{w} , and a small positive constant $\bar{\varepsilon}$.

LEMMA 6 *Suppose that negotiations reach a round \bar{n} such that $\beta_{\bar{n}} \leq \bar{\beta}\eta^d$ and $w_{\bar{n}} \leq \bar{w}\eta$. There exist functions $\bar{W}(\bar{N})$ and $k(\bar{N})$ of \bar{N} (only) such that the following holds. Let \mathcal{S} denote the event that the agent chooses at all rounds $n \in \{\bar{n} + 1, \dots, \bar{n} + \bar{N}\}$ contracts such that $y_n = O(\eta^{d/4})$, $\beta_n \leq \beta_{\bar{n}}\bar{\varepsilon}^{-(n-\bar{n})}$, and $w_n \leq \bar{W}(\bar{N})\eta$. For η small enough, the probability of \mathcal{S} is greater than $1 - k(\bar{N})\bar{\varepsilon}$.*

We now modify the analysis of Part I to study blocks consisting of \bar{N} rounds, indexed by $\bar{n} + 1$ to $\bar{n} + \bar{N}$,

where \bar{N} will be determined shortly. The first such block starts with $\bar{n} = n_K$, the second of these blocks starts with $\bar{n} = n_K + \bar{N}$, etc. These blocks are different from those of Part I, because the number \bar{N} of rounds in each block is fixed and, unlike the blocks of Part I, H 's utility at the end of each block is not precisely controlled.

The analysis of Part I is modified as follows. First, notice that P 's IC constraint at round \bar{n} , looking ahead over the next \bar{N} rounds, implies that

$$\beta_{\bar{n}} a \left\{ (1 - \mu_{\bar{n}})(e_{\bar{n}} - u_H(\bar{n})) + \mu_{\bar{n}}(e_{\bar{n}} - E[u_H(\bar{n} + \bar{N})]) \right\} \geq \beta_{\bar{n}} \mu_{\bar{n}} D \eta \bar{N} - \beta_{\bar{n}} \delta_Q k(\bar{N}) \bar{\varepsilon},$$

where $\mu_{\bar{n}}$ is the probability, seen from round \bar{n} , that H rejects all H -efficient contracts between rounds \bar{n} and $\bar{n} + \bar{N}$. The argument for this equation is the same as in Part I, the only difference being that we are now taking the expectation of $u_H(\bar{n} + \bar{N})$ because we do not know its value (before, we had precisely defined the end of the block as the first time that u_H crosses some level, but now \bar{N} is exogenous). The lower bound D on the loss is valid with probability $1 - k(\bar{N})\bar{\varepsilon}$ because conditional on \mathcal{S} occurring, w_n is small throughout the block by Lemmas 6 and the lower bound on the loss follows from Lemma 16. On the complement of \mathcal{S} , whatever cost is incurred by P conditional on facing H is bounded above by some constant, by compactness of the contract space. The difference between D and that constant is captured by $\delta_Q > 0$ (last term of the above equation), which is independent of η , \bar{N} , and $\bar{\varepsilon}$.

This implies that

$$\mu_{\bar{n}} \leq \frac{a(e_{\bar{n}} - u_H(\bar{n})) + \delta_Q k(\bar{N}) \bar{\varepsilon}}{a(E[u_H(\bar{n} + \bar{N})] - u_H(\bar{n})) + D \eta \bar{N}} \leq \frac{a \bar{W} + \delta_Q k(\bar{N}) \bar{\varepsilon}}{D \bar{N}},$$

where the second inequality comes from i) the fact that $u_H(n)$ is nondecreasing across all paths which implies, taking expectations, that $E u_H(\bar{n} + \bar{N}) \geq u_H(\bar{n})$ and ii) $\bar{w}_{\bar{n}} = e_{\bar{n}} - u_H(\bar{n}) \leq \bar{W} \eta$ (this inequality holds for all blocks of Part II, without loss of generality, see Remark 2 below). Now let $\mu_{\bar{n}}^{\mathcal{S}}$ (resp. $\mu_{\bar{n}}^{\mathcal{B}}$) denote the probability that H rejects all H -efficient contracts, *conditional* on event \mathcal{S} (resp. conditional on its complement, \mathcal{B}), and let $p_{\mathcal{S}}$ (resp. $p_{\mathcal{B}}$) the probability of \mathcal{S} (\mathcal{B}). We have $\mu_{\bar{n}} = p_{\mathcal{S}} \mu_{\bar{n}}^{\mathcal{S}} + p_{\mathcal{B}} \mu_{\bar{n}}^{\mathcal{B}}$. Since $p_{\mathcal{S}} \geq 1 - k(\bar{N}) \bar{\varepsilon}$, we conclude that

$$\mu_{\bar{n}}^{\mathcal{S}} \leq \frac{(a \bar{W} + \delta_Q k(\bar{N}) \bar{\varepsilon})(1 + k(\bar{N}) \bar{\varepsilon})}{D \bar{N}}. \quad (27)$$

We now choose $\bar{\varepsilon}$ and \bar{N} so that this ratio is less than $\frac{1}{2}$: first choose \bar{N} so that $\frac{a \bar{W}}{D \bar{N}} < \frac{1}{8}$, then choose $\bar{\varepsilon}$ so that $k(\bar{N}) \bar{\varepsilon} < 1$ and $\delta_Q k(\bar{N}) \bar{\varepsilon} < a \bar{W}$ so that the numerator of (27) is less than $4a \bar{W}$.

Proceeding as in Part I, there must exist a pushdown choice sequence contained in event \mathcal{S} such that the ex post probability that H has not chosen an H -efficient contract is weakly less than $\mu_{\bar{n}}^{\mathcal{S}}$. Along that sequence, i) y_n is of order $O(\eta^{d/4})$, and ii) $\beta_{\bar{n} + \bar{N}} \leq \frac{\beta_{\bar{n}}}{2}$. We have thus built a choice sequence over \bar{N} rounds, starting from \bar{n} , such that y_n and β_n stay small, and β_n ends up *smaller* than at the beginning of the block.⁷⁵

Starting from round n_K , we build a sequence of \bar{N} -sized blocks as described above. Because \bar{w}_n converges to zero (by part ii) of Lemma 14), it will eventually cross $\frac{D \eta}{2a}$. Let N denote the first round at which \bar{w}_N drops below that threshold. From (52) of Lemma 11 (Appendix D), we have $w_{n+1}(1 - b\beta_{n+1}) \geq w_n - \eta y_n$, or

$$w_{n+1} - w_n \geq b\beta_{n+1} w_{n+1} - \eta y_n. \quad (28)$$

⁷⁵Notice that β_n can increase up to $\beta_{\bar{n}} \bar{\varepsilon}^{-\bar{N}}$ along such a block. However, because \bar{N} is fixed, it still remains of order $O(\eta^d)$ along the sequence, and drops in any case below $\beta_{\bar{n}}/2$ when round $\bar{n} + \bar{N}$ is reached, for the pushdown sequence.

The blocks were constructed in such a way that $y_n = O(\eta^{d/4})$ and $\beta_n = O(\eta^d)$ at each round of each block. Applying these observations to (28) at round $N - 1$, we obtain

$$w_N - w_{N-1} \geq -o(\eta).$$

Finally, we have

$$\begin{aligned} \bar{w}_N - \bar{w}_{N-1} &= (w_N - w_{N-1}) + (\max\{e_k : k \leq N\} - e_N) - (\max\{e_k : k \leq N-1\} - e_{N-1}) \\ &\geq -o(\eta) - (e_N - e_{N-1}), \end{aligned}$$

since the different of maxima is nonnegative. From (48), the difference $e_N - e_{N-1}$ is bounded above by $\frac{\alpha\beta_{N-1}}{1-\beta_{N-1}}w_{N-1}$ which is $o(\eta)$. Since $\bar{w}_{N-1} > \frac{\eta D}{2a}$ by definition of N , we conclude that⁷⁶

$$\bar{w}_N \geq \bar{w}_{N-1} - o(\eta) \geq \frac{\eta D}{2a} - o(\eta) \geq \frac{\eta D}{3a}.$$

This concludes the proof of Proposition 4 and implies that we have reached a round N such that i) \bar{w}_N is above $\hat{w}\eta$, for some $\hat{w} = \frac{D}{3a} > 0$ independent of η , and ii) $\beta_N = O(\eta^d)$.⁷⁷ Part III will show that this is impossible.

REMARK 2 *It is a priori possible that \bar{w}_n goes above $\bar{W}\eta$ at the end of some block. If that happens, the bound $D = D(2\varepsilon)$ need not be valid for the next block. At the end of such a block, should it occur, β_n is of order $\eta^d \leq \beta_0$. We can restart the blocks of Part I as if n were the initial round. Since β decreases along the blocks of Part I, we have to reach again a round at which \bar{w}_n drops below $\bar{W}\eta$. At that point, we necessarily have $\beta_n \leq \eta^d$. Because \bar{w}_n converges to zero along any sequence (by Lemma 14), and thus also along the sequences constructed through Parts I and II, the back and forth between blocks of Part I and Part II has to stop in finite time at some round N of the type above, i.e., with $\bar{w}_N \in (\hat{w}\eta, \frac{\eta D}{2a})$ and $\beta_N \leq \eta^d$. The logic of the argument is explained in more detail in Remark 3.*

PART III: ASYMPTOTIC ARGUMENT

The purpose of this section is to establish the following proposition:

PROPOSITION 5 *There exist a constant $\check{w} > 0$ and a threshold $\check{\eta} > 0$ such that the following holds for all $\eta < \check{\eta}$: if one reaches a round N for which $\beta_N \leq \eta^d$ and $\bar{w}_N \leq \frac{\eta D}{2a}$, then $\bar{w}_N \leq \check{w}\eta^{1+d}$.*

The proof proceeds in three steps:

1. Show, starting from round N , that one can build a choice sequence along which β_n is decreasing and, at each round, a simplified version of P's ex ante IC constraint (before the agent chooses R_{n+1}) is also satisfied ex post (after R_{n+1} is chosen). This step is achieved by Lemma 7;
2. Show that along such a sequence, one must necessarily have $w_n \leq \hat{c}\eta\beta_n$ for all $n \geq N$, where $\hat{c} > 0$ is independent of η (Proposition 6);

⁷⁶The last inequality must hold for η small enough, since $o(\eta)$ is negligible compared to η .

⁷⁷As a reminder, the notation $\beta_N = O(\eta^d)$ means that β_N is bounded above by $\check{\beta}\eta^d$ for some positive constant $\check{\beta}$ as η becomes arbitrarily small. But β_N could be much smaller than this bound, too.

3. Show that $\bar{w}_N - w_N = O(\eta^{1+2d})$ (Proposition 7).

Combining these steps (with 2. applied to $n = N$) along with the fact that $\beta_N \leq \eta^d$ then proves Proposition 5.⁷⁸

To express P's IC constraint, recall from Part II that

$$w_n a \beta_n \geq \sum_{R_{n+1} \in (M_n \cup \{R_n\}) \cap \mathcal{H}} \mu_n^L(R_{n+1}) [\beta_n \mu_n(R_{n+1}) \eta D + (1 - \beta_n) \eta (Q(R_{n+1}) - Q(E_L(R_n)))] \quad (29)$$

where $\mu_n(R_{n+1}) = \mu_n^H(R_{n+1}) / \mu_n^L(R_{n+1})$ and D is the lower bound on the loss on H introduced earlier.⁷⁹

In particular, the RHS of (29) is a convex combination of terms indexed by R_{n+1} , and there must exist $R_{n+1} \in (M_n \cup \{R_n\}) \cap \mathcal{H}$ such that

$$w_n a \beta_n \geq \beta_n \mu_n(R_{n+1}) \eta D + (1 - \beta_n) \eta (Q(R_{n+1}) - Q(E_L(R_n))). \quad (30)$$

Therefore, there exists a choice sequence that satisfies (30) for all $n \geq N$. In what follows we entirely focus on that sequence, which will be called a **regular** choice sequence.

We have $Q(R_{n+1}) - Q(E_L(R_n)) \geq Q(E_L(R_{n+1})) - Q(E_L(R_n)) \geq -k\beta_{n+1}w_{n+1}$, where the second inequality comes from (50) of Lemma 11 (Appendix D). Letting $\mu_n = \mu_n(R_{n+1})$, (30) implies that

$$\beta_n w_n a \geq \beta_n \mu_n \eta D - \eta k \beta_{n+1} w_{n+1},$$

which may be re-expressed as

$$\mu_n \leq \frac{w_n a}{\eta D} + k \frac{\beta_{n+1}}{\beta_n D} w_{n+1}. \quad (31)$$

The first step, in order to exploit this equation, is to show that β_n remains small for all $n \geq N$. This is achieved by the following lemma, which guarantees that β_n is actually decreasing along the regular sequence (see Appendix F for the proof).

LEMMA 7 *There exist $\hat{\eta} > 0$ and $\hat{w} > 0$ such that, for $\eta < \hat{\eta}$ i) β_n is decreasing in n and bounded above by η^d , ii) $\mu_n \leq 3/4$ for all $n \geq N$, and iii) $w_n \leq \hat{w}\eta$ for all $n \geq N$.*

Part ii) of Lemma 7 implies that the second term in the right-hand side of (31) is of order w_{n+1} (since $\beta_{n+1} \sim \beta_n \mu_n$ from Bayes rule, see (38)) and is thus negligible compared to the first term, of order $\frac{w_n}{\eta}$, because w_{n+1} is bounded above by $w_n \left(1 + \frac{\alpha \beta_n}{1 - \beta_n}\right)$ (see (49) in Appendix D) and $\eta \ll 1$. Therefore, by slightly increasing a , whose specific value does not matter in any case for the proof, we get

$$\mu_n \leq \frac{w_n a}{\eta D}, \quad (IC_n^{LL}) \quad (32)$$

Moreover, (30) also implies that

$$\beta_n w_n a \geq (1 - \beta_n) \eta (Q(R_{n+1}) - Q(E_L(R_n))) \quad (IC_n^{LH}) \quad (33)$$

⁷⁸Indeed, we have $\bar{w}_N \leq w_n + O(\eta^{1+2d}) \leq \hat{c}\eta\beta_N + O(\eta^{1+2d}) \leq \hat{c}\eta^{1+d} + O(\eta^{1+2d})$, where the first inequality comes from Step 3., the second comes from Step 2. applied to $n = N$, and the third one comes from the assumption that $\beta_N \leq \eta^d$. Taking \tilde{w} slightly above \hat{c} then yields the proposition.

⁷⁹This bound holds only for w_n smaller than 2ε . However, Lemma 7 shows that w_n is $O(\eta)$, which guarantees that the bound applies. More precisely, the proof of Lemma 7 shows inductively, starting from round N , that at each subsequent round n less than some round M defined in the proof, w_n is small enough to guarantee that the loss bound D and equation (29) can be used for the next round. This shows that the loss bound D is valid for all rounds $n \leq M$. Since M is in fact shown to be infinite, the loss bound D is valid for all rounds $n \geq N$.

Relaxation

The previous arguments have shown that any PBE must entail, given the assumptions of Proposition 5, a regular choice sequence satisfying (32) and (33) for all $n \geq N$. Moreover, Appendix D contains a list of inequalities which must also be satisfied at all these rounds. Finally, w_n converges to zero along that regular sequence, by Proposition 2.

From now on, we entirely focus on all these inequalities, which arise along the regular sequence. Those inequalities involve the contracts chosen along that sequence as well as the beliefs β_n 's and likelihood ratios μ_n 's for $n \geq N$. We work with these inequalities in complete isolation from the rest of the PBE. The objective here is to prove some properties of the *contracts* involved in this sequence by modifying the *beliefs* β_n and likelihood ratios μ_n .

The transformation works as follows. Starting with round N , one increases the likelihood ratio μ_N so as to satisfy (32) as an equality. The interpretation of this change is that H becomes relatively more likely than before the change, conditional on observing the contract R_{N+1} arising in the regular sequence. We maintain standard Bayesian updating, but applied to this new likelihood ratio, so that β_{N+1} is given by

$$\beta_{N+1} = \frac{\beta_N \mu_N}{\beta_N \mu_N + (1 - \beta_N)}.$$

With only that change, the posteriors β_n 's at all rounds $n \geq N + 1$ of the regular sequence are weakly increased as a result of Bayesian updating. Therefore, the inequality (33) at all rounds $n \geq N + 1$ is preserved (in fact, looser) along the regular sequence. After this is done, one can increase the likelihood ratio μ_{N+1} pertaining to the contract R_{N+2} chosen in the regular sequence at round $N + 1$, so as to make IC_{N+1}^{LL} tight. All beliefs β_n for $n \geq N + 2$ are then weakly increased as a result of Bayesian updating. This preserves the inequalities (33) for $n \geq N + 2$ and does not perturb (33) and (32) at round N . Increasing μ_n inductively for all $n \geq N$ along the regular choice sequence, the new $\{\mu_n\}_{n \geq N}$ and $\{\beta_n\}_{n \geq N+1}$ satisfy (33) as well as

$$\mu_n = \frac{w_n a}{\eta D}. \quad (34)$$

Moreover, because all the posteriors β_n 's have been weakly increased compared to the initial belief sequence, while the contracts involved in the regular sequence are left completely intact by the transformation, all the inequalities appearing in Appendix D still hold, because they only get looser for higher values of β_n .⁸⁰

This transformation leads to a new sequence of beliefs and likelihood ratios which, along with the contracts of the regular sequence satisfy all the previous inequalities and now satisfy (32) as equality (34). Because the contracts have not been changed, moreover, w_n must still converge to zero. From (32), this implies that the sequence of μ_n 's also goes to zero and, hence, that β_n goes to zero even after the relaxation.

We are now ready to perform the Steps 2 and 3 needed to prove Proposition 5. In Appendix D (Lemma 11), it is shown that

$$u_H(E_H(R_{n+1})) - u_H(E_H(R_n)) \geq -\hat{b} \beta_{n+1} w_{n+1} \quad (35)$$

for some constant $\hat{b} > 0$. As in Part II, we use the following relation between w_n and w_{n+1} (see (54)):

$$w_{n+1} = w_n - \eta y_n + \eta w_{n+1} + (1 - \eta)(u_H(E_H(R_{n+1})) - u_H(E_H(R_n))).$$

⁸⁰The only inequality which does not get looser is (52). However, this inequality was only used in Part II, is not involved in the derivation of any other inequality of Appendix D and does not appear anywhere in the analysis of regular sequences.

Combining this with (35) yields

$$(1 - \eta)w_{n+1} \geq w_n - \eta y_n - \hat{b}\beta_{n+1}w_{n+1}. \quad (36)$$

Multiplying both sides of (36) by $\frac{a}{\eta D}$ and using (34), we obtain for $n \geq N$

$$\mu_{n+1} \geq (1 - \eta)\mu_{n+1} \geq \mu_n - \frac{a}{D}y_n - \tilde{b}\beta_{n+1} \quad (37)$$

for some constant $\tilde{b} > 0$. (The term \tilde{b} comes from the inequality $w_n \leq \hat{w}\eta$ for $n \geq N$ in Part iii) of Lemma 7.)

The Bayesian updating equation

$$\beta_{n+1} = \frac{\beta_n \mu_n}{\beta_n \mu_n + (1 - \beta_n)}$$

implies that⁸¹

$$\frac{\beta_{n+1}}{\beta_n} \geq \mu_n \geq \frac{\beta_{n+1}}{\beta_n} - \mu_n \beta_n + \mu_n O(\beta_n^2) \geq \frac{\beta_{n+1}}{\beta_n} - \beta_{n+1} + o(\beta_{n+1}). \quad (38)$$

Lemma 17 of Appendix F shows that $y_n^2 \leq \frac{\bar{A}\beta_{n+1}}{1-\beta_0}$. Intuitively, this equation means that the loss on L in round n , which is of order ηy_n^2 , must be smaller than the gain on H , which is of order $\beta_n w_n$ (i.e., the probability of facing H times the maximum gain).⁸²

Combining this upper bound for y_n^2 with (37) and (38), we obtain the following dynamic equation for β_n , for all $n \geq N$:

$$\frac{\beta_{n+2}}{\beta_{n+1}} \geq \frac{\beta_{n+1}}{\beta_n} - c\sqrt{\beta_{n+1}} - (1 + \tilde{b})\beta_{n+1} \quad (39)$$

where $c = \frac{a}{D}\sqrt{\frac{\bar{A}}{1-\beta_0}}$. For β_{n+1} small enough, the last term is negligible compared to the penultimate term. Therefore, by slightly increasing the value of c , whose precise value does not affect the proof, we obtain

$$\frac{\beta_{n+2}}{\beta_{n+1}} \geq \frac{\beta_{n+1}}{\beta_n} - c\sqrt{\beta_{n+1}}. \quad (40)$$

Let, for all n , $q_n = \frac{\beta_{n+1}}{\beta_n}$. We have $\prod_0^n q_k = \frac{\beta_{n+1}}{\beta_0}$ and (40) may be rewritten as

$$q_{n+1} \geq q_n - c'\sqrt{\prod_0^n q_k} \quad (41)$$

where $c' = \sqrt{\beta_0}c$. Since q_n is proportional to μ_n and, hence, w_n , it converges to zero as $n \rightarrow +\infty$.

PROPOSITION 6 *Let $\hat{c} = \frac{4Dc^2}{a}$. Along the regular choice sequence, we have $w_n \leq \hat{c}\eta\beta_n$ for all $n \geq N$.*

Proof. The proposition is based on the following two lemmas, proved in Appendix F.

LEMMA 8 *Suppose that there exists a round $\hat{N} > N$ such that*

$$\beta_{\hat{N}+1} \geq 4c^2\beta_{\hat{N}}^2, \quad (42)$$

$$\beta_{\hat{N}}^{1/4} \leq \frac{1}{2\sqrt{c}}, \quad (43)$$

Then,

$$\liminf_{n \rightarrow +\infty} \frac{q_{n+1}}{q_n} \geq 1.$$

⁸¹We have $\frac{\beta_{n+1}}{\beta_n} = \mu_n \frac{1}{1-\beta_n(1-\mu_n)} = \mu_n(1 + \beta_n(1-\mu_n)) + \mu_n O(\beta_n^2)$. Rearranging yields the second inequality.

⁸²Dividing by η , we get $y_n^2 \leq \chi\beta_n w_n / \eta$ for some constant χ . Since w_n / η is proportional to μ_n and $\mu_n \beta_n$ is roughly equal to β_{n+1} , this provides some intuition for how the equation was derived.

LEMMA 9 Suppose that $\{q_n\}$ is a strictly positive sequence such that

$$q_n - q_{n+1} \leq c' \sqrt{\Pi_0^n q_k}$$

for some $c' > 0$ and $\liminf_n \frac{q_{n+1}}{q_n} \geq 1$. Then, $\{q_n\}$ does not converge to zero.

To conclude the proof of Proposition 6, suppose by contradiction that there exists $\hat{N} \geq N$ such that $w_{\hat{N}} > \hat{c}\eta\beta_{\hat{N}}$. From (34), this implies that $\mu_{\hat{N}} > \frac{a}{D}\hat{c}\beta_{\hat{N}}$, and from the first inequality of (38) this implies that $\beta_{\hat{N}+1} \geq \frac{a\hat{c}}{D}\beta_{\hat{N}}^2$. From the definitions of c and \hat{c} , this shows that (42) holds for \hat{N} . Moreover, from Lemma 7, $\beta_{\hat{N}}$ clearly satisfies (43) for η small enough. Therefore the hypotheses of Lemma 8 are satisfied and, hence, $\liminf_{n \rightarrow +\infty} \frac{q_{n+1}}{q_n} \geq 1$. Combining this with Lemma 9 then implies that w_n cannot converge to zero, which contradicts Proposition 2, since w_n converges to zero along any choice sequence. ■

PROPOSITION 7 There exist $\hat{w} > 0$ and $\bar{\eta} > 0$ such that $\bar{w}_N - w_N \leq \hat{w}\eta^{1+2d}$ for all for $\eta \leq \bar{\eta}$.

Proof. Recalling the definition of \bar{C}_H as the H -efficient contract that provides H with its asymptotic utility $\lim_n u_H(n)$ (see Proposition 2), we have

$$u_H(\bar{C}_H) - u_H(E_H(R_N)) = \sum_{n \geq N} u_H(E_H(R_{n+1})) - u_H(E_H(R_n)) \leq 2\alpha \sum_{n \geq N} \beta_n w_n,$$

where the last inequality comes from (48). From Proposition 6, we have $w_n \leq \hat{c}\eta\beta_n$ for all $n \geq N$. Therefore,

$$u_H(\bar{C}_H) - u_H(E_H(R_N)) \leq \tilde{K}\eta \sum_{n \geq N} \beta_n^2,$$

where $\tilde{K} = 2\alpha\hat{c}$.

We have $\beta_{n+1} \leq 2\mu_n\beta_n = \frac{2aw_n}{D\eta}\beta_n$, by (34). Using again the inequality $\frac{w_n}{\eta} \leq \hat{c}\beta_n$, which holds for all $n \geq N$, and letting $\hat{K} = \frac{2a\hat{c}}{D}$, we have $\beta_{n+1} \leq \hat{K}\beta_n^2$. For $\beta_N < \frac{1}{\sqrt{2\hat{K}}}$, this implies that $\beta_{n+1} \leq \frac{\beta_n}{\sqrt{2}}$ and, hence, that

$$\beta_n \leq \beta_N(\sqrt{2})^{-(n-N)}.$$

Therefore,

$$\sum_{n \geq N} \beta_n^2 \leq \sum_{n \geq N} 2^{-(n-N)} \beta_N^2 = 2\beta_N^2.$$

We then obtain

$$u_H(\bar{C}_H) - u_H(E_H(R_N)) \leq 2\tilde{K}\eta\beta_N^2 \leq 2\tilde{K}\eta^{1+2d}, \quad (44)$$

where the last inequality comes from the fact that $\beta_N \leq \eta^d$.

To conclude, note that $\bar{w}_N - w_N = \max\{e_k : k \leq N\} - e_N = \max\{u_H(E_H(k)) : k \leq N\} - u_H(E_H(R_N))$. Since $\max\{u_H(E_H(k)) : k \leq N\} \leq u_H(\bar{C}_H)$, by an argument that is similar to the proof of Lemma 14), (44) yields the result. ■

References

- ABREU, D., GUL, F. “Bargaining and Reputation,” *Econometrica*, Vol. 68, pp. 85–117.
- ATAKAN, A., EKMEKCI, M. (2012) “Reputation in Long-Run Relationships,” *Review of Economic Studies*, Vol. 79, pp. 451–480.
- AUSUBEL, L., DENECKERE, R. (1989) “Reputation in bargaining and durable goods monopoly,” *Econometrica*, Vol. 57, pp. 511–531.
- BATTAGLINI, M. (2007) “Optimality and Renegotiation in Dynamic Contracting,” *Games and Economic Behavior*, Vol. 60, pp. 213–246.
- BEAUDRY, P., POITEVIN, M. (1993) “Signalling and Renegotiation in Contractual Relationships,” *Econometrica*, Vol. 61, pp. 745–782.
- BERNHEIM, B.D., RAY, D. (1989) “Collective Dynamic Consistency in Repeated Games,” *Games and Economic Behavior*, Vol. 1, pp. 295–326.
- BESTER, L., STRAUSS, R. (2001) “Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case,” *Econometrica*, Vol. 69, pp. 1077–1098.
- BRENNAN, J., WATSON, J. (2013) “The Renegotiation-Proofness Principle and Costly Renegotiation,” Working Paper, UCSD.
- BOARD, S., PYCIA, M. (2013) “Outside Options and the Failure of the Coase Conjecture,” *Forthcoming, American Economic Review*.
- COASE, R. H. (1972) “Durability and monopoly,” *Journal of Law and Economics*, Vol. 15, pp. 143–149.
- CRIPPS, M., DEKEL, E., AND W. PESENDORFER (2005) “Reputation with Equal Discounting in Repeated Games with Strictly Conflicting Interests,” *Journal of Economic Theory*, Vol. 121, pp. 259–272.
- DENECKERE, R., LIANG, M-Y (2006) “Bargaining with Interdependent Values,” *Econometrica*, Vol. 74, pp. 1309–1364.
- DEWATRIPONT, M. (1989) “Renegotiation and Information Revelation over Time: The Case of Optimal Labor Contracts,” *Quarterly Journal of Economics*, Vol. 104, pp. 589–619.
- FARRELL, J., MASKIN, E. (1989) “Renegotiation in Repeated Games,” *Games and Economic Behavior*, Vol. 1, pp. 327–360.
- FUCHS, W., SKRZYPACZ, A. (2010) “Bargaining with Arrival of New Traders,” *American Economic Review*, Vol. 100, pp. 802–836.
- FUDENBERG, D., LEVINE, D., AND J. TIROLE (1985) “Infinite Horizon Models of Bargaining with One-Sided Incomplete Information,” in *Game Theoretic Models of Bargaining*, ed. by A. Roth. Cambridge University Press, Cambridge, UK.

- FUDENBERG, D., LEVINE, D. (1989) "Reputation and Equilibrium Selection in Games with a Patient Player," *Econometrica*, Vol. 57, pp. 759–778.
- FUDENBERG, D., TIROLE, J. (1990) "Moral Hazard and Renegotiation in Agency Contracts," *Econometrica*, Vol. 58, No. 6, pp. 1279–1319.
- FUDENBERG, D., TIROLE, J. (1991) *Game Theory*, MIT Press.
- GROMB D. (1994) "Renegotiation in Debt Contracts," Working Paper, London Business School.
- GUL, F., SONNENSCHN, L., AND R. WILSON (1986), "Foundations of Dynamic Monopoly and the Coase Conjecture," *Journal of Economic Theory*, Vol. 39, pp. 155–190.
- HARRIS, C. (1985) "Existence and Characterization of Perfect Equilibrium in Games of Perfect Information," *Econometrica*, Vol. 53, No. 3, pp. 613–628.
- HART, O., TIROLE, J. (1988) "Contract Renegotiation and Coasian Dynamics," *Review of Economic Studies*, Vol. 55, No. 4, pp. 509–540.
- LAFFONT, J.-J., TIROLE, J. (1990) "Adverse Selection and Renegotiation in Procurement," *Review of Economic Studies*, Vol. 57, No. 4, pp. 597–625.
- MAESTRI, L. (2015) "Dynamic Contracting under Adverse Selection and Renegotiation," Working Paper, FGV/EPGE.
- MASKIN, E., TIROLE, J. (1992) "The Principal-Agent Relationship with an Informed Principal, II: Common Values," *Econometrica*, Vol. 60, pp. 1–42.
- MATTHEWS, S. (1995) "Renegotiation of Sales Contracts," *Econometrica*, Vol. 63, pp. 567–590.
- RUBINSTEIN, A. (1982) "Perfect Equilibrium in a Bargaining Model," *Econometrica*, Vol. 50, pp. 97–110.
- RUBINSTEIN, A. (1985) "A Bargaining Model with Incomplete Information about Time Preferences," *Econometrica*, Vol. 53, pp. 1151–1172.
- RUBINSTEIN, A., WOLINSKY, A. (1992) "Renegotiation-Proof Implementation and Time Preferences," *American Economic Review*, Vol. 82, pp. 600–614.
- SCHMIDT, K. (1993) "Reputation and Equilibrium Characterization in Repeated Games with Conflicting Interests," *Econometrica*, Vol. 61, pp. 325–351.
- SKRETA, V. (2006) "Sequential Optimal Mechanisms," *Review of Economic Studies*, Vol. 73, pp. 1085–1111.
- SOBEL, J., TAKAHASHI, I. (1983) "A Multistage Model of Bargaining," *Review of Economic Studies*, Vol. 50, pp. 411–426.
- STRULOVICI, B. (2011) "Renegotiation-Proof Contracts with Moral Hazard and Persistent Private Information," Working Paper, Northwestern University.
- STRULOVICI, B. (2013) "Renegotiation and Persistent Information," Working Paper, Northwestern University.

WANG, G.H. (1998) “Bargaining over a Menu of Wage Contracts,” *Review of Economic Studies*, Vol. 65, pp. 295–305.

D Inequalities (For Online Appendix)

LEMMA 10 (REGULARITY BOUNDS) *There exist positive constants $\underline{a}, \underline{a}, \underline{b}, \underline{b}$ such that for any $C, \hat{C} \in \mathcal{E}_H$ such that $u_H(C) < u_H(\hat{C})$, we have*

$$\underline{a}(u_H(\hat{C}) - u_H(C)) \leq Q(\hat{C}) - Q(C) \leq \underline{a}(u_H(\hat{C}) - u_H(C)) \quad (45)$$

$$\underline{b}(Q(\hat{D}) - Q(D)) \leq Q(\hat{C}) - Q(C) \leq \underline{b}(Q(\hat{D}) - Q(D)), \quad (46)$$

where D (resp. \hat{D}) is the L -efficient contract that gives H the same utility as C (resp. \hat{C}).

Proof. Consider two contracts C and \hat{C} on \mathcal{E}_H ordered as in the statement of the lemma. The efficiency curve \mathcal{E}_H can be parameterized by a univariate parameter λ such that, letting $C(\lambda) = (x_1(\lambda), x_2(\lambda))$ denote the H -efficient contract corresponding to parameter λ , the map $\lambda \mapsto C(\lambda)$ is continuous, one-to-one, and onto from the parameter set Λ (a compact interval of \mathbb{R}) to \mathcal{E}_H . We can assume without loss that Λ contains $[0, 1]$ and that $C(0) = C$ and $C(1) = \hat{C}$. We choose the parametrization to be regular, i.e., such that $\lambda \mapsto C(\lambda)$, seen as a function from Λ to \mathbb{R}^2 , is smooth and does not go “too slow” or “too fast” along \mathcal{E}_H .⁸³ We have

$$\begin{aligned} Q(\hat{C}) - Q(C) &= \int_0^1 \frac{dQ(x_1(\lambda), x_2(\lambda))}{d\lambda} \cdot dC(\lambda) \\ &= \int_0^1 \left(\frac{\partial Q(C(\lambda))}{\partial x_1} \frac{dx_1}{d\lambda} + \frac{\partial Q(C(\lambda))}{\partial x_2} \frac{dx_2}{d\lambda} \right) d\lambda. \end{aligned}$$

Similarly, we have

$$\begin{aligned} u_H(\hat{C}) - u_H(C) &= \int_0^1 \frac{du_H(x_1(\lambda), x_2(\lambda))}{d\lambda} \cdot dC(\lambda) \\ &= \int_0^1 \left(\frac{\partial u_H(C(\lambda))}{\partial x_1} \frac{dx_1}{d\lambda} + \frac{\partial u_H(C(\lambda))}{\partial x_2} \frac{dx_2}{d\lambda} \right) d\lambda. \end{aligned}$$

By assumption, the partial derivatives of Q and u_H are strictly positive and continuous on the compact domain \mathcal{C} , and hence bounded below away from zero as well as bounded above. Therefore, there exist positive constants $\underline{a} < \underline{a}$ such that $\underline{a} \frac{\partial u_H}{\partial x_i} \leq \frac{\partial Q}{\partial x_i} \leq \underline{a} \frac{\partial u_H}{\partial x_i}$ for $i = 1, 2$. Using these inequalities into the previous integral representations of $Q(\hat{C}) - Q(C)$ and $u_H(\hat{C}) - u_H(C)$ then shows (45).

For the second part of the lemma, consider the parameterizations of \mathcal{E}_H and \mathcal{E}_L for which the parameter corresponds to the utility that each contract gives to H (thus, $u_H(C(\lambda)) = \lambda$), with elements $C(\lambda) = (x_1^H(\lambda), x_2^H(\lambda))$ for \mathcal{E}_H and $D(\lambda) = (x_1^L(\lambda), x_2^L(\lambda))$ for \mathcal{E}_L . Because the partial derivatives of u_H are strictly positive on the compact domain \mathcal{C} and because the curves \mathcal{E}_θ are both nondecreasing in \mathcal{C} , that parameterization is well defined and regular (in the sense of the previous paragraph) for both curves. Consider two contracts C and \hat{C} of \mathcal{E}_H with provide H with utilities $u_H < \hat{u}_H$ and let D and \hat{D} denote the contracts of \mathcal{E}_L corresponding to utilities u_H and \hat{u}_H . Repeating the argument of the previous paragraph, we have

$$Q(\hat{C}) - Q(C) = \int_{u_H}^{\hat{u}_H} \left(\frac{\partial Q(C(\lambda))}{\partial x_1} \frac{dx_1^H}{d\lambda} + \frac{\partial Q(C(\lambda))}{\partial x_2} \frac{dx_2^H}{d\lambda} \right) d\lambda$$

⁸³Formally, this means that the norm of the gradient of the function $\lambda \mapsto C(\lambda)$ is uniformly bounded below and above by strictly positive constants.

and

$$Q(\hat{D}) - Q(D) = \int_{u_H}^{\hat{u}_H} \left(\frac{\partial Q(D(\lambda))}{\partial x_1} \frac{dx_1^L}{d\lambda} + \frac{\partial Q(D(\lambda))}{\partial x_2} \frac{dx_2^L}{d\lambda} \right) d\lambda.$$

Because the parameterizations are regular and the curves are nondecreasing, there must exist positive constants $\underline{x} < \bar{x}$ such that $0 < \underline{x} \max\{dx_1^H/d\lambda, dx_2^H/d\lambda\} \leq \max\{dx_1^L/d\lambda, dx_2^L/d\lambda\} \leq \bar{x} \max\{dx_1^H/d\lambda, dx_2^H/d\lambda\}$. Moreover, since Q has strictly positive derivatives, bounded below away from zero and bounded above, there also exist positive constants $\underline{q} < \bar{q}$ such that $\underline{q} \partial Q(C(\lambda))/\partial x_i \leq \partial Q(D(\lambda))/\partial x_i \leq \bar{q} \partial Q(C(\lambda))/\partial x_i$ for all $\lambda \in [u_H, \hat{u}_H]$ and $i = 1, 2$. Combining these inequalities with the previous integral representations implies, as is easily checked, that there exist positive constants $\underline{b} < b$ such that

$$\underline{b}(Q(\hat{D}) - Q(D)) \leq Q(\hat{C}) - Q(C) \leq b(Q(\hat{D}) - Q(D)),$$

which concludes the proof. \blacksquare

For the next result, let Q_θ denote P's expected continuation cost at the beginning of round n , conditional on facing type θ . (We omit dependence on n for simplicity).

LEMMA 11 (INCENTIVE BOUNDS) *Given any PBE and choice sequence $\{R_n\}$, there exist positive constants α, γ, \hat{b} , and b , such that*

$$Q_L \leq Q(E_L(R_n)) + \frac{\beta_n}{(1 - \beta_n)} a w_n, \quad (47)$$

$$u_H(E_H(R_{n+1})) - u_H(E_H(R_n)) \leq \frac{\alpha \beta_n}{1 - \beta_n} w_n, \quad (48)$$

$$w_{n+1} \leq w_n \left(1 + \frac{\alpha \beta_n}{1 - \beta_n} \right), \quad (49)$$

$$u_L(R_n) - u_L(R_{n+1}) \leq \gamma \beta_{n+1} w_{n+1}, \quad (50)$$

$$u_H(E_H(R_{n+1})) - u_H(E_H(R_n)) \geq -\hat{b} \beta_{n+1} w_{n+1}, \quad (51)$$

$$w_{n+1}(1 - b \beta_{n+1}) \geq w_n - \eta y_n. \quad (52)$$

Proof. Lemma 2 implies that

$$\beta_n Q_H + (1 - \beta_n) Q_L \leq \beta_n Q(E_H(R_n)) + (1 - \beta_n) Q(E_L(R_n)).$$

Moreover, Q_H is bounded below by the cost of the H -efficient contract $C_H(n)$ that provides utility $u_H(n)$ to H , since that is the cheapest way of providing H with his continuation utility (by convexity of the cost function Q). This implies that $Q_L \leq Q(E_L(R_n)) + \frac{\beta_n}{1 - \beta_n} (Q(E_H(R_n)) - Q(C_H(n)))$. The contracts $E_H(R_n)$ and $C_H(n)$ both lie on \mathcal{E}_H . Equation (45) implies that $Q(E_H(R_n)) - Q(C_H(n)) \leq a(u_H(E_H(R_n)) - u_H(n)) = a w_n$. This shows (47).

From (47), R_{n+1} cannot give L a utility greater than the L -efficient contract that costs $Q(E_L(R_n)) + \frac{\alpha \beta_n}{1 - \beta_n} w_n$. This implies that $Q(E_L(R_{n+1})) - Q(E_L(R_n))$ is bounded above by $\frac{\alpha \beta_n}{1 - \beta_n} w_n$. Combining this with (46) yields⁸⁴

$$Q(E_H(R_{n+1})) - Q(E_H(R_n)) \leq \frac{\alpha b \beta_n}{1 - \beta_n} w_n.$$

⁸⁴Equation 46 applies if $Q(E_H(R_{n+1})) - Q(E_H(R_n)) \geq 0$. In the opposite case, the inequality holds trivially since the left-hand side is negative and the right-hand side is positive.

This, along with the first part of (45) yields (48). We have

$$\begin{aligned} w_{n+1} &= u_H(E_H(R_{n+1})) - u_H(n+1) = [u_H(E_H(R_{n+1})) - u_H(E_H(R_n))] + u_H(E_H(R_n)) - u_H(n+1) \\ &\leq [u_H(E_H(R_{n+1})) - u_H(E_H(R_n))] + u_H(E_H(R_n)) - u_H(n) \\ &\leq w_n \left(\frac{\alpha\beta_n}{1-\beta_n} + 1 \right) \end{aligned}$$

where the first inequality comes from the monotonicity of $u_H(n)$ in n , and the second inequality comes from (48). This shows (49).

Because L can hold on forever to R_n , his continuation utility $u_L(n)$ is bounded below by $u_L(R_n)$. At round $n+1$, P 's expected cost conditional on facing L is bounded above by $Q(E_L(R_{n+1})) + \frac{\beta_{n+1}}{1-\beta_{n+1}}aw_{n+1}$, from (47) applied to round $n+1$. By the same argument that yielded (45), there exists $\alpha_L > 0$ such that $u_L(E) - u_L(E') \leq \alpha_L(Q(E) - Q(E'))$ for all $E, E' \in \mathcal{E}_L$. Therefore, the highest utility which may be achieved at that cost is bounded above by $u_L(R_{n+1}) + \hat{a}\beta_{n+1}/(1-\beta_{n+1})w_{n+1}$, for some proportionality constant \hat{a} , and

$$u_L(R_n) \leq u_L(n) \leq u_L(n+1) \leq u_L(R_{n+1}) + \hat{a}\beta_{n+1}/(1-\beta_{n+1})w_{n+1},$$

which yields (50).

In general, $u_H(E_H(R_{n+1}) - u_H(E_H(R_n)))$ may be negative. To provide a lower bound for that case, we use the relations

$$u_L(E_L(R_n)) - u_L(E_L(R_{n+1})) = u_L(R_n) - u_L(R_{n+1}) \leq \hat{a}\beta_{n+1}/(1-\beta_{n+1})w_{n+1}, \quad (53)$$

where the equality simply comes from the definition of $E_L(R_n)$ and $E_L(R_{n+1})$ and the inequality comes from (50). Replicating the argument that yielded (45), but using \mathcal{E}_L instead of \mathcal{E}_H and the pair of functions (u_H, u_L) instead of (u_H, Q) , we get

$$u_H(E_L(R_n)) - u_H(E_L(R_{n+1})) \leq \tilde{\alpha} [u_L(E_L(R_n)) - u_L(E_L(R_{n+1}))]$$

for some positive constant $\tilde{\alpha}$. Since $u_H(E_L(R_n)) = u_H(E_H(R_n))$ and $u_H(E_L(R_{n+1})) = u_H(E_H(R_{n+1}))$, the previous equation combined with (53) proves (51).⁸⁵

For the last equation, subtracting $u_H(E_H(R_n))$ from (19) and rearranging (recalling that $w_n = u_H(E_H(R_n)) - u_H(n)$) leads, along any choice sequence, to

$$w_{n+1} = w_n - \eta y_n + \eta w_{n+1} + (1-\eta)(u_H(E_H(R_{n+1})) - u_H(E_H(R_n))). \quad (54)$$

Combining this with (51) yields

$$w_{n+1} - w_n \geq \eta w_{n+1} - \eta y_n - b\beta_{n+1}w_{n+1}$$

and hence (52). ■

⁸⁵Intuitively, that equation comes from two observations. First, L 's utility from the current contract, R_n , cannot decrease by too much between consecutive rounds. Indeed, recall that β_{n+1} is the probability of facing H in round $n+1$, while w_{n+1} is a measure of the maximum rent that P can extract from H at round $n+1$. If the product $\beta_{n+1}w_{n+1}$ is small, it means that, comes round $n+1$, P has very little incentive to extract rents from H , which implies, intuitively, that his continuation strategy must be similar to what he would do if he only faced L , namely to jump to the L efficient contract $E_L(R_{n+1})$, which gives L utility $u_L(R_{n+1})$. Anticipating this, however, L is willing to forgo the current contract R_n only if R_{n+1} gives him a utility that is not much lower than R_n . The second observation is that the H -efficient contracts $E_H(R_n)$ and $E_H(R_{n+1})$ are constructed based on the utility that L gets from R_n and R_{n+1} . Using a suitable Lipschitz relation between utility differences then yields (35).

LEMMA 12 (GEOMETRIC BOUND) *There exists $\underline{q} > 0$ such that for any C on \mathcal{E}_L and R in \mathcal{H} such that $u_L(R) = u_L(C)$,*

$$Q(R) - Q(C) \geq \underline{q}(u_H(C) - u_H(R))^2.$$

Before proving this result, it is useful to give some intuition. Given isoutility curve of L , moving along this curve affects both H 's utility and P 's cost. The idea is to bound the cost difference for any two points on the curve in terms of the corresponding differences in H 's utility. At any regular point, one would expect to obtain a Lipschitz bound, just like the ones obtained in Lemma 10. However, because C is on L 's efficiency curve, this efficiency curve is tangent to P 's isocost curve at C . Therefore, moving slightly away from C along L 's isoutility curve has only a second-order effect on the cost, while it has a first-order effect on H 's utility (since H and L 's isoutility curves satisfy the single crossing property and are thus not tangent). This yields (locally) the sharper bound provided by Lemma 12, and the global bound can be obtained by choosing the bounding constant \underline{q} appropriately. *Proof.* Fix some $C \in \mathcal{E}_L$ and consider the referential centered at C whose x -axis is the common tangent of u_L and Q at C , oriented towards \mathcal{H} , and whose y -axis is the normal vector pointing northeast in \mathcal{C} . The components of a contract, in this referential, are denoted x_t and x_n , respectively, with C being the origin. We parameterize the contract set $\mathcal{U}_L(C) = \{R \in \mathcal{C} : u_L(\tilde{C}) = u_L(C)\}$ in terms of x_t : $\{C(x_t) = (x_t, x_n(x_t))\}$. With this parameterization, we have $C(0) = C$ and $C(x_t) \in \mathcal{H}$ if and only if $x_t \geq 0$.⁸⁶

Let $Q(x_t) = Q(C(x_t))$ and $u_H(x_t) = u_H(C(x_t))$ denote the cost and utility of H along $\mathcal{U}_L(C)$, as a function of the parameter x_t . By L -efficiency of C , we have $Q'(0) = 0$.⁸⁷ Because u_L is concave and Q is convex, the iso-utility curve of u_L going through C is convex and corresponds to positive values of x_n , while the isocost curve going through C is concave and corresponds to negative values of x_n . Moreover, by assumption at least one of these curves has a nonzero curvature at C . In the (x_t, x_n) space, this means that either $d^2 u_L / dx_t^2 > 0$ or $d^2 Q / dx_t^2 < 0$. We wish to show the existence of a constant $\hat{q} > 0$ such that $Q(x_t) - Q(0) \geq \hat{q}x_t^2$ for x_t in a right neighborhood of 0. Suppose first that $d^2 u_L / dx_t^2 > 0$. This implies that $x_n(x_t) \geq q_x x_t^2$ for some $q_x > 0$ and x_t in a neighborhood of zero. Therefore, $Q(x_t) \geq Q(0) + q_x \|\nabla Q(C)\| x_t^2$ for that neighborhood. Now suppose that $d^2 Q / dx_t^2 < 0$. In that case, let $D(x_t)$ denote the contract of the isocost curve with x -value x_t in the new referential (hence, just below $C(x_t)$ in the new referential), so that $Q(D(x_t)) = Q(C)$ for all x_t . By tangency of the curves, we have $\|C(x_t) - D(x_t)\| = o(x_t)$. Moreover, $Q(x_t) = Q(C(x_t)) = Q(D(x_t) + \nabla Q(D(x_t)) \cdot (C(x_t) - D(x_t)) + O(\|C(x_t) - D(x_t)\|^2))$ by a standard Taylor expansion. Finally, for x_t in a neighborhood of 0, $\nabla Q(D(x_t)) = \nabla Q(C) + O(\|D(x_t) - C\|) = \nabla Q(C) + o(x_t)$. Combining this, we get $Q(x_t) = Q(D(x_t)) + \nabla Q(C) \cdot (C(x_t) - D(x_t)) + o(x_t)(\|C(x_t) - D(x_t)\| + \|D(x_t) - C\|)$. Since $d^2 Q / dx_t^2 < 0$, the y -value of $D(x_t)$ in the new referential satisfies $x_n^D(x_t) \leq -\hat{q}_x x_t^2$ for some $\hat{q}_x > 0$. Hence, $\nabla Q(C) \cdot (C(x_t) - D(x_t)) \geq \tilde{q}_x (x_n(x_t) - x_n^D(x_t)) \geq \tilde{q}_x x_t^2$ for some positive constants \tilde{q}_x, \hat{q}_x . Combining all this implies that

$$Q(x_t) \geq Q(C) + \hat{q}x_t^2 + o(x_t^2), \tag{55}$$

proving the result for that case too.

⁸⁶The parameterization is well defined, because $\mathcal{U}_L(C)$ can only have one point for each x_t , by strict monotonicity of u_L in the original coordinates (x_1, x_2) and the fact that increasing x_n corresponds to increasing both x_1 and x_2 and at least one of these increases is strict, since the normal vector defining x_n points northeastwards.

⁸⁷Formally, we have $Q'(x_t) = \frac{\partial Q}{\partial x_1} \frac{dx_1}{dx_t} + \frac{\partial Q}{\partial x_2} \frac{dx_2}{dx_t}$. Since C is L -efficient, Q and u_L are tangent at C . This implies that the tangent vector $(dx_1/dx_t, dx_2/dx_t)$ is orthogonal to the normal vector $(\partial Q/\partial x_1, \partial Q/\partial x_2)$ at C .

By compactness and convexity of $\mathcal{U}_L(C)$, moreover, \hat{q} may be chosen small enough so that the inequality

$$Q(x_t) - Q(0) \geq \hat{q}x_t^2$$

holds for all nonnegative x_t . Since u_H has bounded derivatives, there must exist $\bar{u} > 0$ such that $|u_H(x_t) - u_H(0)| \leq \bar{u}x_t$ (the single-crossing property between u_H and u_L imply that $u_H(x_t) \leq u_H(0)$ for all $x_t \geq 0$ and that $\nabla u_H(C) \cdot (C(x_t) - C) \neq 0$ for x_t in a neighborhood of 0). Combining these inequalities, there exists $\underline{q}(C) > 0$ such that

$$Q(C_\lambda) - Q(C) \geq \underline{q}(C) (u_H(C) - u_H(C_\lambda))^2.$$

Moreover, $\underline{q}(C)$ can clearly be chosen to vary continuously in $C \in \mathcal{E}_L$.⁸⁸ By compactness of \mathcal{E}_L , $\underline{q} = \min_{C \in \mathcal{E}_L} \underline{q}(C)$ is strictly positive and yields the desired inequality. ■

LEMMA 13 *There exist positive constants k_2 and k_3 such that*

$$y_n^2 \leq k_2[Q(R_{n+1}) - Q(E_L(R_n))] + k_3(\max\{(\beta_n w_n / (1 - \beta_n))^2, (\beta_{n+1} w_{n+1})^2\} + \beta_{n+1} w_{n+1}) \quad (56)$$

Proof. We have

$$\begin{aligned} y_n^2 &= [(u_H(E_H(R_n)) - u_H(E_H(R_{n+1}))) + (u_H(E_H(R_{n+1})) - u_H(R_{n+1}))]^2 \\ &\leq 2[u_H(E_H(R_n)) - u_H(E_H(R_{n+1}))]^2 + 2[u_H(E_H(R_{n+1})) - u_H(R_{n+1})]^2 \\ &\leq k_1(\max\{\beta_n w_n / (1 - \beta_n), \beta_{n+1} w_{n+1}\})^2 + 2[u_H(E_H(R_{n+1})) - u_H(R_{n+1})]^2 \\ &\leq k_1 \max\{(\beta_n w_n / (1 - \beta_n))^2, (\beta_{n+1} w_{n+1})^2\} + k_2[Q(R_{n+1}) - Q(E_L(R_{n+1}))] \\ &= k_1 \max\{(\beta_n w_n / (1 - \beta_n))^2, (\beta_{n+1} w_{n+1})^2\} + k_2[Q(E_L(R_n)) - Q(E_L(R_{n+1}))] + k_2[Q(R_{n+1}) - Q(E_L(R_n))]. \end{aligned}$$

The first inequality is standard ($(a+b)^2 \leq 2a^2 + 2b^2$). The second inequality comes from (48) and (51), which taken together imply an upper bound on $|u_H(E_H(R_n)) - u_H(E_H(R_{n+1}))|$. The third inequality comes from the equality $u_H(E_H(R_{n+1})) = u_H(E_L(R_{n+1}))$ and Lemma 12 applied to the contracts $C = E_L(R_{n+1})$ and $R = R_{n+1}$. The difference $Q(E_L(R_n)) - Q(E_L(R_{n+1}))$ is bounded above in proportion to $u_L(R_n) - u_L(R_{n+1})$ (by a simple transposition to \mathcal{E}_L of the proof of (45)), and that latter difference is bounded above by $\gamma\beta_{n+1}w_{n+1}$, from (50). This shows that

$$y_n^2 \leq k_2[Q(R_{n+1}) - Q(E_L(R_n))] + k_3(\max\{(\beta_n w_n / (1 - \beta_n))^2, (\beta_{n+1} w_{n+1})^2\} + \beta_{n+1} w_{n+1})$$

which yields the result. ■

E Proofs for Parts I and II (For Online Appendix)

PROOF OF LEMMA 4 Fix any choice sequence and let n_0 denote the first round along that sequence such that $\bar{w}_{n_0} \leq \varepsilon$. By construction, $\bar{w}_{n_0-1} > \varepsilon$. From Lemma 5 (whose proof, in the main text, is independent of this lemma), we have $u_H(n_0) \leq u_H(n_0 - 1) + \eta\Delta_H$. Therefore,

$$\bar{w}_{n_0} \geq \bar{w}_{n_0-1} + u_H(n_0 - 1) - u_H(n_0) \geq \varepsilon - \eta\Delta_H.$$

Since we can always select a choice sequence along which β_n is weakly decreasing, we also get $\beta_{n_0} \leq \beta_0$. ■

⁸⁸Indeed, all the constants involved in the previous steps are based on the curvature of the isoutility and isocost curves at C , which only involve the second derivative of the utility and cost functions at C . These functions were assumed to be C^2 over \mathcal{C} .

LEMMA 14 *i) For any round n_0 , $\tilde{\epsilon} > 0$, and choice sequence, there exists a round $n > n_0$ such that $u_H(R_n) \geq \max\{u_H(E_H(R_m)) : m \leq n_0\} - \tilde{\epsilon}$. ii) The augmented rent index $\bar{w}_n = \max\{u_H(E_H(R_m)) : m \leq n\} - u_H(n)$ converges to zero as n goes to infinity, along any choice sequence.*

Proof. i) Fix $\tilde{\epsilon} > 0$. Proposition 2 guarantees that, along any choice sequence, R_n converges to an L -efficient \bar{C}_L . Continuity of $u_H(\cdot)$ implies that there exists a round \tilde{n} such that $u_H(R_n) \geq u_H(\bar{C}_L) - \tilde{\epsilon}$ for all $n \geq \tilde{n}$. Therefore, it suffices to show that $u_H(\bar{C}_L) \geq \max\{u_H(E_H(R_m)) : m \leq n_0\}$ for all n_0 . Equivalently, we must show that $u_H(\bar{C}_L) \geq \max\{u_H(E_L(R_m)) : m \leq n_0\}$ for all n_0 since, by construction, $E_H(R)$ and $E_L(R)$ give the same utility to H for any $R \in \mathcal{H}$. For contracts C, C' on the L -efficiency curve \mathcal{E}_L , $u_H(C) \leq u_H(C')$ if and only if $u_L(C) \leq u_L(C')$. Therefore, it suffices to show that $u_L(\bar{C}_L) \geq \max_{m \in \mathbb{N}}\{u_L(E_L(R_m))\}$. By construction, $u_L(E_L(R)) = u_L(R)$ for all $R \in \mathcal{H}$, since $E_L(R)$ is the L -efficient contract that gives L the same utility as R . Therefore, we have reduced the problem to showing that

$$u_L(\bar{C}_L) \geq \max_{m \in \mathbb{N}}\{u_L(R_m)\}.$$

We recall that for all n , $u_L(n) \geq u_L(R_n)$ since holding on to R_n is always a feasible strategy for L , and that $u_L(n)$ is nondecreasing in n for all choice sequences (see Lemma 5; the argument there also applies to L). Since R_n converges to \bar{C}_L , $u_L(n)$ must converge to $u_L(\bar{C}_L)$. Finally, because $u_L(n)$ is nondecreasing, we get

$$u_L(R_n) \leq u_L(n) \leq u_L(\bar{C}_L),$$

which concludes the proof of i). To prove ii), it suffices to notice that $\max\{u_H(E_H(R_m)) : m \leq n\}$ and $u_H(R_n)$ both converge to $u_H(\bar{C}_L)$, by the previous argument. ■

LEMMA 15 *There exists a pushdown sequence at Block 1.*

Proof. Let $\mu^\theta(\{\tilde{R}_n\})$ denote the probability, conditional on facing type θ , of observing choice sequence $\{\tilde{R}_n\}$ until \hat{u}_1 is reached. By definition, summing over all choice sequence with elements in \mathcal{H} and truncated at the first round at which H 's continuation utility reaches \hat{u}_1 , we have $\sum_{\{\tilde{R}_n\}} \mu^H(\{\tilde{R}_n\}) = \mu_0$. Because L always chooses contracts in \mathcal{H} , we also have $\sum_{\{\tilde{R}_n\}} \mu^L(\{\tilde{R}_n\}) = 1$. These two equations immediately imply that there exists a choice sequence $\{R_n^0\}$ such that $\mu^H(\{R_n^0\})/\mu^L(\{R_n^0\}) \leq \mu_0$. Conditional on observing that choice sequence, the posterior is given by Bayesian updating

$$\hat{\beta}_1 = \frac{\mu^H(\{R_n^0\})\hat{\beta}_0}{\mu^H(\{R_n^0\})\hat{\beta}_0 + \mu^L(\{R_n^0\})(1 - \hat{\beta}_0)}.$$

Dividing by $\mu^L(\{R_n^0\})$ and using that $\mu^H(\{R_n^0\})/\mu^L(\{R_n^0\}) \leq \mu_0$ yields the result. ■

Let $u_H = u_H(R_0)$ and, for any $\tilde{\epsilon} \geq 0$,

$$D(\tilde{\epsilon}) = \inf\{Q(C) - Q(E) : C \in \mathcal{H}, E \in \mathcal{E}_H : u_H \leq u_H(E) \leq u_H(E_H(C)) + \tilde{\epsilon}\}. \quad (57)$$

$D(\tilde{\epsilon})$ is nonincreasing in $\tilde{\epsilon}$, as a higher $\tilde{\epsilon}$ merely increases the set of (C, E) pairs over which the objective is minimized. Because R_0 is regular, the contracts C arising in (57) are bounded away from \mathcal{E}_H for $\tilde{\epsilon}$ small enough, and this implies that $D(\tilde{\epsilon})$ is strictly positive for $\tilde{\epsilon}$ small enough. Intuitively, C and E must provide almost the same utility to H , and E is a strictly cheaper way than C of doing so. For such values of $\tilde{\epsilon}$, $D(\tilde{\epsilon})$ defines a lower bound on the inefficiency of contracts in \mathcal{H} conditional on facing H .

LEMMA 16 *If at the beginning of any block k , $w_{n_{k-1}} \leq \varepsilon$, then for all rounds n of block k ,*

$$Q(R_n) \geq Q(E_H(R_{n_{k-1}})) + D(\varepsilon)$$

Proof. Let C denote the L -efficient contract that gives H utility $u_H(n_{k-1})$. Since $u_H(n)$ is nondecreasing, we have for any round n of block k , $u_H(n_{k-1}) \leq u_H(n)$. From part iv) of Proposition 1, we have $u_H(n) \leq u_H(E_H(R_n))$. This implies that R_n must cost weakly more than C ; otherwise, we have $u_L(E_L(R_n)) < u_L(C)$ and hence $u_H(E_H(R_n)) < u_H(E_H(C))$, which implies that $u_H(n) \leq u_H(E_H(R_n)) < u_H(E_H(C)) = u_H(n_{k-1})$, a contradiction. So $Q(R_n) \geq Q(C)$. By assumption, we have $u_H(E_H(R_{n_{k-1}})) - u_H(C) = w_{n_{k-1}} \leq \varepsilon$. By definition of $D(\varepsilon)$, this implies that $Q(C) \geq Q(E_H(R_{n_{k-1}})) + D(\varepsilon)$. Since $Q(R_n) \geq Q(C)$, this proves the lemma. \blacksquare

We will consider ε such that $D(2\varepsilon) > 0$, we let $D = D(2\varepsilon)$ denote the lower bound on the loss that is used throughout the proof.

REMARK 3 *In principle, one could reach a block k for which w_{k-1} , and hence \hat{w}_{k-1} , is greater than 2ε , which would imply that the lower bound D on the loss is not guaranteed to hold for that block. If that is the case, however, Lemma 4 guarantees that one can find a later round n for which \bar{w}_n lies in $(\varepsilon/2, \varepsilon)$, and one can restart the analysis from that round (i.e., this is our new “ n_0 ”). Moreover, $\hat{\beta}_{k-1} \leq \beta_0$, so the two conclusions of Lemma 4 hold. Re-starting Part I from the new round n_0 , one may encounter a block for which this problem arises again, in which case one re-initialize the analysis again, starting from a yet later round. Since \bar{w}_n converges to zero along any choice sequence as n goes to infinity, by Lemma 14, there can only be finitely many such initializations: there must exist a round n_0 such that i) $\bar{w}_{n_0} \in (\varepsilon/2, \varepsilon)$, ii) $\beta_{n_0} \leq \beta_0$, and iii) \hat{w}_k remains below 2ε for all blocks constructed from n_0 .⁸⁹*

PROOF OF LEMMA 6

Consider any round n and contract R_{n+1} in $M_n \cup \{R_n\}$. If $\mu_n^L(R_{n+1}) \geq \bar{\varepsilon}$, then $\mu_n(R_{n+1}) \leq \frac{1}{\bar{\varepsilon}}$ and, hence, $\beta_{n+1} \leq \frac{\beta_n}{\bar{\varepsilon}}$, since $\beta_{n+1} \sim \mu_n(R_{n+1})\beta_n$ by Bayesian updating (see (38); the term $\beta_{n+1}\beta_n$ can be neglected). The set of contracts R_{n+1} for which $\mu_n^L(R_{n+1}) < \bar{\varepsilon}$ has probability at most $G\bar{\varepsilon}$, where G is the upper bound on the size of the menu. Therefore, with probability at least $1 - G\bar{\varepsilon}$,

$$\beta_{n+1} \leq \frac{\beta_n}{\bar{\varepsilon}}$$

At round \bar{n} , we have $\beta_{\bar{n}} \leq \eta^d$ and $w_{\bar{n}} \leq \bar{w}\eta$. Therefore, $\beta_{\bar{n}+1} \leq \eta^d/\bar{\varepsilon}$ with probability at least $1 - G\bar{\varepsilon}$. From (49), $w_{n+1} \leq w_n \left(1 + \frac{\alpha\beta_n}{1-\beta_n}\right)$. Therefore, we also have $w_{\bar{n}+1} \leq k_1\eta$ for some constant k_1 . This implies that with probability at least $1 - G\bar{\varepsilon}$, the lower bound $D = D(2\varepsilon)$ on the loss is valid for round $\bar{n} + 1$, because $w_{n+1} \leq 2\varepsilon$. The previous reasoning can be applied by induction to rounds $n = \bar{n}, \dots, \bar{n} + \bar{N} - 1$. It implies that with probability $1 - k(\bar{N})\bar{\varepsilon}$, we have

$$\beta_n \leq (\bar{\varepsilon})^{-\bar{N}} \bar{\beta} \eta^d \tag{58}$$

$$w_n \leq \bar{W}(\bar{N})\eta \tag{59}$$

for all $n \in \{\bar{n}, \dots, \bar{n} + \bar{N}\}$, for some constants $k(\bar{N})$ and $\bar{W}(\bar{N})$ independent of $\bar{\varepsilon}$ and η .

⁸⁹Indeed, the contrapositive is that there exists a choice sequence such that for all n , there is a round $n' > n$ for which $w_{n'} \geq 2\varepsilon$, which clearly contradicts the convergence of w_n to zero along all choice sequences.

Consider any choice sequence such that β_n and w_n satisfy the above inequalities throughout the block, which occur with probability $1 - k(\bar{N})\bar{\varepsilon}$. There remains to show the claim that $y_n = O(\eta^{d/4})$ throughout the block for such sequences. We begin by showing the result for round $n = \bar{n}$. The first step is to show that $Q(R_{n+1}) - Q(E_L(R_n))$ must be of order $O\left(\frac{\beta_n}{\mu_n^L(R_{n+1})}\right)$ for that round. If each term in the sum entering P's IC constraint (26) is nonnegative, this result comes from the inequality⁹⁰

$$w_n a \beta_n \geq \mu_n^L(R_{n+1}) [\beta_n \mu_n(R_{n+1}) \eta D + (1 - \beta_n) \eta (Q(R_{n+1}) - Q(E_L(R_n)))]$$

which implies that

$$Q(R_{n+1}) - Q(E_L(R_n)) \leq \frac{a w_n}{\eta(1 - \beta_n)} \frac{\beta_n}{\mu_n^L(R_{n+1})}. \quad (60)$$

In general, while some terms

$$\mu_n^L(R_{n+1})(1 - \beta_n) \eta (Q(R_{n+1}) - Q(E_L(R_n))) \quad (61)$$

involved in the sum of (26) may be negative, they can only be very slightly so: indeed, we have

$$Q(R_{n+1}) - Q(E_L(R_n)) \geq Q(E_L(R_{n+1})) - Q(E_L(R_n)) \geq -k \beta_{n+1} w_{n+1}, \quad (62)$$

where the second inequality comes (50) of Lemma 11. Moreover, $w_{n+1} \leq w_n \left(1 + \frac{\alpha \beta_n}{1 - \beta_n}\right)$, from (49), and $\mu_n^L(R_{n+1}) \beta_{n+1}$ is of order β_n . Therefore, the lower bound of (60) is of order $w_n \beta_n$, and hence each negative term which may arise in (26) are of order $\eta w_n \beta_n$. Since there are at most G of them, we conclude that (60) holds up to a term of order $\eta \beta_n$, which is negligible compared to the first term.

From (60), we see that for all R_{n+1} such that $\mu_n^L(R_{n+1}) \geq \sqrt{\beta_n}$, the difference $Q(R_{n+1}) - Q(E_L(R_n))$ is at most of order $\sqrt{\beta_n}$. By Lemma 13 of the Appendix, this implies that y_n is $O(\eta^{d/4})$: Indeed, y_n^2 is bounded above by terms proportional to $Q(R_{n+1}) - Q(E_L(R_n))$ and a term proportional to $\max\{(\beta_n w_n / (1 - \beta_n))^2, (\beta_{n+1} w_{n+1})^2\} + \beta_{n+1} w_{n+1}$. The first term is of order $\sqrt{\beta_n} = O(\eta^{d/2})$, while the latter is of order $\beta_n \eta$, and is thus negligible compared to the first. Moreover, the set of contracts R_{n+1} for which $\mu_n^L(R_{n+1}) < \sqrt{\beta_n}$, is negligible: it arises with probability at most $G\sqrt{\beta_n}$. Since $\sqrt{\beta_n} = O(\eta^{d/2})$ is small compared to $\bar{\varepsilon}$, for η small enough, we conclude that with probability $1 - O(\bar{\varepsilon})$, $\beta_n = O(\eta^d)$, $w_n = O(\eta)$ and $y_n = O(\eta^{d/4})$ for round \bar{n} and, by induction, for all rounds of the block. \blacksquare

F Proofs for Part III (For Online Appendix)

PROOF OF LEMMA 7

By assumption, $\beta_N \leq \eta^d$ so β_N becomes arbitrarily small as η gets small. We recall equation (49) from Lemma 11:

$$w_{n+1} \leq w_n \left(1 + \frac{\alpha \beta_n}{1 - \beta_n}\right), \quad (63)$$

where $\alpha > 0$. From Bayesian updating, we have $\beta_{n+1} = \frac{\mu_n \beta_n}{\mu_n \beta_n + (1 - \beta_n)}$. Since β_N is arbitrarily small, the denominator is arbitrarily close to 1 for $n = N$. More generally we have

$$\beta_{n+1} \leq \mu_n \beta_n (1 + \epsilon) \quad (64)$$

⁹⁰The inequality holds because each term in the sum is nonnegative, and $w_n a \beta_n$ is bigger than the sum.

where ϵ is a small positive constant, as long as β_n remains small. At N , we have $\beta_N \leq \eta^d$ and $w_N = e_N - u_N \leq \bar{e}_N - u_N \leq \frac{\eta D}{2a}$, which implies from (63) that $w_{N+1} \leq \frac{\eta D}{2a}(1 + \alpha\beta_n/(1 - \beta_n))$. From (31), this implies that $\mu_N \leq \frac{1}{2} + O(\eta) \leq \frac{3}{5}$.

Consider the first round $M > N$ for which $\mu_M \geq 3/4$. The probability β_n is decreasing⁹¹ until at least round M . Proceeding by induction, from round N to round M , the previous inequalities imply that

$$w_{N+m} \leq w_N \prod_{i=1}^m (1 + \alpha(1 + \epsilon)\beta_{N+i}) \quad (65)$$

and

$$\beta_{N+i} \leq \beta_N \prod_{j=0}^{i-1} (\mu_{N+j}(1 + \epsilon)), \quad (66)$$

and hence also that (64) is valid for all rounds $n \in \{N, \dots, M\}$. From (66), we have

$$\beta_{N+i} \leq \left(\frac{3(1 + \epsilon)}{4} \right)^i \beta_N$$

Therefore, (65) implies that

$$w_M \leq w_N \prod_{i=1}^{M-N} \left(1 + \alpha(1 + \epsilon)\eta^d \left(\frac{3(1 + \epsilon)}{4} \right)^i \right)$$

The product

$$\prod_{i=1}^{\infty} \left(1 + \alpha(1 + \epsilon)\eta^d \left(\frac{3(1 + \epsilon)}{4} \right)^i \right) \quad (67)$$

is finite for η small enough, and converges to 1 as η goes to zero.⁹² Therefore, for η small, w_M is bounded above by $\frac{5}{4}w_N \leq \frac{5\eta D}{8a}$. From (31), this implies that μ_M is bounded above by $5/8 + O(\eta) < 3/4$, so M cannot be finite. This shows that for η below some threshold $\hat{\eta}$, μ_n is bounded above by $3/4$ for all $n \geq N$ and, from (64), that β_n is decreasing. Since w_n is bounded above by $\frac{3}{2}w_N$ and $w_N \leq \frac{\eta D}{2a}$, the last claim follows easily. \blacksquare

LEMMA 17 *There exists a positive constant \bar{A} such that*

$$y_n^2 \leq \frac{\bar{A}\beta_{n+1}}{1 - \beta_0} \quad (68)$$

Proof. Equation (33) implies that $Q(R_{n+1}) - Q(E_L(R_n)) \leq \frac{\beta_n w_n a}{\eta(1 - \beta_0)}$, since $\beta_n \leq \beta_0$. This, along with (34), yields⁹³

$$Q(R_{n+1}) - Q(E_L(R_n)) \leq \frac{D\beta_{n+1}}{1 - \beta_0}.$$

⁹¹This comes from the Bayesian updating equation $\beta_{n+1} = \frac{\mu_n \beta_n}{\mu_n \beta_n + (1 - \beta_n)}$, which is nondecreasing in μ_n . Taking $\mu_n = 1$ shows that $\beta_{n+1} \leq \beta_n$ as long as $\mu_n \leq 1$.

⁹²Indeed, taking the logarithm of that product, we obtain a sequence that is approximately geometric with geometric factor $3/4$ and, hence converges, uniformly in η . Moreover, each term of the sequence is of order η^d , which converges to 0 as η goes to zero. This implies that all partial sums converge to zero and, by uniform convergence, that the sequence converges to zero as well. By continuity of the exponential function, the product itself thus converges to 1 as η goes to zero.

⁹³We are using $\beta_{n+1} \geq \mu_n \beta_n$, which comes from the first inequality of (38).

Combining this inequality with Lemma 13, we get

$$y_n^2 \leq k_2 \frac{D\beta_{n+1}}{1-\beta_0} + k_3(\max\{(\beta_n w_n/(1-\beta_n))^2, (\beta_{n+1} w_{n+1})^2\} + \beta_{n+1} w_{n+1}).$$

Since $w_{n+1} \leq \frac{\eta D}{2a} \ll 1$, the last term is negligible compared to β_{n+1} . Taking \bar{A} slightly greater than $k_2 D$ proves the lemma. \blacksquare

PROOF OF LEMMA 8

Taking the square root of (42) and multiplying the result by $\frac{\sqrt{\beta_{\hat{N}+1}}}{\beta_{\hat{N}}}$, we get

$$\frac{\beta_{\hat{N}+1}}{\beta_{\hat{N}}} \geq 2c\sqrt{\beta_{\hat{N}+1}}.$$

Combining this with (40) yields

$$\frac{\beta_{\hat{N}+2}}{\beta_{\hat{N}+1}} \geq c\sqrt{\beta_{\hat{N}+1}}.$$

Taking the square root of this expression and dividing both sides by $\sqrt{\beta_{\hat{N}+1}}$, we get

$$\frac{\sqrt{\beta_{\hat{N}+2}}}{\beta_{\hat{N}+1}} \geq \frac{\sqrt{c}}{\beta_{\hat{N}+1}^{1/4}} \quad (69)$$

Combining this with (43) (and using that $\beta_{\hat{N}+1} \leq \beta_{\hat{N}}$) shows that (42) holds at round $\hat{N} + 1$. Since β_n is non-increasing in n for $n \geq \hat{N}$ and hence satisfies (43) for all $n \geq \hat{N}$, we can apply the previous argument by induction to conclude that (42) and (69) hold for all $n \geq \hat{N}$. Multiplying (40) by $\frac{\beta_n}{\beta_{n+1}}$, we obtain

$$\frac{q_{n+1}}{q_n} \geq 1 - \frac{c\beta_n}{\sqrt{\beta_{n+1}}}.$$

From (69) applied to round n (instead of $\hat{N} + 1$), the last term is bounded above by $c(\beta_n^{1/4}/\sqrt{c})$, which converges to zero as n goes to infinity. \blacksquare

PROOF OF LEMMA 9

Suppose by contradiction that $\{q_n\}$ converges to zero. This along with the second assumption of the lemma implies the existence, for any fixed $\varepsilon > 0$, of an integer \bar{N} such that i) $\frac{q_{n+1}}{q_n} \geq 1 - \varepsilon$ and ii) $q_n \leq q_{\bar{N}} \leq \varepsilon$ for all $n \geq \bar{N}$.⁹⁴ Convergence of $\{q_n\}$ to zero also implies that $\max_N \Pi_0^N q_k$ is bounded above by some constant $\bar{\Pi}$. Letting $\tilde{\varepsilon} = \sqrt{q_{\bar{N}}}$, we have $\Pi_{\bar{N}+1}^{\bar{N}+k} q_k \leq \tilde{\varepsilon}^{2k}$ for all integers $k \geq 1$. Therefore, for any integer $K \geq 1$, we have

$$q_{\bar{N}+K} = q_{\bar{N}+K} - q_\infty = \sum_{n \geq \bar{N}+K} (q_n - q_{n+1}) \leq \tilde{c} \tilde{\varepsilon}^K \sum_{k=0} \tilde{\varepsilon}^k,$$

where $\tilde{c} = c'\sqrt{\bar{\Pi}}$ and the last inequality comes from the first hypothesis of the lemma. Taking $K = 3$ and using that $\sum_{k \geq 0} \tilde{\varepsilon}^k = 1/(1 - \tilde{\varepsilon})$, this yields

$$q_{\bar{N}+3} \leq \frac{c'}{1 - \tilde{\varepsilon}} q_{\bar{N}}^{3/2} \leq 2c' q_{\bar{N}}^{3/2}. \quad (70)$$

⁹⁴Indeed, there exist N_1 such that i) holds for all $n \geq N_1$ and N_2 such that $q_n \leq \varepsilon$ for all $n \geq N_2$. Letting $N = \max\{N_1, N_2\}$, any $\bar{N} \in \arg \max_{n \geq N} \{q_n\}$ satisfies conditions i) and ii).

Applying inequality i) above to $n = \bar{N}$, $\bar{N} + 1$, and $\bar{N} + 2$, yields

$$q_{\bar{N}+3} \geq q_{\bar{N}}(1 - \varepsilon)^3. \quad (71)$$

Combining (70) and (71), we get $(1 - \varepsilon)^3 \leq 2c'q_{\bar{N}}^{1/2} \leq 2c'\varepsilon^{1/2}$, which is impossible if we choose ε small enough. This yields the desired contradiction. \blacksquare

G Proof of Theorem 2, Statement B (For Online Appendix)

Fix an initial belief $\beta_0 \in (0, 1)$ and suppose without loss that $R_0 \in \mathcal{H}$. We start by showing that the probability $p_{\mathcal{H}}$ that H ends up with a contract in \mathcal{H} converges to zero as η goes to zero. Let $\hat{Q}(u, p)$ denote the minimal expected cost of providing an expected utility u to H with a contract distribution that puts probability at least p on contracts lying in \mathcal{H} . We have $\hat{Q}(u_H(E_H(R_0)), 0) = Q(E_H(R_0))$, and $\hat{Q}(u_H(E_H(R_0)), p)$ is strictly increasing for p in a neighborhood of zero because contracts in \mathcal{H} are inefficient for H .⁹⁵ Statement A of Theorem 2 guarantees that H must get a utility arbitrarily close to $u_H(E_H(R_0))$ and that the cost to P conditional on facing H must be arbitrarily close to $Q(E_H(R_0))$ as η goes to zero.⁹⁶ For any $\varepsilon > 0$, this implies that there exists a threshold $\tilde{\eta}(\varepsilon)$ such that $p_{\mathcal{H}} < \varepsilon$ for all PBEs corresponding to any $\eta < \tilde{\eta}(\varepsilon)$.

For the remainder of the proof, we fix some small⁹⁷ $\varepsilon > 0$ and focus on η 's below the threshold $\tilde{\eta} = \tilde{\eta}(\varepsilon^4)$, so that $p_{\mathcal{H}} \leq \varepsilon^4$.

From Statement A, there exists a threshold $\hat{\eta}$ such that θ 's expected utility at the beginning of the game is bounded below by $v_{\theta}(\varepsilon) = u_{\theta}(E_{\theta}(R_0)) - \varepsilon^4$ for all η 's below that threshold.⁹⁸ Moreover, the cheapest contract $E_{\theta}(v_{\theta}(\varepsilon))$ that provides this utility costs $Q(E_{\theta}(R_0)) - O(\varepsilon^4)$ and lies within ε^4 of $E_{\theta}(R_0)$. We also recall from Lemma 2 that P 's expected cost is bounded above by $\beta_0 Q(E_H(R_0)) + (1 - \beta_0)Q(E_L(R_0))$ for any η and PBE. Fix a PBE associated with some $\eta \leq \min\{\hat{\eta}, \tilde{\eta}\}$ and let Q_{θ} denote P 's expected cost conditional on facing θ and u_{θ} denote θ 's expected utility in that PBE. The previous observations imply that $|Q_{\theta} - Q(E_{\theta}(R_0))| = O(\varepsilon^4)$ (see Footnote 96).

Let E_{ε} denote the L -efficient contract that gives H a utility of $u_H(v_H(\varepsilon))$. That contract lies within $O(\varepsilon^4)$ of $E_L(R_0)$. Part iv) of Proposition 1 implies that H never accepts a contract R' such that $u_H(R') < u_H(E_{\varepsilon})$. Moreover, L as well would reject such a contract since it would reveal his type and lead to a lower utility, from Part i) of Proposition 1. Thus, such a contract does not arise in equilibrium. Let $\mathcal{C}(\varepsilon) = \{R' \in \mathcal{C} : u_H(R') \geq u_H(E_{\varepsilon})\}$ denote the set of contracts which may arise in equilibrium. Also let $B(\theta)$ denote the ε -ball of \mathbb{R}^2 centered at $E_{\theta}(R_0)$.

The set of contracts which θ may end up with can be split between $B(\theta)$ and its complement $C(\theta) = \mathcal{C}(\varepsilon) \setminus B(\theta)$. Let p_{θ} denote the probability that θ ends up with a contract in $C(\theta)$, and let u_{θ}^B and u_{θ}^C denote the expected probabilities of θ conditional on ending up with a contract in $B(\theta)$ and $C(\theta)$, respectively. Our objective is to establish that $p_{\theta} \leq \varepsilon$ for η small enough.

⁹⁵In fact, Lemma 12 (which can be reproduced for \mathcal{E}_H instead of \mathcal{E}_L) already implies this for contracts that must lie outside of any fixed ball centered at $E_H(R_0)$.

⁹⁶Indeed, each type gets a utility arbitrarily close to $u_{\theta}(E_{\theta}(R_0))$ but P 's expected cost is bounded above by $\beta_0 Q(E_H(R_0)) + (1 - \beta_0)Q(E_L(R_0))$, from Lemma 2. Since $E_{\theta}(R_0)$'s are efficient, the claim follows.

⁹⁷It suffices to show the claim for all ε small enough, as it immediately implies that claim for higher values of ε .

⁹⁸We have proved the result for $\theta = H$, and the result is also trivially true for L , without the ε^4 , since $u_L(E_L(R_0)) = u_L(R_0)$ is a lower bound on L 's utility.

From above, we already know that $p_{\mathcal{H}} = O(\varepsilon^4)$. Consider the probability p that H ends up with a contract in $D(H) = C(H) \setminus \mathcal{H}$. We show that p is of order ε^2 . Because these contracts lie in $\mathcal{C}(\varepsilon) \setminus \mathcal{H}$, they must provide H with utility at least $u_H(R_0) - \varepsilon^4$. However, because they are outside of $B(H)$, they must cost at least $Q(E_H(R_0)) + q_H \varepsilon^2 + O(\varepsilon^4)$ from Lemma 12 (again, applying it to \mathcal{E}_H instead of \mathcal{E}_L) where $q_H > 0$. Therefore, the expected cost \tilde{Q}_H conditional on contracts being in $D(H)$ is bounded below by $Q_H + q_H \varepsilon^2 + O(\varepsilon^4)$. Moreover, the expected cost \hat{Q}_H conditional on the contracts being in $B(H)$ is bounded below by $Q_H - O(\varepsilon^4)$. Let $Q_{\mathcal{H}}$ denote P's expected cost conditional on the joint event that $\theta = H$ and that H ends up with a contract in \mathcal{H} . From $Q_H = p\tilde{Q}_H + p_{\mathcal{H}}Q_{\mathcal{H}} + (1 - p - p_{\mathcal{H}})\hat{Q}_H$ we get

$$p = \frac{Q_H - \hat{Q}_H + p_{\mathcal{H}}(Q_{\mathcal{H}} - \hat{Q}_H)}{\tilde{Q}_H - \hat{Q}_H}$$

Since the numerator is bounded above by $O(\varepsilon^4)$ while the denominator is bounded below by a factor ε^2 , we conclude that $p = O(\varepsilon^2)$ and, hence, that $p_H = p + p_{\mathcal{H}} \leq k_H \varepsilon^2$ for some $k_H > 0$.

There remains to show the result for p_L . We repeat the argument of the previous paragraph. Any contract in $\mathcal{C}(\varepsilon)$ provides L with utility at least $u_L(E_L(R_0)) - O(\varepsilon^4)$. However, because contracts in $C(L)$ lie outside of $B(L)$, they must cost at least $Q(E_L(R_0)) + q_L \varepsilon^2 + O(\varepsilon^4)$. Therefore, the expected cost \tilde{Q}_L conditional on contracts being in $C(L)$ is bounded below by $Q_L + q_L \varepsilon^2 + O(\varepsilon^4)$ where $q_L > 0$. Moreover, the expected cost \hat{Q}_L conditional on the contracts being in $B(L)$ is bounded below by $Q_L - O(\varepsilon^4)$. From $Q_L = p_L\tilde{Q}_L + (1 - p_L)\hat{Q}_L$ we get

$$p_L = \frac{Q_L - \hat{Q}_L}{\tilde{Q}_L - \hat{Q}_L}$$

Since the numerator is bounded above by $O(\varepsilon^4)$ while the denominator is bounded below by a factor ε^2 , we conclude that $p_L \leq k_L \varepsilon^2$ for some $k_L > 0$.

For ε small enough, $\max\{k_L \varepsilon^2, k_H \varepsilon^2\} \leq \varepsilon$. The threshold $\min\{\tilde{\eta}, \hat{\eta}\}$ delivers the conclusions of Statement B.

H Notation

- $u_{\theta}(n)$: type θ 's continuation utility at the beginning of round n .
- $E_{\theta}(R)$: If R is the H -Rent configuration, $E_L(R)$ is the L -efficient contract that gives L the same utility as R and $E_H(R)$ is the H -efficient contract that gives H the same utility as $E_L(R)$ (see the definition preceding Theorem 2).
- $w_n = u_H(E_H(R_n)) - u_H(n)$.
- $\bar{w}_n = \max\{u_H(E_H(m)) : m \leq n\} - u_H(n)$.
- $y_n = u_H(E_H(R_n)) - u_H(R_{n+1})$.
- $\bar{C}_{\theta} = \lim_{n \rightarrow +\infty} u_{\theta}(n)$.
- $a_{\eta} = o(\eta)$ means that $\lim_{\eta \rightarrow 0} \frac{|a_{\eta}|}{\eta} = 0$.
- $a_{\eta} = O(\eta)$ means that there exists $A > 0$ such that $\limsup_{\eta \rightarrow 0} \frac{|a_{\eta}|}{\eta} < A$.