

Contract Negotiation and the Coase Conjecture*

Bruno Strulovici
Northwestern University

December 20, 2013

Abstract

This paper analyzes an explicit protocol of contract negotiation between a principal who has all the bargaining power and an agent with a privately known type, and provides a foundation for renegotiation-proof contracts in such environments. The model extends the framework of the Coase conjecture to situations in which the seller and the buyer must determine the quantity or the quality of the good being sold. All equilibria converge to the same type-specific contracts as renegotiation frictions become negligible. Those contracts are efficient and the principal extracts a strictly positive share of the gain from negotiation.

1 Introduction

In the standard model of the durable-good monopolist, any sale is efficient and definitive: buyer and seller cannot both benefit from modifying the price of the sale. In richer contractual environments, however, a signed contract may be inefficient. For example, the parties may benefit from increasing the quantity of the good initially sold, or by agreeing on a different quality of that good.

This issue is particularly important when the buyer holds private information, because his willingness to sign some contract is informative of his type, and may thus reveal some inefficiency of the

*Acknowledgements: This paper arose from my earlier work on renegotiation-proof contracts with private information (Strulovici, 2011 and 2013). I benefited from conversations with Eddie Dekel, Mehmet Ekmecki, Drew Fudenberg, Jim Malcomson, Paul Milgrom, John Moore, Wojciech Olszewski, Alessandro Pavan, Marcin Peski, Ariel Rubinstein, Larry Samuelson, Joel Watson, and Asher Wolinsky, Alex Wolitzky's discussion at the 2013 Cowles Conference in Economic Theory, as well as from the excellent research assistance provided by Mikhail Safronov. Early versions of this project were presented at Princeton University, Caltech, UCSD, Oxford University, the Ecole 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 (2013), the Asian Meeting of the Econometric Society (2013), Columbia University, Western University, and the University of Toronto. 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.

contract just signed. This, in turn, prompts the seller to propose a new contract, and may distort the buyer’s ex ante incentives to accept any given contract.

Contract renegotiation with a privately informed agent has traditionally been studied from two different angles. The first approach is axiomatic, and focuses on “renegotiation-proof” contracts.¹ It essentially *assumes* that renegotiation leads to an efficient contract, even when one party holds private information. The second approach focuses on simple renegotiation protocols, in which the principal gets a single shot at renegotiating the contract, by making a take-it-or-leave-it offer. This approach typically results in *inefficient* contracts.²

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.³ This paper studies, instead, a more flexible negotiation protocol in which the principal can propose a new contract whenever he desires to do so, and in particular after learning any new information regarding the agent’s type. Put differently, the principal cannot commit *not* to renegotiate a contract.

While such flexibility seems necessary to guarantee ex post efficiency, proving that it is sufficient raises complex issues. To appreciate the difficulty, consider again the standard durable-good monopolist. Efficiency, in that context, means that the good is sold without delay, and was established by Gul, Sonnenschein, and Wilson (1986) as the discount rate, or breakdown probability, goes to zero.⁴ The proof is sophisticated even in this simple contractual environment, where each contract

¹See Dewatripont (1989), Maskin and Tirole (1992), Battaglini (2007), Maestri (2012), and Strulovici (2011, 2013). A similar approach has been used to study renegotiation in repeated games with complete information by Bernheim and Ray (1989) and Farrell and Maskin (1989).

²See 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.

³In the model studied here, imposing any *finite* number k of negotiation opportunities results in the full commitment outcome, as the friction parameter η goes to zero: the principal simply passes the first $k - 1$ opportunities to negotiate the contract, and then proposes the full commitment allocation.

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

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 (i.e., by inefficiently delaying the sale) 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, changes in quantities or qualities may be added), 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.

This paper analyzes a flexible negotiation protocol with the following properties. At each round, the principal can propose a menu of contracts (round zero starts with a default contract, which may correspond to the absence of a prior relationship, some status quo, or some unmodeled previous play). The agent then chooses a contract from that menu, or holds on to the last accepted contract. At the end of each round, negotiations exogenously break down with a fixed probability η , in which case the last accepted contract is implemented. The breakdown probability captures negotiation frictions: when it is equal to 1, the protocol reduces to full commitment, and the principal typically distorts the allocation of some type of the agent, creating some ex post inefficiency. The model focuses on a binary type structure, which satisfies a 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 same is true for the standard Coase conjecture in the gap and no gap cases, as long as there is no atom at the bottom of the lower end of the distribution in the no gap case).

The main finding of the paper is that, as η goes to zero, all PBE outcomes of the negotiation game converge to an immediate, separating agreement, in which each type of the agent gets an ex-post efficient and, hence, renegotiation-proof contract (i.e., P could not benefit from renegotiating the outcome of the game even if he was given an extra opportunity to do so after negotiations have broken down). Another interpretation is that renegotiation provides a dynamic implementation, without commitment, of efficient allocations. All PBEs converge to the same type-specific contracts, which are straightforward to characterize or determine graphically. Unlike the full commitment case, these contracts are independent of the initial belief. Moreover, the principal always extracts, in equilibrium, *some* surplus from negotiation. The closer to efficiency the initial contract is for one type of the agent, and the more surplus the principal can extract from the other type. To illustrate this point, consider a situation in which the initial contract is almost efficient for L , say. In that case, there is little room for renegotiation with L . This implies, intuitively, that H cannot gain

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.

much from pretending to be L , and receives almost no rent. More precisely, H 's rent converges to zero as the initial contract becomes efficient for L .

This rent-extraction pattern stands in sharp contrast with the surplus allocation obtained in the standard Coase conjecture with binary types. In that setting, the seller receives a profit $v_L - c$ from the sale, where v_L is the valuation of the low-type buyer and c is the seller's marginal cost; the high-valuation buyer gets a rent $v_H - v_L$ from the sale; the initial contract corresponds to the no-trade allocation (i.e., the seller holds the good). As one reduces v_L , keeping other parameters unchanged, the surplus from trade conditional on facing L , $v_L - c$, gets smaller, and thus the initial no-trade contract becomes more efficient for L . However, this does not improve the seller's ability to extract surplus from H . In fact, the seller's expected profit, $v_L - c$, diminishes with v_L , converges to zero and is discontinuous as v_L reaches c , since in that case the seller can credibly charge H its valuation v_H . This discontinuity of the standard Coase conjecture does not arise here. This suggests that the seller's inability to extract any rent in the standard Coase conjecture depends crucially on the very specific contractual environment (binary sale, price posting) in which it has been analyzed.

The fact that almost efficient contracts are proposed immediately implies that, *in equilibrium*, renegotiation plays a relatively minor role, even though the *possibility* of renegotiation has a decisive impact on the the contracts that are initially proposed. 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 the very first contracts that are proposed.

At a broader level, this paper analyzes a dynamic screening problem and bears some resemblance with the literature on reputation, in which some players are trying to determine whether other players have a "commitment" type.⁶ Compared to this literature, the present analysis 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 large, because it includes the last accepted contract, in addition to the principal's belief, and iii) all types of the agent are strategic. The richer state space, in particular, requires specific tools, and the use of a number of inequalities which combine the nonlinear geometry of the problem (as captured by the agent's and the principal's utility over the contract space) with the incentives of the players.

Finally, a separate contribution of the paper is to establish the existence of a PBE for a negotiation game with a (relatively) rich contract space. In the present setting, backward induction techniques

⁶See Fudenberg and Levine (1989), Schmidt (1993), Abreu and Gul (2000), Cripps et al. (2005), and Atakan and Ekmekci (2012).

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.

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 .

The agent 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 $Q : \mathcal{C} \rightarrow \mathbb{R}$. It is assumed throughout that i) u_L , u_H and Q are strictly increasing and twice continuously differentiable, ii) for each θ , u_θ is concave and its iso-level curves are not linear on any convex subset of their domain (but they can be arbitrarily close to being linear),⁷ and iii) Q is convex.

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

The functions u_L and u_H satisfy a standard single-crossing property: 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 the efficiency curves. 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. \mathcal{H} will denote the set of contracts in the *H-Rent* configuration. To rule out pathological cases, it is assumed that the efficiency curves $\mathcal{E}_L, \mathcal{E}_H$ are upward sloping. To guarantee the applicability of differential techniques, it is further assumed that for any contract C in the *H-Rent* or *L-Rent* configuration, the θ -efficient contract that gives θ the same utility as C lies on \mathcal{E}_θ .⁸

This setting is represented on Figure 1 in the context of a trade application (other applications

⁷This assumption, satisfied in the applications discussed below, holds if iso-utility curves have a nonzero curvature (however, the curvature can be arbitrarily close to zero). It is used in Lemma 12 to guarantee that at any tangency point of an iso-utility and an iso-cost curve, the curves do not stay “stuck” together as one moves away from the tangency point. The curvature at $x_0 \in \mathbb{R}$ of a real-valued function $f : x \mapsto f(x)$ is equal to $|f''(x_0)|/(1+(f'(x_0))^2)^{3/2}$.

⁸If the utility and cost functions are defined on some upper orthant $\mathcal{O} = [\ell_1, +\infty) \times [\ell_2, +\infty)$ containing \mathcal{C} , one can always extend \mathcal{C} to a larger subset \mathcal{C}' of \mathcal{O} so as to satisfy this assumption. This is for example achieved by taking \mathcal{C}' to be the smallest rectangle including \mathcal{C} and such that \mathcal{E}_L (resp. \mathcal{E}_H) hits the boundary of \mathcal{C}' on its right (resp. upper) edge.

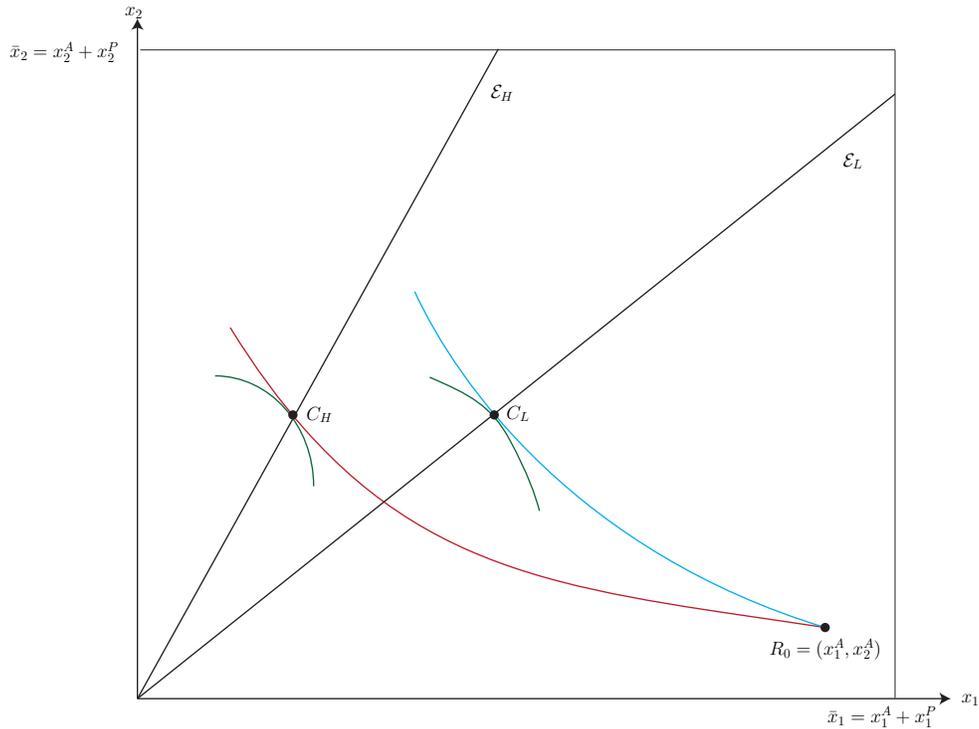


Figure 1: Setting (trade interpretation)

are described later in this section). \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: the agent’s type is privately known, initially, and $\beta_0 = Pr(\theta = H)$ characterizes P’s prior about θ . The game starts with an initial contract $R_0 \in \mathcal{C}$ representing some status quo, the absence of a prior relationship, or the result of some unmodeled earlier play. There are countably many potential rounds, indexed by $n \in \mathbb{N}$. At each round n , P can propose a menu M_n of contracts in \mathcal{C} . The number of contracts in M_n is bounded by some constant $G \geq 2$ that is arbitrary but fixed throughout the game.⁹ The agent chooses a contract in M_n or holds on to

⁹There 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. One could for example imagine that several almost identical contracts are used as a communication device between the principal and the agent. Theorem 2 concerns

the last accepted contract, R_n . Any mixed strategy over the choice set $M_n \cup \{R_n\}$ is allowed. The contract R_{n+1} that is selected by the agent becomes the new reference. At the end of each round, renegotiation breaks down with probability $\eta \in (0, 1]$ and the last accepted contract, R_{n+1} , is implemented. Otherwise, negotiation moves on to the next round.

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

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

while P's expected cost is

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

The parameter η represents the *negotiation friction*. The objective of this paper is to characterize the PBEs of the game as the friction η goes to zero.

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

For any contract $R \in \mathcal{C}$, let $(E_H(R), E_L(R))$ 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: if 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. 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 $[0, \bar{x}_1] \times [0, \bar{x}_2]$.

THEOREM 2 *Consider any initial contract R_0 and belief β_0 , and fix any $\varepsilon > 0$. There exists $\bar{\eta}(\varepsilon) > 0$ such that the following statements hold for any $\eta \leq \bar{\eta}(\varepsilon)$ 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 distance¹⁰ ε of $E_\theta(R_0)$ when renegotiation breaks down is greater than $1 - \varepsilon$.*

all equilibria, and the continuation of any equilibrium, at any round, is typically incentive inefficient. Even if an equilibrium is incentive efficient seen from round 0, there is no guarantee that its continuation equilibria are incentive efficient. This (along with several other setting differences) prevents the application of Bester and Strausz's modified revelation principle for incentive efficient contracts without commitment.

¹⁰The statement holds for any norm on \mathbb{R}^2 .

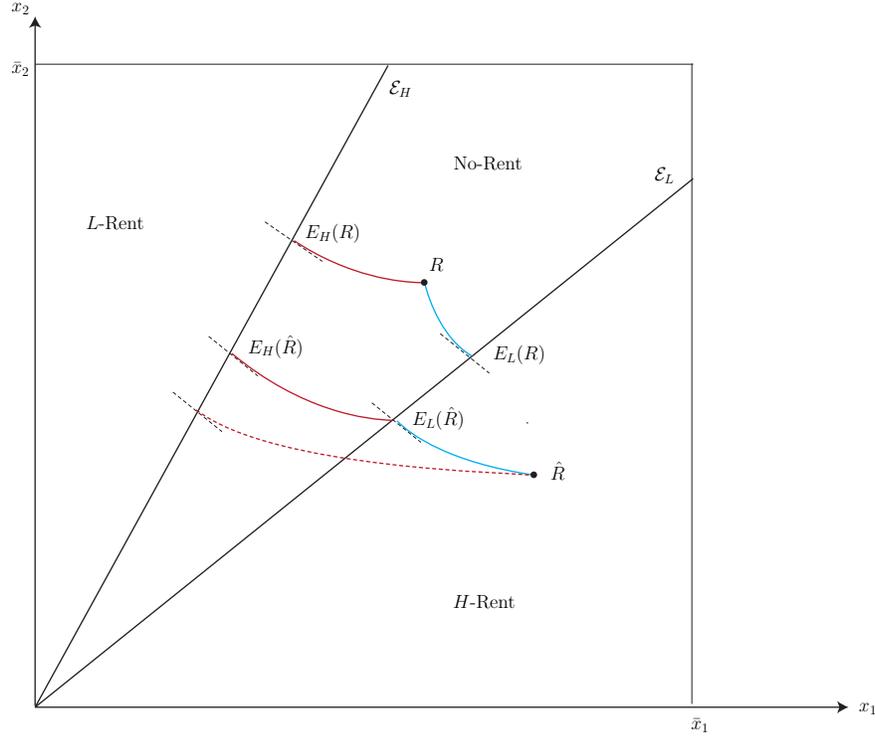


Figure 2: Renegotiation outcomes

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. The statement is a relatively straightforward consequence of Statement A, to which the quasi-totality of the proof is devoted.

Theorem 2 implies that P always gets some of the 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. Even when R_0 is in the H -Rent configuration, P gets all the surplus from negotiation that he would obtain if he were facing L , but also extracts some additional surplus in case he is facing H . The surplus extracted from H may increase as L 's utility change so as to make the initial contract closer to efficiency. At the limit, if R_0 is L -efficient, P extracts all the surplus from renegotiating R_0 with H .

Applications

1. **Durable Good Monopolist.** The agent, 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 cash holding, and u is his

concave utility function.¹¹ 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_2 captures how much wealth "P leaves to A".¹²

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 utility $u_\theta(C) = \theta\psi(-x_2) + x_1$, where the cost of effort, ψ , is increasing in its argument, and θ is a worker-specific cost of effort. 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 $u(x_1, x_2) = v(x_1) + \theta v(x_2)$ or $u(x_1, x_2) = v(x_1) + E[w(x_2, \tilde{\rho})|\theta]$ where $\tilde{\rho}$ is a taste shock whose distribution FOSD increases in θ , and w is supermodular, so that $E[\partial w/\partial x(x_2, \tilde{\rho})|\theta]$ is increasing in θ .¹³ 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 has convex preferences, and Q is the negative of a utility function representing those preferences. R_0 denotes the agent's initial holdings of the goods.

OVERVIEW OF THE PROOFS

The remaining of this section provides a short introduction and discussion of the proofs in the paper.

The proof of Theorem 2 begins with three cases of increasing complexity.

1) If the principal assigns probability 1 to one type of the agent, the problem reduces to a bargaining problem with perfect information in which P has all the bargaining power.¹⁴ P extracts all the surplus from renegotiation and the outcome is reached in a single round.

¹¹The iso-level curves of u_θ do not have linear pieces as long as the second derivative \bar{u} is strictly negative, as is easily checked. A similar condition applies to ψ in the labor application.

¹²P'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 .

¹³This application is explored in detailed by Strulovici (2013).

¹⁴There is an important difference, in analyzing off-equilibrium behavior, between perfect information and putting

2) If the initial contract R_0 is efficient for one type, say L , P cannot gain anything from negotiating with L , but extracts all the surplus from H . This outcome is again reached in a single round. Intuitively, P 's inability to gain anything from L gives him some commitment power with respect to H : H knows that there is nothing to be gained from mimicking L , since L 's allocation won't move. This intuition is simplistic, however: couldn't *both* types gain from P 's attempt to extract surplus from H ? While showing that P extracting all the gain from renegotiation with H constitutes *an* equilibrium, showing that this is the *only* equilibrium is less straightforward.

3) When R_0 is in the No-Rent configuration, the situation is similar to Case 2), because the efficient contract that P wants to propose to each type is unattractive to the other type. The main difficulty is to show that P extracting all surplus from renegotiation is the unique equilibrium. The proof relies on Cases 1) and 2), and is more complex than either of these cases. The argument derives upper bounds on P 's expected cost over all PBEs, and works out its implications for the continuation utility of each type.

Those preliminary results hold for any friction level. The hardest part of the proof, by far, is to show that efficiency obtains as the breakdown probability η goes to zero, when R_0 is in the L - or H -Rent configuration. This result is clearly false for arbitrary η : if $\eta = 1$, in particular, the situation amounts to full commitment, and the optimum proposal involves distorting one of the two types. If, for example, R_0 is the H -rent configuration, P proposes an inefficient contract to L in order to reduce H 's rent; L receives an inefficient allocation does not extract any rent, while H gets a positive rent and an efficient allocation. These observations are standard in the contracting literature with commitment, and their proofs are omitted.

Coasian dynamics and its limits

The first part of the argument is similar to the proof of the Coase conjecture (Gul et al., 1986). Suppose that R_0 is in the H -Rent configuration. To avoid giving too much rent to H , P tries to exploit H 's concern about facing an early inefficient breakdown, and H 's resulting willingness to forego some of his rents. However, P is also harmed by such breakdowns. For the scheme to be attractive, H must therefore accept the low-rent contracts with sufficiently high probability. This, however, quickly leads to a very low posterior probability of facing H , which makes P want to jump to the efficient contract for L that entails leaving all the rent to H . As the breakdown probability goes to zero, this mechanism forces P to leave all the rent to H , as in the Coase conjecture.

A major complication arises, however, owing to the richer contract space and bargaining possibilities: P does not have to wait for H to accept low-rent contracts in order to reduce the breakdown probability 1 on one type of the agent. This difference does not affect the analysis, as shown in the proof of Proposition 1.

inefficiency: he can also propose new contracts to L (and H) that reduce this inefficiency. Unlike the standard analysis of the Coase conjecture, in which the seller is reduced to posting a single price at each period, here the principal can propose arbitrarily many contracts. In particular, P could have L accept gradually more efficient contracts for L , and H randomize between those contracts and low-rent H -efficient contracts. Thus, the situation is neither one of a “gap” or “no gap” configuration: here the gap may diminish gradually as the principal proposes contracts that are increasingly more efficient for L . Addressing this difficulty lies at the core of the proof, whose strategy is described at the beginning of Section 4.¹⁵

Equilibrium existence and the lack of backward induction

Because of the richness of the contract space, it is impossible to use backward induction to characterize equilibrium behavior near efficiency (unlike in the Coase conjecture with binary types, where setting the price at L 's valuation is optimal for low enough probabilities of facing H). This prevents the use of constructive proofs for the existence of an equilibrium.

Instead, equilibrium existence is shown in two steps. In the first step, the setting is simplified to a game between P and H , and L 's behavior is taken as given: in period n , P proposes two contracts, R_{n+1} and C_n ; H randomizes between the two contracts, while L accepts R_{n+1} with probability 1. This simpler auxiliary game satisfies the assumptions of the main theorem in Harris (1985), which guarantees that it has an equilibrium. The second step builds on the equilibrium of the auxiliary game to construct a PBE of the original game.

3 Results holding for all friction levels

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 $E_\theta(R_0)$.*

¹⁵To pursue the comparison with the Coase conjecture, the seller could propose *probabilities* of getting the good: the high valuation buyer would buy the good at a high price with some probability. With the remaining probability he, like the low-valuation buyer, would buy a small probability of getting the good, which would gradually increase over time.

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 .

LEMMA 1 (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 2 *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 of generality, 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 2, 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, and ii) R_n converges to an L -efficient contract, denoted \bar{C}_L .

4 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 2, 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 $u_H(E_H(0)) - \varepsilon$. (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.¹⁶

The difference $w_0 = u_H(E_H(0)) - u_H(0)$ 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 $u_H(E_H(0))$ 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

¹⁶Without loss of generality, we focus on ε small enough so that the constant $D(2\varepsilon)$ defined in the Appendix (see (41)), is strictly positive.

- a) $\beta_{\tilde{N}} \geq \eta^d$ and $w_0 \leq \hat{w}\sqrt{\eta}$, where $\hat{w} > 0$ is exogenous, or
- 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 $\bar{w}_N \geq \underline{w}\eta$ and $\beta_N \leq \eta^d$, for some exogenous constants a, D, \underline{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 \underline{w} , which rules out Case b). (Part III).

Once Statement A of Theorem 2 has been proven, showing Statement B is relatively straightforward. The implication is shown in Appendix F.

PART I: MACRO LEVEL

The strategy of the proof is to build a sequence of *blocks* (each consisting of finitely many rounds), and 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, and 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 ε . The ulterior parts (Parts II and III) of the proof establish that $\hat{\beta}_K > \eta^d$ is the only viable 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 3 *If $u_H(E_H(R_0)) > u_H(0) + \varepsilon$, 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 at the round n_0 guaranteed by Lemma 3, so that $\bar{w}_{n_0} = \bar{e}_{n_0} - u_H(n_0) \in (\varepsilon/2, \varepsilon)$. If we show that such a round n_0 cannot exist for η small enough, this will imply, by Lemma 3, that $u_H(E_H(R_0))$ could not have been greater than $u_H(0) + \varepsilon$, and hence prove Statement A of Theorem 2. 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 is a fixed threshold, greater than 1, to be determined shortly. Also 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 . n_1 defines the last round of Block 1. Because $\hat{u}_1 < \hat{e}_0$, Lemma 14, in the Appendix, guarantees that n_1 is finite with probability 1, as it implies that $u_H(R_n)$ must eventually exceed any utility level \hat{u}_1 such that $\hat{u}_1 < \max\{u_H(E_H(R_m)) : m \leq n_0\}$, along any choice sequence, as n gets large enough.

To get a control on how much β must have dropped within Block 1, let μ_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, in the Appendix, shows that there must exist a *pushdown* 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

$$\hat{\beta}_1 \leq \frac{\hat{\beta}_0 \mu_0}{\hat{\beta}_0 \mu_0 + (1 - \hat{\beta}_0)}. \quad (1)$$

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 1. 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 4 *Along any choice sequence, H 's continuation utility at round n , $u_H(n)$, 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 equation¹⁷

$$u_H(n) = \eta u_H(R_{n+1}) + (1 - \eta)u_H(n + 1). \quad (2)$$

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 $m \geq n$, $u_H(n + 1)$ 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 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 $\hat{\beta}_0(1 - \mu_0)a(\hat{e}_0 - \hat{u}_0)$ for some Lipschitz constant $a > 0$. Indeed, $\hat{\beta}_0(1 - \mu_0)$ is the joint 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 contract that give him at least his continuation utility,¹⁸ and because that continuation utility is nondecreasing, by Lemma 4, the smallest utility that P can give H when singling him out 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 Appendix (Lemma 10).

Similarly, the expected net gain made after round n_1 , but seen from round n_0 , is bounded above by $\hat{\beta}_0\mu_0a(\hat{e}_0 - \hat{u}_1)$, because $\hat{\beta}_0\mu_0$ is the joint probability of facing H and of 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 $E_H(R_{n_0})$. The Appendix (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 is without loss of generality (see Remark 3 in the Appendix). 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 4. 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.

¹⁷More 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 (2) then follows for all contracts that are optimal for H in round n .

¹⁸Indeed, 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).

This implies that the breakdown probability is bounded below by¹⁹

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

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 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 yields²⁰

$$\beta a [(\hat{e}_0 - \hat{u}_0)(1 - \mu_0) + (\hat{e}_0 - \hat{u}_1)\mu_0] \geq \beta \mu_0 D \frac{\hat{u}_1 - \hat{u}_0}{\Delta_H}. \quad (3)$$

Recall that \hat{u}_1 was defined in terms of an arbitrary threshold t . We now define t precisely by the following equation²¹

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

With this value of t , we have

$$\mu_0 \leq \frac{a}{a + D/\Delta_H} \frac{\hat{e}_0 - \hat{u}_0}{\hat{u}_1 - \hat{u}_0} = t^{-1}.$$

Combining this inequality with (1) implies that, upon observing the constructed pushdown choice sequence until round n_1 , the posterior $\hat{\beta}_1$ satisfies

$$\hat{\beta}_1 \leq \frac{\mu_0 \hat{\beta}_0}{\mu_0 \hat{\beta}_0 + (1 - \hat{\beta}_0)} \leq \hat{\beta}_0 \frac{t^{-1}}{\hat{\beta}_0 t^{-1} + (1 - \hat{\beta}_0)} = g \hat{\beta}_0.$$

where $g = \frac{t^{-1}}{\hat{\beta}_0 t^{-1} + (1 - \hat{\beta}_0)}$. (We also use the inequality $\hat{\beta}_0 \leq \beta_0$.) Because $t^{-1} < 1$, g is strictly less than 1. We have thus achieved our 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 1.²² 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 4. The

¹⁹The 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 $x \leq 0$.

²⁰For 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 23.

²¹ D is defined independently of t (and of this entire block construction), so there is no circularity in the definition.

²²The 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.

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 μ_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 $\hat{\beta}_2 \leq \frac{\hat{\beta}_1 \mu_1}{\hat{\beta}_1 \mu_1 + (1 - \hat{\beta}_1)}$. Let n_2 denote the round at which \hat{u}_2 is first exceeded. By a similar analysis to the used for the first block, we have

$$\hat{\beta}_2 \leq \frac{\mu_1 \hat{\beta}_1}{\mu_1 \hat{\beta}_1 + (1 - \hat{\beta}_1)} \leq \hat{\beta}_1 \frac{t^{-1}}{\hat{\beta}_1 t^{-1} + (1 - \hat{\beta}_1)} \leq g^2 \hat{\beta}_0.$$

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. Upon observing the pushdown choice sequence across blocks 1 to k , we get

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

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 $\bar{W} > \max\{\frac{t-1}{t}(1 + \Delta_H), \frac{\hat{W}}{t \Delta_H}\}$, where \hat{W} is an arbitrarily large constant.²³ 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 $\eta < \tilde{\eta}$.*

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

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

²³The number of rounds in each block $k \leq K$ is bounded below by $\frac{\hat{u}_k - \hat{u}_{k-1}}{\Delta_H \eta} \geq \frac{1}{t \Delta_H \eta}(\hat{e}_{k-1} - \hat{u}_{k-1}) \geq \frac{\bar{W} \eta}{t \eta \Delta_H} = \bar{W}$, which can be made arbitrarily large by choosing \hat{W} appropriately. The reason for choosing $\bar{W} > \frac{t-1}{t}(1 + \Delta_H)$ is explained at the beginning of Part II.

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

$$\hat{e}_K - \hat{u}_0 \leq \left(\frac{t}{t-1}\right)^K \bar{W}\eta. \quad (4)$$

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 $\hat{e}_0 - \hat{u}_0 \in (\varepsilon/2, \varepsilon)$.

PART II: MICRO LEVEL

This part, and the next, establish 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 that²⁴

$$\hat{w}_K = \hat{e}_K - \hat{u}_K \geq \bar{W}\eta \quad (5)$$

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 4, we obtain, for the augmented index evaluated at round n_K ,²⁵

$$\hat{w}_K = \hat{e}_K - u_H(n_K) \geq (\bar{W} - \Delta_H)\eta \geq \eta. \quad (6)$$

If $\hat{w}_K \leq \frac{D\eta}{2a}$, Proposition 5 (Part III) implies that $\hat{w}_K \leq \hat{w}\eta^{1+d}$, which contradicts (6) for η small enough, ruling out this case.

The objective of the present part is to solve the remaining case, for which $\hat{w}_K \in \left(\frac{D\eta}{2a}; \bar{W}\eta\right)$. 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.

²⁴This inequality comes 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 $\hat{e}_K \geq \hat{e}_{K-1}$.

²⁵The 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 $\hat{u}_K + \Delta_H\eta$.

PROPOSITION 4 *Let $N \geq n_K$ denote the first round for which $\bar{w}_N \leq \frac{\eta D}{2a}$. Then,*

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

The objective of this part is, therefore, to build a bridge between Parts I and III, showing that 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 expressing P's IC constraint, at each round n . 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} . Also let $E_\theta(n) = E_\theta(R_n)$. Because P can always jump to $(E_L(n), E_H(n))$ by Lemma 1, P's IC constraint implies, as explained below, that

$$\begin{aligned} w_n a \beta_n &\geq \sum_{R_{n+1} \in (M_n \cup \{R_n\}) \cap \mathcal{H}} \beta_n \mu_n^H(R_{n+1}) \eta D + (1 - \beta_n) \mu_n^L(R_{n+1}) \eta (Q(R_{n+1}) - Q(E_L(n))) & (7) \\ &= \sum_{R_{n+1} \in (M_n \cup \{R_n\}) \cap \mathcal{H}} \mu_n^L(R_{n+1}) [\beta_n \mu_n(R_{n+1}) \eta D + (1 - \beta_n) \eta (Q(R_{n+1}) - Q(E_L(n)))], & (8) \end{aligned}$$

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 11.²⁶

The left-hand side of (7) 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 11, in the Appendix, this gain is bounded above by $a(u_H(E_H(n)) - u_H(n)) = a w_n$ (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 $u_H(n)$).²⁷ 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, in case a breakdown occurs at the end of round n . This loss is bounded by D as long as $w_n \leq 2\epsilon$, 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.

Proposition 4 is based on the following lemma. Fix any positive integer \bar{N} , positive constants $\bar{\beta}$ and \bar{w} , as well as a small positive $\bar{\epsilon}$. Let $y_n = u_H(E_H(n)) - u_H(R_{n+1})$. The quantity y_n

²⁶We 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.

²⁷This is an upper bound on the gain, since C_n is the cheapest way of providing H with his continuation utility.

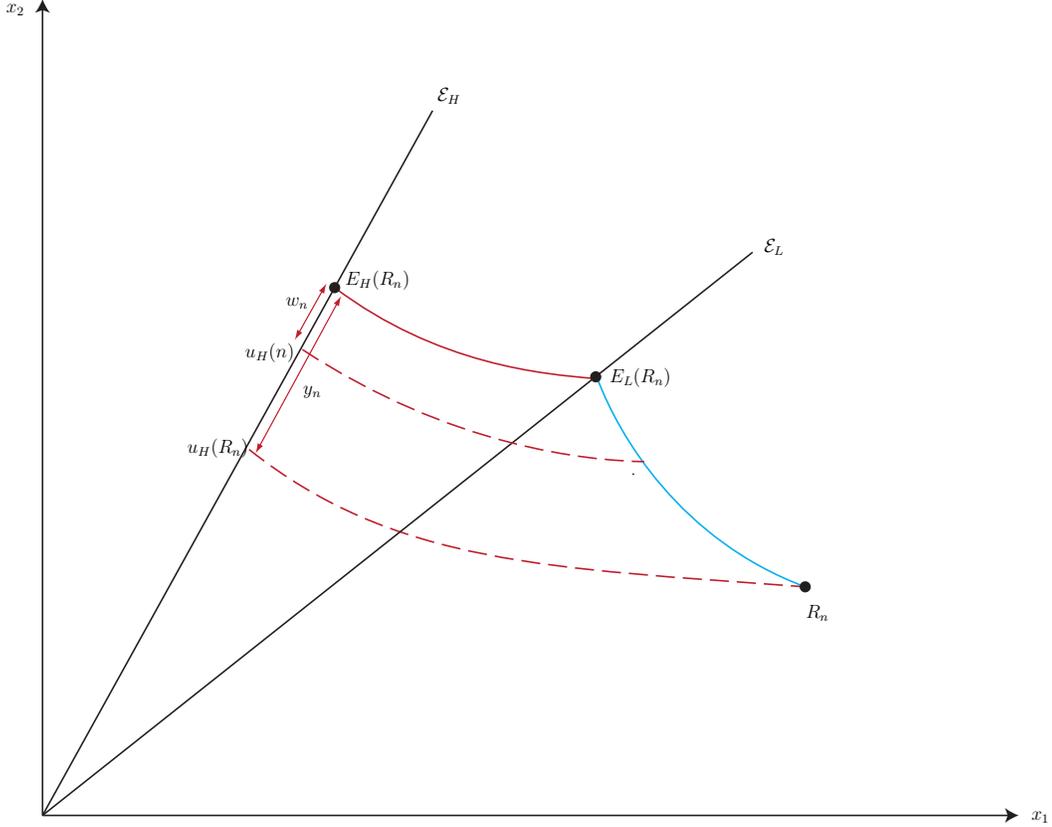


Figure 3: Concepts

represents 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 increments of w_n and makes sure that we do not overshoot the threshold $\frac{\eta D}{2a}$ by too much. Indeed, subtracting $u_H(E_H(n))$ from (2) and rearranging (and recalling that $w_n = u_H(E_H(n)) - u_H(n)$) leads, along any choice sequence, to²⁸

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

These concepts are represented on Figure 3

LEMMA 5 Consider a round \bar{n} such that $\beta_{\bar{n}} \leq \bar{\beta}\eta^d$ and $w_{\bar{n}} \leq \bar{w}\eta$ and fix some positive integer \bar{N} and a small number $\bar{\varepsilon} > 0$. There exist $\bar{W}(\bar{N})$ and $k(\bar{N})$ (independent of η and $\bar{\varepsilon}$) such that the following holds. Let \mathcal{S} denote the event that the agent chooses contracts at rounds $n \in \{\bar{n} + 1, \dots, \bar{n} + \bar{N}\}$

²⁸Equation (38) in the Appendix shows that $w_{n+1}(1 - b\beta_{n+1}) \geq w_n - \eta y_n$, which is simpler to work with.

such that $y_n = O(\eta^{d/4})$, $\beta_n \leq \beta_{\bar{n}}\bar{\varepsilon}^{-(n-\bar{n})}$, and $w_n \leq \bar{W}(\bar{N})$ for all rounds $n \in \{\bar{n}, \dots, \bar{n} + \bar{N}\}$. For η small enough, the probability of \mathcal{S} is greater than $1 - k(\bar{N})\bar{\varepsilon}$.

We will 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, must satisfy

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

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 before, 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 loss D is valid with probability $1 - k(\bar{N})\bar{\varepsilon}$ because w_n is small throughout the block with that probability, by Lemmas 5, and the lower bound on the loss then follows from Lemma 16. For the remaining probability, whatever cost is incurred by P , conditional on facing H , is bounded above by some constant. The difference between D and that constant is taken into account by the last term, where $\delta_Q > 0$ is independent of η, \bar{N} , and $\bar{\varepsilon}$.

This implies that

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

where the second inequality comes from the fact that u_H is nondecreasing across all paths which implies, taking expectations, that $E u_H(\bar{n} + \bar{N}) \geq u_H(\bar{n})$, and from the fact that $\bar{w}_{\bar{n}} \leq \bar{W}\eta$ (this inequality holds for all blocks of Part II, without loss of generality, see Remark 1 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 the 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

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

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 $\delta_Qk(\bar{N})\bar{\varepsilon} < a\bar{W}$.

Proceeding as in Part I, there must exist a pushdown choice sequence within \mathcal{S} such that the ex post probability that H has not chosen an H -efficient contract is weakly less than $\mu_{\bar{n}}^{\mathcal{S}}$. Therefore, along that sequence, we have i) y_n small, and ii) $\beta_{\bar{n}+\bar{N}} \leq \frac{\beta_{\bar{n}}}{2}$: in other words, we have built a 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.²⁹

Starting from round n_K , we build a sequence of blocks of this kind. Because \bar{w}_n converges to zero (Lemma 14, part ii)), it will eventually cross $\frac{\eta D}{2a}$. Let N denote the first round at which \bar{w}_N drops below that threshold. From (38), in the Appendix, we have

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

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 round $N - 1$, we obtain

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

Finally, we have

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

The difference in parentheses is bounded above by $\frac{\alpha\beta_{N-1}}{1-\beta_{N-1}}w_{N-1} = o(\eta)$, from (34). Since $\bar{w}_{N-1} > \frac{\eta D}{2a}$, by construction of N , we conclude that

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

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 $\hat{w} > 0$ independent of η , and ii) $\beta_N = O(\eta^d)$. Part III will show that this is impossible.

REMARK 1 *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 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*

²⁹Notice that β_n can increase up to $\beta_{\bar{n}}\varepsilon^{-\bar{N}}$ 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.

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 $\beta_N \leq \eta^d$. Proposition 5 of Part III shows that such a round cannot exist, proving the contradiction.

PART III: ASYMPTOTIC LEVEL

The goal of this section is to show the following proposition:

PROPOSITION 5 *There exists a constant $\hat{w} > 0$ such that if one reaches a round N for which $\beta_N \leq \eta^d$ and $\bar{w}_N \leq \frac{\eta D}{2a}$, then $\bar{w}_N \leq \hat{w}\eta^{1+d}$ for all η below some threshold.*

The proof proceeds in three steps: 1) Show that, starting from round N , one can build a choice sequence along which β_n is decreasing and a simplified version of P's ex ante IC constraint at each round is also satisfied ex post; 2) Show that along such a sequence, one must necessarily have $w_n \leq \eta\beta_n c$ for all $n \geq N$, for some $c > 0$ that depends on a and D but not on η ; 3) Conclude with the desired inequality for \bar{w}_N , by showing that $\bar{w}_N - w_N = O(\eta^{1+2d})$ (Proposition 7). Combining the last two observations, along with the fact that $\beta_N \leq \eta^d$, proves Proposition 5.

The idea of the proof is again to derive, for a given PBE, an equation for the dynamics of the posterior belief β_n , based on an incentive compatibility condition for P. This time, however, there are no blocks: the equation is used for each single round, and exploits the losses on *both* types. For η small enough, this equation is then showed to contradict the convergence of w_n to zero, which has to hold for any PBE, by Proposition 2.

In the Appendix (Lemma 11), it is shown that

$$u_H(E_H(n+1)) - u_H(E_H(n)) \geq -\hat{b}\beta_{n+1}w_{n+1} \quad (9)$$

for some constant $\hat{b} > 0$. 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 $E_H(n)$ and $E_H(n+1)$ are constructed based on the utility that L gets from R_n and R_{n+1} . A simple Lipschitz property, established in the Appendix, then yields (9).

We also use the following relation between w_n and w_{n+1} (this is (39), in the Appendix):

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

Combining this with inequality (9) yields

$$(1 - \eta)w_{n+1} \geq w_n - \eta y_n - \hat{b}\beta_{n+1}w_{n+1}. \quad (10)$$

In Part I, we focused on *pushdown* sequences, in order to relate P's ex ante incentive compatibility constraint with his ex post belief's about the agent. Here, similarly, we need to focus on particular choices by the agent that play a similar role.

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

$$w_n a \beta_n \geq \sum_{R_{n+1} \in (M_n \cup \{R_n\}) \cap \mathcal{H}} \mu_n^L(R_{n+1}) [\beta_n \mu_n(R_{n+1}) \eta D + (1 - \beta_n) \eta (Q(R_{n+1}) - Q(E_L(n)))], \quad (11)$$

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 obtained in Lemma 11.

In particular, the RHS of (11) 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

$$w_n a \beta_n \geq \beta_n \mu_n(R_{n+1}) \eta D + (1 - \beta_n) \eta (Q(R_{n+1}) - Q(E_L(n))). \quad (12)$$

Therefore, there exists a choice sequence that satisfies (12) for all $n \geq N$. In what follows we entirely focus on that sequence, which will be called a **regular** choice sequence.

We split the IC constraint (12) into two parts, in a way that allows us to modify P's belief β_n to achieve tractability. The first step is to note that $Q(R_{n+1}) - Q(E_L(n)) \geq Q(E_L(n+1)) - Q(E_L(n)) \geq -k\beta_{n+1}w_{n+1}$, where the second inequality comes (36) of Lemma 11 in the Appendix. To simplify notation, let $\mu_n = \mu_n(R_{n+1})$. Equation (12) 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

$$\mu_n \leq \frac{w_n a}{\eta D} + k \frac{\beta_{n+1}}{\beta_n D} w_{n+1}. \quad (13)$$

The first step, in order to exploit this equation, is to show that β_n is small and decreasing along that sequence for $n \geq N$. More precisely, we will prove the following lemma.

LEMMA 6 *There exists $\hat{\eta} > 0$ and $\hat{w} > 0$ such that for $\eta < \hat{\eta}$ and $n \geq N$, i) β_n is decreasing in n , ii) μ_n is bounded above by $3/4$, and iii) w_n is bounded above by $\hat{w}\eta$.*

Lemma 6 implies that $\frac{\beta_{n+1}}{\beta_n} \leq \mu_n(1 + \epsilon) \leq \frac{3}{4}(1 + \epsilon)$. Hence, the second term in the right-hand side of equation (13) is of order w_{n+1} , while its first term is of order $\frac{w_n}{\eta}$. Since w_{n+1} is bounded above

by $w_n \left(1 + \frac{\alpha\beta_n}{1-\beta_n}\right)$, from (35) in the Appendix, the second term is negligible compared to the first one. Therefore, by slightly increasing a , whose specific value does not matter in any case for the proof, we get

$$\mu_n \leq \frac{w_n a}{\eta D}, \quad (IC_n^{LL}) \quad (14)$$

Moreover, (12) also implies that

$$\beta_n w_n a \geq (1 - \beta_n) \eta (Q(R_{n+1}) - Q(E_L(n))) \quad (IC_n^{LH}) \quad (15)$$

Relaxed Problem

We now introduce a relaxed problem, for rounds $n \geq N$, constructed in two steps. First, P's incentive constraint at each round $n \geq N$ is weakened to the coupled inequalities (14) and (15). Clearly, if the initial sequence of contracts proposed by P and choices made by the agent formed a PBE, then the problem in which P's IC constraint is replaced by these two weaker equations must also have a solution.

The second step is as follows. If, starting from round N , one decreases $\mu_N^L(R_{N+1})$ so that $\mu_N = \mu_N^L(R_{N+1}) = \frac{\mu_N^H(R_{n+1})}{\mu_N^L(R_{n+1})}$ is increased so as to satisfy (14) as an equality,³⁰ then the posteriors β_n 's at all rounds $n \geq N + 1$ are increased as a result, by Bayesian updating. Therefore, P's IC constraint (15) at future rounds is preserved (in fact, looser) along the choice sequence of interest.³¹ After this is done, one can increase $\mu_{N+1}^L(R_{N+2})$ (by decreasing $\mu_N^L(R_{N+2})$) so as to make IC_{N+1}^{LL} tight, preserving IC constraints (15) for $n \geq N + 2$, (without perturbing P's IC constraints at round N), etc. To simplify the notation, let $\mu_n = \mu_n^L(R_{n+1})$. Increasing μ_n inductively for all $n \geq N$, this shows that the regular choice sequence, along with the new mixing probabilities for L and resulting belief sequence $\{\beta_n\}_{n \geq N+1}$ is a solution to the relaxed problem with

$$\mu_n = \frac{w_n a}{\eta D} \quad (16)$$

The resulting regular sequence, with the new mixing probabilities and beliefs, is called the *relaxed* version of the initial sequence. Since only the agent's mixing probabilities were changed, the contracts proposed, the sequence w_n , and the limits \bar{C}_L, \bar{C}_H are unchanged. In particular, H 's and L 's incentives are unchanged. Since, also, w_n goes to zero, (14) implies that μ_n goes to zero and, hence, that β_n goes to zero along the relaxation. It is worth emphasizing that the contracts

³⁰The interpretation would be that the principal erroneously becomes too optimistic, conditional on observing R_{n+1} , about the posterior probability of facing H .

³¹At this point, we totally ignore the IC constraints outside of the regular sequence. Our point is only to show that a contradiction along that specific choice sequence.

have *not* changed. We are still focusing on the same regular choice sequence, and merely observing that a sequence of coupled inequalities, which had to hold before, implies a new sequence of paired equality and inequality, which are easier to work with.

As just argued, the relaxed problem, which consists of the incentive compatibility and dynamic equations arising along the relaxed regular sequence, must have a solution. The remainder of the proof shows that this is impossible, which will yield the desired contradiction. We do this by building a dynamic equation for β_n , in the relaxed problem, and show from that equation that w_n cannot converge to zero along the relaxed regular sequence, contradicting Proposition 2.

Multiplying both sides of (10) by $\frac{a}{\eta D}$, and using (16), we obtain for $n \geq N$

$$\mu_{n+1} \geq (1 - \eta)\mu_{n+1} = \mu_n - \frac{a}{D}y_n - \tilde{b}\beta_{n+1}. \quad (17)$$

for some constant $\tilde{b} > 0$ (also using that $w_n \leq \hat{w}$ for $n \geq N$, from Part iii) of Lemma 6).

The Bayesian updating equation

$$\beta_{n+1} = \frac{\beta_n \mu^H(R_{n+1})}{\beta_n \mu^H(R_{n+1}) + (1 - \beta_n) \mu^L(R_{n+1})} = \frac{\beta_n \mu_n}{\beta_n \mu_n + (1 - \beta_n)}$$

implies that³²

$$\frac{\beta_{n+1}}{\beta_n} \geq \mu_n \geq \frac{\beta_{n+1}}{\beta_n} - \mu_n \beta_n + \mu_n O(\beta_n^2) \geq \frac{\beta_{n+1}}{\beta_n} - \beta_{n+1} + o(\beta_{n+1}). \quad (18)$$

Equation (49) in the appendix 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).³³

Combining the upper bound on y_n^2 with (17) and (18), we obtain the following dynamic equation for β_n , for all $n \geq N$:

$$\frac{\beta_{n+2}}{\beta_{n+1}} \geq \frac{\beta_{n+1}}{\beta_n} - c\sqrt{\beta_{n+1}} - (1 + \tilde{b})\beta_{n+1} \quad (19)$$

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 next to last term, because $\sqrt{\beta_{n+1}} \ll c$. Therefore, by slightly increasing the value of c , whose precise value does not affect the proof, we obtain

$$\frac{\beta_{n+2}}{\beta_{n+1}} \geq \frac{\beta_{n+1}}{\beta_n} - c\sqrt{\beta_{n+1}}. \quad (20)$$

³²We 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.

³³Dividing by η , we get $y_n^2 \leq C\beta_n w_n / \eta$ for some constant C . Since w_n / η is proportional to μ_n and $\mu_n \beta_n$ is roughly equal to β_{n+1} , this gives some idea for how the equation was derived.

Letm for all n , $q_n = \frac{\beta_{n+1}}{\beta_n}$. We have $\prod_0^n q_k = \frac{\beta_{n+1}}{\beta_0}$. (20) may be rewritten as

$$q_{n+1} \geq q_n - c' \sqrt{\prod_0^n q_k} \quad (21)$$

where $c' = \sqrt{\beta_0}c$. Because q_n is proportional to w_n , it must converge to zero as n goes to infinity. The rest of this section shows, however, that this is impossible. Let $\hat{c} = \frac{4c^2D}{a}$.

PROPOSITION 6 *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, which are proved in the Appendix.

LEMMA 7 *Suppose that there exist $\hat{N} > N$ such that*

$$\beta_{\hat{N}+1} \geq 4c^2\beta_{\hat{N}}^2, \quad (22)$$

$$\beta_{\hat{N}}^{1/4} \leq \frac{1}{2\sqrt{c}}, \quad (23)$$

Then,

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

LEMMA 8 *Suppose that $\{q_n\}$ is a strictly positive sequence such that*

$$q_n - q_{n+1} \leq c' \sqrt{\prod_0^n q_k}$$

and $\liminf_n q_{n+1}/q_n \geq 1$. Then, $\{q_n\}$ does not converge to zero.

To conclude the proof of Proposition 6, suppose that there exists $\hat{N} \geq N$ such that $w_{\hat{N}} > \hat{c}\eta\beta_{\hat{N}}$. From (16), this implies that $\mu_{\hat{N}} > \frac{a}{D}\hat{c}\beta_{\hat{N}}$, and from the first inequality of (18), this implies (using the definition of \hat{c}), that equation (22) holds for \hat{N} . Moreover, from Lemma 6, $\beta_{\hat{N}}$ clearly satisfies (23), for η small enough. Therefore the hypotheses of Lemma 7 are satisfied and, hence, $\liminf_{n \rightarrow +\infty} \frac{q_{n+1}}{q_n} \geq 1$. Combining this with Lemma 8 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

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

where the last inequality comes from (34). From Proposition 6, we have $w_n \leq \hat{c}\eta\beta_n$ for all $n \geq N$. Therefore,

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

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

We have $\beta_{n+1} \leq 2\mu_n\beta_n = \frac{2aw_n}{D\eta}$, by (16). Using again the inequality $\frac{w_n}{\eta} \leq \hat{c}\beta_n$, which holds for all $n \geq N$, we have, letting $\hat{K} = \frac{2a\hat{c}}{D}$,

$$\beta_n \leq \beta_N \prod_{k=N+1}^{n-1} (\hat{K}\beta_k) \leq \beta_N (\hat{K}\beta_N)^{n-N}.$$

For $\beta_N < \frac{1}{\sqrt{2\hat{K}}}$, this implies that

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

We then obtain

$$u_H(\bar{C}_H) - u_H(E_H(N)) \leq 2\tilde{K}\eta\beta_N^2 \leq 2\tilde{K}\eta^{1+2d}, \quad (24)$$

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

To conclude, note that $\bar{w}_N - w_N = \max\{e_k : k \leq N\} - e_N = \max\{u_H(E_H(k)) : k \leq N\} - u_H(E_H(N))$. 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), (24) yields the result. \blacksquare

5 Conclusion

Whether it concerns trade, production, labor, finance, or any other economic activity, contract negotiation is central to economic analysis and has been widely studied from both pure and applied perspectives. However, much remains to be discovered 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 this common situation, private information is endogenously revealed through negotiation, even in the absence of any exogenous information arrival, and this information affects ongoing negotiations, providing a dynamic interaction between beliefs and contracts.

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 indeed essential to generate a foundation for renegotiation-proof contracts.³⁴ While the analysis has focused on a specific negotiation protocol (a feature common to many foundations of cooperative concepts, from the foundations of the Nash bargaining solution to the foundations of the Coase conjecture), it would be useful to explore, in future research, more general protocols of negotiation. A natural conjecture is that contract negotiation should lead to efficient outcomes as long as it is unconstrained in the above sense and does not entail additional frictions, such as explicit renegotiation costs.³⁵

In the present setting, where the efficiency “gap” is endogenous, varies over time, and converges to zero, it is perhaps surprising that ex post efficiency should obtain for all equilibria, including non-stationary ones, of the negotiation game, as frictions vanish. It would be natural to study extensions of these results to more than two types or contract dimensions and to independent values imposed on the model.³⁶ While those extensions do not seem straightforward, the techniques developed here should prove helpful to analyze them. One might also wish to consider shared bargaining power between the principal and the agent.³⁷ While important, this extension is not needed to

³⁴For 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.

³⁵See, e.g., Brennan and Watson (2013).

³⁶Deneckere and Liang (2006) provide an in-depth analysis of the interdependent value case, for a binary sale.

³⁷Ausubel 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.

provide a foundation for the renegotiation-proof contracts, most common in the literature,³⁸ that are *proposed by the principal*. In fact, any negotiation protocol in which the bargaining power is shared by the parties should typically *not* yield the renegotiation-proof contracts that are optimal for the principal. That point is straightforward to see in the case of symmetric information.

The analysis was focused on a single “delivery” time, at which the contract is implemented. To provide a foundation for multi-period renegotiation-proof contracts, one should consider a general model with multiple “physical” events, i.e., times at which payments or efforts are made or new exogenous information arrives (e.g., if the type of the agent is persistent but not constant). A renegotiation protocol like the one studied in this paper could be inserted between physical events, and negotiation would pertain to continuation contracts for the remaining horizon. For example, if good deliveries or monetary transfers occur at integer times $t = 1, 2, \dots$, the renegotiation rounds between dates $t - 1$ and t would occur at times $\tau_n^t = t - \frac{1}{2n}$, 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.³⁹

³⁸See, e.g., Dewatripont (1989), Maskin and Tirole (1992), and Battaglini (2007).

³⁹The time structure described here allows unrestricted renegotiation between physical events, as there are infinitely 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.

Appendices

A Proof of Theorem 1

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, and 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 by 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 $R_0 \in \mathcal{H}$, where $\mathcal{H} \subset \mathcal{C}$ is the set of contracts in the H -Rent configuration, and a belief $\beta \in (0, 1)$. For this auxiliary game, β is just a parameter of the payoff functions and devoid of its interpretation as a belief.

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

$$u_L(R_{n+1}) \geq u_L(R_n) \quad (25)$$

$$u_L(R_{n+1}) \geq u_L(C_n) \quad (26)$$

$$u_H(C_n) \geq u_H(R_n). \quad (27)$$

H then chooses a number $\mu_n \in [0, 1]$. The interpretation 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 just an action deterministically affecting payoffs.

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

$$\begin{aligned} \mathcal{Q}(\{R_{n+1}, C_n\}, \{\mu_n\}) &= \sum_{n \geq 0} Q(C_n) \beta (1 - \eta)^n (1 - \mu_n) \prod_{k=0}^{n-1} \mu_k \\ &\quad + \sum_{n \geq 0} Q(R_{n+1}) (\beta (1 - \eta)^n \eta \prod_{k=0}^n \mu_k + (1 - \beta) (1 - \eta)^n \eta), \end{aligned} \quad (28)$$

while H 's payoff is

$$\mathcal{V}(\{R_{n+1}, C_n\}, \{\mu_n\}) = \sum_{n \geq 0} u_H(C_n) (1 - \eta)^n (1 - \mu_n) \prod_{k=0}^{n-1} \mu_k + \sum_{n \geq 0} u_H(R_{n+1}) (1 - \eta)^n \eta \prod_{k=0}^n \mu_k. \quad (29)$$

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 .

LEMMA 9 *For any initial R_0 and $\beta \in (0, 1)$, there exists a perfect equilibrium of the auxiliary game*

Proof. To apply Theorem 1 of Harris (1985), we need to check Assumptions 1–5 of that theorem. The payoff function of the principal is simply the negative of his cost, \mathcal{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 (25) and (27), 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 $-\mathcal{Q}$ and \mathcal{V} are clearly continuous on their domain \mathcal{S}_f , so Assumption 5 is satisfied as well. The result follows. ■

REMARK 2 *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*

$$V_H(\beta) = \sum_{n \geq 0} (1 - \eta)^n \eta u_H(R_{n+1}). \quad (30)$$

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.⁴⁰ By construction, \bar{R}_{n+1} and \bar{C}_n satisfy (25)–(27). 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.

⁴⁰If 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 .⁴¹ There is always 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 .⁴² 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 2) of the auxiliary game starting in period 0 with belief $\tilde{\beta}_0 = \beta_n$ and at some fictitious contract \tilde{R} in the L -Rent configuration such that \hat{R}_{n+1} and \bar{C}_n satisfy equations (27) and (25), 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 (26), 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) less than $u_L(R_{n+1})$ and, hence, less than his continuation utility if he chooses R_{n+1} .⁴³ 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 above⁴⁴ by $u_L(R_n)$, which is weakly dominated by accepting R_{n+1} by (25) (guaranteeing that L 's continuation utility is bounded below by $u_L(R_{n+1})$).

Now consider the optimality of H 's strategy, still assuming that P follows the prescribed sequence of contracts. From (27), $u_H(C_n) \geq u_H(R_n)$. 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 (29), and by perfection of the auxiliary equilibrium, the strategy $\{\mu_n\}$ is a best response to the sequence of contracts.

Now consider 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

⁴¹That 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 .

⁴²By construction, $\bar{R}_{n+1} \in \mathcal{H}$, \hat{C}_n is H -efficient, and the contracts satisfy conditions (25), (26), and (27), so the contract pair is feasible for the auxiliary games.

⁴³That utility is always weakly higher than $u_L(R_{n+1})$, since L can always hold on to R_{n+1} .

⁴⁴Indeed, 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.

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 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 the probability μ_n coming from the auxiliary equilibrium in which \bar{R}_{n+1} and \hat{C}_n are proposed is optimal.⁴⁵

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 (25), (26), and (27). As shown by Lemma 2 (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 Proofs of Section 3

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 $u < \bar{u}$. 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,⁴⁶ 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 the deviation is profitable.⁴⁷ Let Q denote the cost of the θ -efficient contract, C , that gives θ utility u . Clearly, any PBE must cost exactly Q , otherwise P has a profitable deviation which is to propose the θ -efficient

⁴⁵Because $u_H(\hat{C}_n) = u_H(\bar{C}_n)$, 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.

⁴⁶See Fudenberg and Tirole (1991), part iii) of the definition.

⁴⁷The continuation play after P's deviation must be a PBE of the corresponding continuation game. Therefore, if

contract that gives θ slightly more than u and costs less than the PBE. Moreover, the only way of achieving Q is to propose 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 the expected cost under 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 $\bar{Q}_H \leq Q$. 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 .⁴⁸ 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, and thus Part ii) of the proposition.

Part iii) Suppose without loss that $Q_L \geq Q_H$, where $Q_\theta = Q(E_\theta(R_0))$ (the other 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} .⁴⁹ 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, the continuation cost is by construction less than \bar{Q} . This means that there is one type of the agent who must be getting a lower payoff if he rejects the contract \tilde{C}_θ meant for him, because P has to be spending weakly no more on his contract than under \tilde{C}_θ (up to ε , which is negligible compared to η). The contract \tilde{C}_θ maximizes that type's utility subject to that cost constraint. Since rejection also 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 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. This shows that, necessarily, $\bar{Q} \leq Q_L$.

Since L cannot get utility less than $u_L(C)$, under any PBE, and Q_L is the cheapest way of providing that 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.

⁴⁸If $\beta = 0$, P does not propose anything new, from i) and L -efficiency of R_0 , and the result trivially holds in that case too.

⁴⁹If 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} .

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 highest expected utility that H gets over all PBEs and beliefs $\beta > 0$. Since the principal 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 $\bar{u}_H = u_H(C_H)$. Suppose by contradiction that $\bar{u}_H > u_H(C_H)$, and consider a PBE that achieves \bar{u}_H (again, the proof is easily adapted if the maximum is not achieved, by considering a PBE that gets very close to providing \bar{u}_H). 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 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 $u_L(C)$. The cost gain on the high type is of order ε compared to $Q(\bar{C}_H)$ and the cost loss 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 $\bar{u}_H = u_H(C_H)$. This immediately implies that L 's maximal utility across all PBEs for $\beta \in (0, 1)$ is $u_L(C_L)$. The proof is a straightforward modification of the end of the proof of Part i).

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 C . 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,⁵⁰ 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 ii), 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 $u_H(0)$ is bounded above by the utility that H achieves with the H -efficient contract that costs Q_L . We now show that $u_H(0)$ is bounded above by $u_H(E_L)$. Suppose not, and consider a PBE that gives H its highest utility, across PBEs and beliefs, $\bar{u}_H > u_H(E_L)$. 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 ii), 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 .⁵¹

PROOF OF LEMMA 2

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 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

⁵⁰As before, one can use a PBE that yields a cost arbitrarily close to \bar{Q} , in case it is not exactly achieved.

⁵¹It 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 .

a weakly higher iso-utility curve of u_L than R_n . Moreover, because H can always accept R_{n+1} , $u_H(n) \geq u_H(R_{n+1}) \geq u_H(E_H(R_n))$, where the last 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 . 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 1. 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 $u_L(E_L(R_{n+1}))$. Therefore, $u_L(E_L(R_{n+1}))$ must be weakly greater than $u_L(R_n)$. This, along with the single-crossing property, implies that $u_H(R_{n+1})$ is strictly greater than $u_H(E_H(R_n))$ and contradicts 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 such that $\beta_n = \beta_N > 0$ and $R_n = R_N$ for all $n \geq N$. Indeed, 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 and corresponding belief subsequence $\{\beta_{n(k)}\}_{k \in \mathbb{N}}$ that converges to $\beta^* > 0$. Let $u_H^* = \sup\{u_H(R_n)\}$. For H to accept R_n with positive probability infinitely often, $u_H(R_n)$ must converge to u_H^* , including along the subsequence $\{n(k)\}$.⁵² However, that implies that proposing the H -efficient contract C_H that gives u_H^* to H is again a strictly profitable deviation: it does not change P 's cost conditional on facing L , but it strictly reduces the cost conditional on facing H (by an amount arbitrarily close to $Q(C_L) - Q(C_H)$, where C_θ is the θ -efficient contract that provides H with utility u_H^*), which happens with a probability arbitrarily close to $\beta^* > 0$.

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 is 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 , 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 $u_L(R_{\tilde{m}})$ 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 $u_L(\bar{C}_L) = u_L^*$. ■

C Inequalities

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

$$a(u_H(\hat{C}) - u_H(C)) \leq Q(\hat{C}) - Q(C) \leq a(u_H(\hat{C}) - u_H(C)) \quad (31)$$

$$Q(E_H(R_n)) - Q(E_H(R_{n+1})) \leq b(Q(E_L(R_n)) - Q(E_L(R_{n+1}))) \quad (32)$$

⁵²Otherwise, there must exist a subsequence of rounds for which $u_H(R_{m+1})$ is bounded above away from u_H^* by some constant $\delta_u > 0$. At the same time, H 's continuation utility, $u_H(m)$, is nondecreasing and becomes arbitrarily close to u_H^* . When it gets within $\varepsilon\eta$ for some ε arbitrarily small, this implies that accepting R_{m+1} causes a loss of order $\eta\delta_u$ (due to the probability of a breakdown) and contradicts the fact that $u_H(m)$ is within $\varepsilon\eta$ of u_H^* .

Proof. The efficiency curve \mathcal{E}_H can be parameterized by a univariate parameter λ . Consider two contracts C and \hat{C} on \mathcal{E}_H as in the statement, and set $\lambda = 0$ for C and $\lambda = 1$ for \hat{C} . We have $Q(\hat{C}) - Q(C) = \int_0^1 \frac{\partial Q}{\partial \lambda} d\lambda$, where $Q(\lambda)$ denotes P 's cost along \mathcal{E}_H , and is continuously differentiable, by our assumptions. Because u_H is strictly increasing and continuously differentiable and \mathcal{E}_H is compact as a closed subset of \mathcal{C} , there exists $a > 0$ such that $\frac{\partial Q}{\partial \lambda} \leq a \frac{\partial u_H}{\partial \lambda}$ for all λ , where $u_H(\lambda)$ similarly denotes the utility of H along \mathcal{E}_H . This implies that

$$Q(\hat{C}) - Q(C) \leq a \int_0^1 \frac{\partial u_H}{\partial \lambda} d\lambda = a(u_H(\hat{C}) - u_H(C)).$$

The first inequality in (31) is shown similarly.

Let $\tau : \mathcal{E}_H \rightarrow \mathcal{E}_L$ denote the one-to-one, increasing map which to any H -efficient contract C associates the L -efficient contract which gives H the same utility as C . The map τ is continuously differentiable, and hence Lipschitz on its (compact) domain. Since Q is also continuously differentiable and strictly increasing, there exists $b > 0$ such that

$$Q(E) - Q(E') \leq b(Q(\tau(E)) - Q(\tau(E')))$$

for all E, E' on \mathcal{E}_H . Applying the result to $E = E_H(R_n)$ and $E' = E_H(R_{n+1})$ shows the lemma. \blacksquare

For the next proposition, 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*

$$Q_L \leq Q(E_L(n)) + \frac{\beta_n}{(1 - \beta_n)} a w_n, \quad (33)$$

$$u_H(E_H(n+1)) - u_H(E_H(n)) \leq \frac{\alpha \beta_n}{1 - \beta_n} w_n, \quad (34)$$

$$w_{n+1} \leq w_n \left(1 + \frac{\alpha \beta_n}{1 - \beta_n} \right), \quad (35)$$

$$u_L(R_n) - u_L(R_{n+1}) \leq \gamma \beta_{n+1} w_{n+1}, \quad (36)$$

$$u_H(E_H(n+1)) - u_H(E_H(n)) \geq -\hat{b} \beta_{n+1} w_{n+1}, \quad (37)$$

$$w_{n+1}(1 - b \beta_{n+1}) \geq w_n - \eta y_n. \quad (38)$$

Proof. Lemma 1 implies that

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

Moreover, Q_H is bounded below (by convexity of the cost function) 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. This implies that $Q_L \leq Q(E_L(n)) + \frac{\beta_n}{1 - \beta_n} (Q(E_H(n)) - Q(C_H(n)))$. The contracts $E_H(n)$ and $C_H(n)$ both lie on \mathcal{E}_H . Equation (31) implies that $Q(E_H(n)) - Q(C_H(n)) \leq a(u_H(E_H(n)) - u_H(n)) = a w_n$. This shows (33).

From (33), R_{n+1} cannot give L a utility greater than the L -efficient contract that costs $Q(E_L(n)) + \frac{a \beta_n}{1 - \beta_n} w_n$. This implies that $Q(E_L(n+1)) - Q(E_L(n))$ is bounded above by $\frac{a \beta_n}{1 - \beta_n} w_n$. Combining this with (32) yields

$$Q(E_H(n+1)) - Q(E_H(n)) \leq \frac{ab \beta_n}{1 - \beta_n} w_n$$

Applying the argument used to prove (32), this time to the inverse of τ , we have

$$u_H(E_H(n+1)) - u_H(E_H(n)) \leq \alpha''(Q(E_H(n+1)) - Q(E_H(n)))$$

for some $\alpha'' > 0$. Combining these inequalities yields (34). We have

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

where the first inequality comes from the monotonicity of $u_H(n)$ in n , and the second inequality comes from (34). This shows (35).

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 $Q(E_L(n+1)) + \frac{\beta_{n+1}}{1-\beta_{n+1}}aw_{n+1}$, from (33). Repeating an earlier differential argument, we have $u_L(E) - u_L(E') \leq \alpha_L(Q(E) - Q(E'))$ for all $E, E' \in \mathcal{E}_L$, for some $\alpha_L > 0$. 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 (36).

In general, $u_H(E_H(n+1)) - u_H(E_H(n))$ may be negative. From (36), we have

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

Using that the map τ introduced earlier has a continuously differentiable inverse, this implies that

$$u_H(E_H(n+1)) - u_H(E_H(n)) \geq -\hat{b}\beta_{n+1}/(1-\beta_{n+1})w_{n+1},$$

for some $\hat{b} > 0$, proving (37).

For the last equation, subtracting $u_H(E_H(n))$ from (2) and rearranging (recalling that $w_n = u_H(E_H(n)) - u_H(n)$) leads, along any choice sequence, to

$$w_{n+1} = w_n - \eta y_n + \eta w_{n+1} + (1-\eta)(u_H(E_H(n+1)) - u_H(E_H(n))). \quad (39)$$

Combining this with 37 yields

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

and the hence (38). ■

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

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

Proof. For each $C \in \mathcal{E}_L$, we parameterize the contract set $\mathcal{U}_L(C) = \{R \in \mathcal{H} : u_L(\tilde{C}) = u_L(C)\}$ as $\{C_\lambda\}$ where $\lambda \in [0, 1]$ and $C_0 = C$. Let $Q(\lambda) = Q(C_\lambda)$ and $u_H(\lambda) = u_H(C_\lambda)$. By efficiency of C , we have $Q'(0) = 0$. Because u_L is concave, Q is convex, and u_L 's iso-level curves are nonlinear, there exists $\hat{q} > 0$ such that $Q(\lambda) - Q(0) \geq \hat{q}\lambda^2$ for λ in a right neighborhood of 0. By compactness and convexity of $\mathcal{U}_L(C)$, moreover, \hat{q} may be chosen small enough so that the inequality holds for all $\lambda \in [0, 1]$. Since u_H has bounded derivatives, there must exist $\bar{u} > 0$ such that $|u_H(\lambda) - u_H(0)| \leq \bar{u}\lambda$ (the single-crossing property between u_H and u_L imply that $u_H(\lambda) \leq u_H(0)$ for all $\lambda \in [0, 1]$). 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 $C \in \mathcal{E}_L$. 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. \blacksquare

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

$$y_n^2 \leq k_2[Q(R_{n+1}) - Q(E_L(n))] + k_3(\max\{(\beta_n w_n / (1 - \beta_n))^2, (\beta_{n+1} w_{n+1})^2\} + \beta_{n+1} w_{n+1}) \quad (40)$$

Proof. We have

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

The first inequality is standard ($(a+b)^2 \leq 2a^2 + 2b^2$). The second inequality comes from (34) and (37), which taken together imply an upper bound on $|u_H(E_H(n)) - u_H(E_H(n+1))|$. The third inequality comes from the equality $u_H(E_H(n+1)) = u_H(E_L(n+1))$ and Lemma 12 applied to the contracts $C = E_L(n+1)$ and $R = R_{n+1}$. The difference $Q(E_L(n)) - Q(E_L(n+1))$ is bounded above in proportion to $u_L(R_n) - u_L(R_{n+1})$ (by a simple transposition of the proof of (31)), and that latter difference is bounded above by $\gamma\beta_{n+1}w_{n+1}$, from (36). This shows that

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

which yields the result. \blacksquare

D Proofs for Parts I and II

PROOF OF LEMMA 3 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 4 (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 $\beta_{n_0} \leq \beta_0$. \blacksquare

LEMMA 14 *i) For any round n_0 , $\tilde{\varepsilon} > 0$, and choice sequence, there exists a round $n > n_0$ such that $u_H(R_n) \geq \max\{u_H(E_H(R_m)) : m \leq n_0\} - \tilde{\varepsilon}$. *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 $\tilde{\varepsilon} > 0$. 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 $u_H(R_n) \geq u_H(\bar{C}_L) - \tilde{\varepsilon}$ 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 line \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 4, the argument also applies to L). Since R_n converges to \bar{C}_L and $u_L(n)$ must converge to $u_L(\bar{C}_L)$. Since $u_L(n)$ 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 $u_H(\bar{C}_L)$, from the previous reasoning. \blacksquare

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 when \hat{u}_1 is reached, 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. \blacksquare

For any $\tilde{\varepsilon}$, let

$$D(\tilde{\varepsilon}) = \min\{Q(C) - Q(E) : C \in \mathcal{H}, E \in \mathcal{E}_H : u_H(E) \leq u_H(E_H(C)) + \tilde{\varepsilon}\}. \quad (41)$$

$D(\tilde{\varepsilon})$ is nonincreasing in $\tilde{\varepsilon}$. Moreover, because \mathcal{E}_H and \mathcal{E}_L are smooth and disjoint curves of the compact domain \mathcal{C} , $D(\tilde{\varepsilon})$ is strictly positive for $\tilde{\varepsilon}$ small enough. For such values of $\tilde{\varepsilon}$, $D(\tilde{\varepsilon})$ 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, this implies that R_n must cost weakly more than C : otherwise, we would have $u_L(E_L(R_n)) < u_L(C)$, which would imply that $u_H(n) \leq u_H(E_H(R_n)) < u_H(E_H(C)) = u_H(n(k-1))$, a contradiction. 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 3 guarantees that one can find a later round for which \bar{w}_n is 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 3 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. However, because \bar{w}_n converges to zero along any choice sequence as n goes to infinity, by Lemma 14, there can only be finitely many initializations: there must be 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 .*

PROOF OF LEMMA 5

We fix throughout some small $\bar{\varepsilon} > 0$. Consider some round n contract R_{n+1} in the menu $M_n \cup \{R_n\}$. If $\mu_n^L(R_{n+1}) \geq \bar{\varepsilon}$, then $\mu_n(R_{n+1}) \leq \frac{1}{\bar{\varepsilon}}$, and hence $\beta_{n+1} \leq \frac{\beta_n}{\bar{\varepsilon}}$. 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 $1 - G\bar{\varepsilon}$.

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

From (35) in the Appendix, $w_{n+1} \leq w_n \left(1 + \frac{\alpha\beta_n}{1-\beta_n}\right)$.

At round \bar{n} , we have $\beta_{\bar{n}} \leq \eta^d$ and $w_{\bar{n}} \leq \bar{w}\eta$. Therefore, $\beta_{\bar{n}+1} \leq \eta^d/\bar{\varepsilon}$ and $w_{\bar{n}+1} \leq k_1\eta$ for some constant k_1 . This implies in particular that with probability $1 - G\bar{\varepsilon}$, 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

$$\beta_n \leq (\bar{\varepsilon})^{-\bar{N}} \bar{\beta} \eta^d \tag{42}$$

$$w_n \leq \bar{W}(\bar{N})\eta \tag{43}$$

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 .

Consider choice sequences such that β_n and w_n satisfy the inequalities above throughout the block, which occur with probability $1 - k(\bar{N})\bar{\varepsilon}$. We now show that, at round $n = \bar{n}$, $Q(R_{n+1}) - Q(E_L(n))$ must be of order $O\left(\frac{\beta_n}{\mu_n^L(R_{n+1})}\right)$. If each term in the sum entering P's IC constraint (11) is nonnegative, this result comes from the inequality⁵³

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

⁵³The inequality holds because each term in the sum is nonnegative, and $w_n a \beta_n$ is bigger than the sum.

which implies that

$$Q(R_{n+1}) - Q(E_L(n)) \leq \frac{aw_n}{\eta(1-\beta_n)} \frac{\beta_n}{\mu_n^L(R_{n+1})}. \quad (44)$$

In general, while some terms $\mu_n^L(R_{n+1})(1-\beta_n)\eta(Q(R_{n+1}) - Q(E_L(n)))$ of that sum may be negative, they can only be very slightly so: we have

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

where the second inequality comes (36) of Lemma 11 in the Appendix. Moreover, $w_{n+1} \leq w_n \left(1 + \frac{\alpha\beta_n}{1-\beta_n}\right)$, from equation (35) in the Appendix, and $\mu_n^L(R_{n+1})\beta_{n+1}$ is of order β_n . Therefore, the lower bound is of order $\eta \times w_n\beta_n$, for each term of the sum that is negative. Since there are at most G of them, we conclude that (44) holds up to a term of order $\eta\beta_n$, which is negligible.

This shows that for all R_{n+1} such that $\mu_n^L(R_{n+1}) \geq \sqrt{\beta_n}$, the difference $Q(R_{n+1}) - Q(E_L(n))$ 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 $Q(R_{n+1}) - Q(E_L(n))$ 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

E Proofs for Part III

PROOF OF LEMMA 6

By assumption, $\beta_N \leq \eta^d$ so β_N becomes arbitrarily small as η gets small. Recall equation (35) from Lemma 11

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

where $\alpha > 0$. By 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, for η small enough,

$$\beta_{n+1} \leq \mu_n\beta_n(1 + \epsilon) \quad (45)$$

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 that $w_{N+1} \leq \frac{\eta D}{2a}(1 + (1 + \epsilon)\alpha\beta_n)$. From (13), 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 decreasing⁵⁴ until at least round M . Proceeding by induction, from round N to round M , the previous inequalities imply that

$$w_{N+m} \leq w_N \prod_{i=1}^m (1 + \alpha(1 + \epsilon)\beta_{N+i}) \quad (46)$$

⁵⁴This 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 $\mu_n \leq 1$.

and

$$\beta_{N+i} \leq \beta_N \prod_{j=0}^{i-1} (\mu_{N+j}(1+\epsilon)), \quad (47)$$

and (45) is valid for all rounds $n \in \{N, \dots, M\}$. From (47), we have

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

Therefore, (46) 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

$$\prod_{i=1}^{\infty} \left(1 + \alpha(1+\epsilon)\eta^d \left(\frac{3(1+\epsilon)}{4} \right)^i \right) \quad (48)$$

is finite for η small enough, and converges to 1 as η goes to zero.⁵⁵ Therefore, for η small, w_M is bounded above by $\frac{5}{4}w_N \leq \frac{5\eta D}{8a}$. From (13), 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 (45), 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*

$$y_n^2 \leq \frac{\bar{A}\beta_{n+1}}{1-\beta_0} \quad (49)$$

Proof. Equation (15) implies that $Q(R_{n+1}) - Q(E_L(n)) \leq \frac{\beta_n w_n a}{\eta(1-\beta_0)}$, since $\beta_n \leq \beta_0$. This, along with (16), yields⁵⁶

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

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 7

⁵⁵Indeed, 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.

⁵⁶We are using $\beta_{n+1} \geq \mu_n \beta_n$, which comes from the first inequality of (18).

Taking the square root of (22) 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 (20) yields

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

and, taking the square root of this expression,

$$\frac{\sqrt{\beta_{\hat{N}+2}}}{\beta_{\hat{N}+1}} \geq \frac{\sqrt{c}}{\beta_{\hat{N}+1}^{1/4}} \quad (50)$$

Combining this with (23) shows that (22) holds at round $\hat{N} + 1$. Since β_n is non-increasing in n for $n \geq \hat{N}$ and hence satisfies (23) for all $n \geq \hat{N}$, we can apply the previous argument by induction to conclude that (22) and (50) hold for all $n \geq \hat{N}$. Multiplying (20) 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 (50) applied to round n (instead of $\hat{N} + 1$), the last term is bounded above by $\frac{c\beta_n^{1/4}}{\sqrt{c}}$, which converges to zero as n goes to infinity. \blacksquare

PROOF OF LEMMA 8

Suppose by contradiction that $\{q_n\}$ converges to zero. This, along with the second assumption of the lemma, implies the existence, for any $\varepsilon > 0$, 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 $n \geq \bar{N}$.⁵⁷ 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

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

where $\tilde{c} = c'\sqrt{\bar{\Pi}}$. Taking $K = 3$, this yields

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

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

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

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

⁵⁷Indeed, 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).

F Proof of Theorem 2, Statement B

Let $Q(u, \theta, \varepsilon)$ denote the cost of the cheapest contract that provides θ with utility at least u and that is *not* within a distance ε of the θ -efficient contract $E_\theta(u)$ that provides θ utility with that utility. By a standard Taylor approximation argument (using convexity of Q and strict concavity of u_θ , see proof of Lemma 12), it is easy to show that $Q(u, \theta, \varepsilon) \geq Q(E_\theta(u)) + \varepsilon^2$.

From Part A, there exists a threshold $\tilde{\eta}(\varepsilon^4)$ such that P must leave an expected utility to each type θ that is bounded below by $u_\theta(E_\theta(R_0)) - \varepsilon^4$ for all η 's below that threshold. The cheapest contract $E_\theta = E_\theta(u_\theta(E_\theta(R_0)) - \varepsilon^4)$ that provides this utility costs $Q(E_\theta(R_0)) - O(\varepsilon^4)$. From Lemma 1, 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. These observations imply Part B): if θ accepted contracts that are ε away from $E_\theta(R_0)$, those would be $\varepsilon + O(\varepsilon^4)$ away from E_θ , and incur an excess cost of order ε^2 above $Q(E_\theta(R_0))$. Therefore, if θ accepted such contracts with probability at least ε , the efficiency loss, compared to accepting $E_\theta(u)$ would be of order ε^3 , compared to a maximal potential gain of order ε^4 . Setting $\bar{\eta}(\varepsilon) = \tilde{\eta}(\varepsilon^4)$ then yields the statement of Part B. ■

G 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).
- $E_\theta(n) = E_\theta(R_n)$.
- $w_n = u_H(E_H(n)) - u_H(n)$.
- $\bar{w}_n = \max\{u_H(E_H(m)) : m \leq n\} - u_H(n)$.
- $y_n = u_H(E_H(n)) - u_H(R_{n+1})$.
- $\bar{C}_\theta = \lim_{n \rightarrow +\infty} u_\theta(n)$.

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.
- BERNHEIM, B.D., RAY, D. (1989) “Collective Dynamic Consistency in Repeated Games,” *Games and Economic Behavior*, Vol. 1, pp. 295–326.
- BESTER, L., STRAUZ, 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*.
- 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., SONNENSCHNEIN, 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.
- MAESTRI, L. (2012) “Dynamic Contracting under Adverse Selection and Renegotiation,” Working Paper, Toulouse School of Economics.
- 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., 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.