Available online at www.sciencedirect.com

Games and Economic Behavior 44 (2003) 98–113

GAMES and Economic Behavior

www.elsevier.com/locate/geb

Moral hazard, renegotiation, and forgetfulness

Francesco Squintani

Department of Economics, University of Rochester, Rochester, NY 14627, USA

Received 19 October 1999

Abstract

When a principal and an agent operate with simple contracts, in equilibrium, renegotiation will occur after the agent takes her action. Since renegotiation makes incentive contracts non-credible, the principal may prefer non-renegotiable monitoring options. The current literature does not fully reconcile these predictions with the observation of simple, non-renegotiated incentive contracts. We model a principal-agent interaction in a social learning framework, and assume that when renegotiation is not observed, agents may forget its feasibility with infinitesimal probability. In the unique stable state of our model, renegotiation occurs with infinitesimal frequency, and second-best simple incentive contracts appear with non-negligible frequency.

JEL classification: C73; D82; D83

Keywords: Moral hazard; Renegotiation; Evolution

1. Introduction

This paper characterizes the stable states of a simple social learning model of a moral hazard game. We introduce a minimal departure from the standard assumptions: if renegotiation is not observed, its feasibility may be forgotten with infinitesimal probability. We suggest that this exercise may be helpful in refining the current literature's predictions with respect to simple moral hazard scenarios.

In the classic formulation (Mirrlees, 1976) of the two-action principal-agent problem, for the relevant parameter values, the agent exerts high effort when offered a second-best contract, which is an incentive scheme that makes her indifferent between exerting high or low effort, or rejecting the offer. However, if the principal can renegotiate the contract after

E-mail address: sqnt@troi.cc.rochester.edu.
URL address: http://www.econ.rochester.edu/Squintani/squintani.html.

0899-8256/03/$ – see front matter
doi:10.1016/S0899-8256(03)00014-9

the agent has chosen her effort level and before output is realized, Fudenberg and Tirole (1990) insightfully show that the second-best contract will be renegotiated, and thus will not elicit high effort. When the parties sign a "simple" contract (i.e., a contract that does not require that the agent reports her action to the principal), one can show that the principal will offer renegotiation in equilibrium. In fact, if contracts do not depend on messages, but just on the outcome, the only means to separate high-effort from low-effort workers is to offer a renegotiation that only the latter will accept.

The current literature does not fully reconcile this prediction with real-world intuition about simple moral hazard scenarios. Casual observation suggests that, while contractual renegotiation is not uncommon in complex interactions (involving, say, large companies, and financial institutions),[1] the possibility of a principal's offer of renegotiation does not make incentive contracts ineffective in simple moral hazard scenarios, such as those in which professional services are offered to private citizens. For example, a contractor building a house does not offer a menu-contract, but rather a simple incentive scheme: if the project is delayed, she is subject to some penalty. Similarly, a lawyer in a suit is paid only if she wins her case. Most professional services are subject to incentive contracts, and while the parties could in principle renegotiate such contracts, it does not seem that this possibility makes the incentives ineffective.[2]

One may offer the conjecture that the players form reputations that include commitment not to renegotiate incentive schemes. While the reputation framework (see Kreps and Wilson, 1982, and Milgrom and Roberts, 1982) plausibly explains the rarity of renegotiation in two-player repeated interactions, this need not be the case for complex societies modeled as random-matching games. In fact, when renegotiation is not observable by a third party, the existence of a reputation equilibrium requires that at least one of the players in the match reports whenever a renegotiation has been offered.[3] But, since renegotiation makes both players better off, neither of them has any incentive to threaten to report it, rather the players will cooperate and secretly renegotiate the contract, to their mutual advantage.

This brings us back to the question of how to explain that simple contracts are renegotiated less often than the current literature would predict. An implicit assumption in the original principal-agent model is that the principal will achieve a higher payoff by designing a (second-best) incentive scheme, than by paying a cost in order to monitor the agent's action. When introducing the possibility of renegotiation, there may instead be scenarios where the second-best outcome dominates monitoring, but implementing the latter is more profitable than the solution of Fudenberg and Tirole (1990). In fact, monitoring is implicitly non-renegotiable when its cost is paid before, or while, the agent

[1] Huberman and Kahn (1988), for instance study the renegotiation of loan contracts between banks and large companies. While the company's assets are contracted to be a collateral of the loan, in case of default, the bank will typically renegotiate the contract, and the threat of takeover will not be carried out.

[2] Following Fudenberg and Tirole (1990), we assume that the contracting and renegotiation stages consist of a take-it-or-leave-it principal's offer, rather than allowing for the possibility of alternating counteroffers. In simple moral hazard scenarios, the contractual gains are likely to be too small for the parties to entertain the possibility of a lengthy and costly sequence of alternating offers.

[3] For a folk theorem in random matching games, see Kandori (1992).

takes her action.[4] Thus, unlike renegotiation-proof contracts, it may elicit high effort. Therefore, monitoring can explain why renegotiation is not pervasive, but it does not account for the prevalence of second-best simple contracts.

This paper argues that we often observe non-renegotiated incentive contracts in simple moral hazard scenarios because not everybody is sophisticated enough to include the possibility of renegotiation in their interpretational model of these interactions. We believe however that such an explanation is insufficient for the purposes of economic theory, unless one can demonstrate a social learning model in which the fraction of non-sophisticated individuals does not vanish over time as players learn about renegotiation, and in which non-renegotiated incentive contracts are observed with non-negligible probability at steady state.

The key features of our evolutionary model are as follows. Each period, players in a large population are randomly paired to play a moral hazard problem. Each player's model of the game may be incomplete, and it coincides with her parent's model unless one of two possibilities occurs. The offspring of a player unaware of renegotiation, who is offered it in her match, will *learn* that renegotiation is possible. At the same time, the offspring of an aware player who is not offered renegotiation, may *forget* with small probability that renegotiation is possible. The players do not observe the outcome of matches involving any player other than their parents.[5] The pairings, and the assignment of the role of principal or agent are anonymous, and independent over time.

Our stability concept entails two separate requirements. The first one is that the population play is "in equilibrium," given the fraction of aware players in the society.[6] The second one is that, given the population play, and the induced aggregate awareness transition, the fraction of aware players is stable over time. We show that the stable frequency of non-renegotiated second-best contracts is non-negligible. Whenever too many players are aware of renegotiation, in fact, aware principals choose to monitor. Their opponents do not observe renegotiation, and thus they may forget it. Forgetful principals offer non-renegotiated second-best contracts. For infinitesimal forgetfulness, the stable frequency of renegotiation is infinitesimal. In fact, the fraction of aware players is stable only when as many aware players are forgetting renegotiation as there are forgetful players recalling it. This requires that forgetful agents observe renegotiation only with small probability, when forgetfulness is infinitesimal.

The paper is presented as follows. The second section presents a traditional treatment of moral hazard with simple contracts. The third section presents and discusses our social learning model in details. The fourth section derives the stability results. Some of the proofs are in Appendix A.

[4] Our homeowner, for instance, may go onto the site to monitor the progress of construction. This form of monitoring is hardly renegotiable. Border and Sobel (1987), instead, consider situations in which the monitoring cost is paid after the player takes her action.

[5] Unlike standard evolutionary analysis, our learning dynamics are not payoff-monotone, so that the diffusion of a strategy in the population depends on how often it is used, not on its payoff. Payoff-monotone dynamics were first analyzed by Nachbar (1990). For a general review of the evolutionary literature, see Fudenberg and Levine (1998), or Weibull (1995).

[6] Note that this requires that the aware players correctly assess the distribution of types in the population.

2. Moral hazard with simple contracts

A principal P may motivate a prospective agent A through an incentive contract C or through the use of a monitoring device M. The agent accepts (y) or refuses (n) the principal's proposal and then takes a privately observed action a, consisting of high effort (H) or low effort (L). That action will influence the probability p_a that a high h or a low l output will result. High output will be more likely under high effort (i.e., $p_H > p_L$). The incentive contract C prescribes the agent's compensation profile (c_h, c_l) as a function of the output realized. When using the monitoring device, the principal pays $k > 0$ to know the action taken by the agent, and compensates her with the profile $M = (m_H, m_L)$ as a function of her observed action.

The utility of the agent is $V = U(c) - e(a)$, where $e(H) = e > e(L) = 0$, and c is the compensation received. The function $U(\cdot)$ is continuous, twice differentiable, and satisfies $U'(\cdot) > 0$ and $U''(\cdot) < 0$, with the normalization $U(0) = 0$. The agent's reservation utility is normalized to 0. The principal is a risk-neutral profit-maximizer.

When the parties sign an incentive contract, the agent's Von Neumann–Morgestern expected utility, and the principal's profit are:

$$V(C, a) = (1 - p_a)U(c_l) + p_a U(c_h) - e(a);$$
$$\Pi(C, a) = (1 - p_a)(l - c_l) + p_a(h - c_h).$$

If the principal proposes, and the agent accepts, a monitoring device, the payoffs are:

$$V'(M, a) = U(m_a) - e(a);$$
$$\Pi'(M, a) = (1 - p_a)l + p_a h - m_a - k.$$

To make the problem non-trivial, we assume that the principal prefers to motivate the agent to work hard, or to monitor her, over letting her shirk and giving her no compensation (we denote that contract by 0). That is, there exists a contract C such that $\Pi(C, H) > \Pi(0, L)$ and $V(C, H) \geqslant V(C, L)$, and there exists a compensation profile M such that $V'(M, H) \geqslant V'(M, L)$ and $\Pi'(M, H) > \Pi(0, L)$. For further reference, we denote this game as G_0.

The game G_1 is obtained by expanding G_0 so as to allow for the possibility that the contract C is renegotiated. A renegotiation R is a new contract, proposed by the principal after the agent's action is taken and before the output is realized, that assigns new compensations (r_l, r_h). The payoffs will be $V(R, a)$ and $\Pi(R, a)$ if the agent accepts R, and $V(C, a)$ and $\Pi(C, a)$ if she rejects R. Unlike the incentive contract, the monitoring option is not renegotiable, because its cost is paid before the agent chooses her effort level.

First we solve game G_0. In equilibrium, it is well known that the principal offers the agent the (second-best) contract C^* that solves for $V(C^*, L) = V(C^*, H) = 0$. To motivate the agent to work hard, the principal proposes a contract that makes the agent weakly better off when she exerts high effort. The most profitable contract gives the agent zero informational rent, and (since U is strictly concave) makes the agent indifferent between exerting low and high effort. In equilibrium, the agent accepts the contract C^* and exerts high effort (H); by construction, $\Pi(C^*, H) > \Pi(0, L)$.

In any equilibrium of game G_1, instead, it is well known that incentive contracts fail to motivate the agent to work hard with unit probability, and thus second best cannot be achieved. If the agent exerted high effort in equilibrium, in fact, the principal would renegotiate the initial contract to fully insure the agent. But, anticipating this, the agent should exert low effort in the first place.

In order to present a complete solution of game G_1, we first need to solve the moral hazard problem with renegotiation and simple contracts, but without monitoring. Our results are summarized in Proposition 1.[7] The solution concept is Perfect Bayesian Equilibrium, but as is customary in the renegotiation literature, we restrict attention to equilibria where, in each subgame starting after an initial contract has been accepted, the players coordinate on a Pareto-undominated equilibrium.

Proposition 1. *On the equilibrium path, the principal initially proposes a contract $\overline{C}$ such that $V(\overline{C}, H) = V(\overline{C}, L) \geqslant 0$. The agent accepts it, and plays H with probability $x[\overline{C}] < 1$. The principal renegotiates offering the contract $R = (U^{-1}(V(\overline{C}, L)), U^{-1}(V(\overline{C}, L)))$. The agent accepts the offer if and only if she has played L. The principal's equilibrium profit $\overline{\Pi}$ is strictly smaller than $\Pi(C^*, H)$. If U satisfies non-decreasing absolute risk aversion, then $\overline{C}$ coincides with the second-best contract C^*, and R coincides with the contract* 0.

If, in equilibrium, the principal opts for monitoring, she will choose a contract:[8]

$$M^* = \left(m_H^*, m_L^*\right) \quad \text{such that} \quad 0 = V'(M^*, H) \geqslant V'(M^*, L).$$

We restrict attention to those cases where $\Pi'(M^*, H) > \overline{\Pi}$: the principal's profit for (optimally) monitoring the agent is higher than the profit in the equilibrium of any subgame following an incentive contract with renegotiation. Thus along the equilibrium path of game G_1, the principal offers M^*, the agent accepts and plays H.

3. Social learning

The players of a continuous large population of players live for two periods. In the second period, they are randomly paired to myopically play game G_1. At the first period they observe their parents' play. Matching and role assignment are anonymous, and independent across generations. After being matched, each player formulates a model of

[7] Our result is closely related to the analysis of Fudenberg and Tirole (1990) on moral hazard and renegotiation with menu-contracts. They show that in equilibrium, the agent chooses low effort with positive probability, and that the optimal contract includes a safe scheme for agents who report low effort, and an incentive scheme for agents who report high effort. As long as the coefficient of absolute risk aversion is non-decreasing, the optimal initial contract gives the agent zero informational rent. Under decreasing absolute risk aversion, the equilibrium rent may be positive.

[8] In fact, by assumption, there exists a compensation profile M such that $V'(M, H) \geqslant V'(M, L)$ and $\Pi'(M, H) > \Pi(0, L)$. By definition, $U(m_H^*) = e$, so M^* maximizes $\Pi(M, H)$, as long as the agent accepts M^* and plays H after doing so. In equilibrium, the agent cannot do otherwise, or else the principal would not have a well-defined best-response. Since for any $m_H > m_H^*$ the agent's sequentially rational response is to accept M and work hard, the profit would be strictly decreasing in m_H for any $m_H < m_H^*$, and would discontinuously drop at m_H^*.

the game, the model may or may not include the possibility that the principal proposes a contractual renegotiation, in the latter case, we say that the player is forgetful. For any time t, we denote as ρ_t the fraction of players aware of renegotiation at time t.

The players unaware of renegotiation must also be unaware that they or their opponents could be aware of it, or else, they could not be unaware of renegotiation in the first place. The players aware of renegotiation are also aware that their opponents could be aware or unaware of it. Since players cannot observe their opponent's state of mind, each aware player formulates a conjecture of the probability of facing an aware opponent. At the same time, each player formulates conjectures on her opponent's play, depending on her conjectured type. Given her conjectures, each player chooses a sequentially rational course of action. We assume that, when presented with an unforeseen renegotiation offer, forgetful agents do not revise their probability assessment of the occurrence of h and l.

Each player enters into play holding her parent's model of the game, except in one of the two following cases. The offspring of a forgetful agent becomes aware of the possibility of renegotiation if (and only if) her parent is offered renegotiation on the path of play.[9] The offspring of an aware agent who does not receive a renegotiation offer will become forgetful with probability ε. The offsprings of players in the role of principal do not directly observe opponents' actions. We thus assume that they maintain their parents' model of the game. At any time t, the society is described by the current fraction of aware players in the population, and by the current population play. The evolution of the society is induced by the aggregate awareness transition given the population play, and by the change of the population play given the awareness transition.

Our stability definition entails two separate requirements. The first one is that, given the stable fraction of aware players ρ, the population play is "in equilibrium" in the sense that all players' conjectures are correct. In particular, this requires that forgetful players' conjectures about the strategy of forgetful players are correct, and that aware players' conjectures about the value of ρ, and the strategies of both forgetful and aware players are correct.[10] As is customary in evolutionary game theory, we restrict attention to symmetric equilibria, and assume that all players with the same roles and states of mind choose the same strategy. Our second requirement is that given equilibrium population play, the fraction of aware players ρ is (asymptotically) stable under the learning dynamics.

In order to represent population play in equilibrium, we shall make use of the Harsanyi model of games of incomplete information with subjective priors (see Harsanyi, 1967). We introduce a state space and let nature choose the state according to a prior distribution. We attach a game to each state, and for each player we introduce a type space that consists of a partition of the state space. The game is completed by assigning to each player beliefs over the realization of the state; for ease of presentation, we specify beliefs as conditional on types, rather than as prior beliefs.

---

[9] In principle, forgetful players may suspect that their model is incorrect also when they are unexpectedly offered an initial contract different from the second-best one. We explore this possibility in the next section.

[10] Forgetful player, on the other hand, cannot correctly assess the fraction of aware players in the population, or the aware players' strategies, as they are not even aware that their opponents could be aware of renegotiation.

Definition 1. Let the Augmented Game $AG(\rho)$ consist of the Harsanyi game with state space $S = \{AA, AF, FA, FF\}$, where

- the nature's choice is $p(AA) = \rho^2$, $p(FA) = p(AF) = (1-\rho)\rho$, $p(FF) = (1-\rho)^2$;
- a copy of the game G_1 originates in the states AA, AF, and a copy of the game G_0 originates in the states FA, FF;
- the principal's types are $P_A := \{AA, AF\}$, and $P_F := \{FA, FF\}$, and the principal's conditional beliefs are $p_P(AA \mid P_A) = \rho$, $p_P(FF \mid P_F) = 1$;
- the agent's types are $A_A := \{AA, FA\}$, and $A_F := \{AF, FF\}$, and the agent's conditional beliefs are $p_A(AA \mid A_A) = \rho$, $p_A(FF \mid A_F) = 1$.

We say that the population play is *in equilibrium* given the fraction of aware players ρ, if it is identified by a Perfect Bayesian Equilibrium $(\sigma_{P_A}, \sigma_{P_F}; \sigma_{A_A}, \sigma_{A_F})$ of the augmented game $AG(\rho)$, such that the players coordinate on a Pareto-undominated equilibrium in each subform starting after an initial contract C has been accepted. We say that the pair (ρ, σ) is *stable* if the population play σ is in equilibrium given the fraction of aware players ρ, and if ρ is (*asymptotically*) stable[11] under the aggregate learning transition induced by σ.

We conclude this section by discussing the equilibrium concept introduced in Definition 1. In order to describe symmetric equilibrium population play, we look at a single representative pair of players. There are four possible combinations of states of mind, and these combinations make up the state space introduced in Definition 1. The first letter of each state represents the state of mind of the principal, and the second the agent's. The letter A stands for aware, while the letter F stands for forgetful. The nature's prior assigns the frequencies of the four combinations of states of mind in the population.

If a player is aware, she knows that either her opponent is aware, or that she is forgetful. In equilibrium, she correctly assesses ρ, the fraction of aware players in the population. These assessments are represented by the conditional beliefs introduced in Definition 1. Since $p_P(AA \mid P_A) = \rho$ and $p_A(AA \mid A_A) = \rho$, conditional on being the aware type, each player believes she is playing against an aware type of opponent with probability ρ. If a player is forgetful, she chooses her course of actions believing to play a game in which it is common knowledge that renegotiation is impossible. In Definition 1, since $p_P(FF \mid P_F) = p_A(FF \mid A_F) = 1$, the forgetful types believe with probability 1 that they are playing against a forgetful type, and thus they play *as if* it were common knowledge that renegotiation were impossible.

The use of subjective priors is required to make sure that forgetful players do not condition their choice on the strategies of aware players in the population. This would be impossible if the forgetful players knew the fraction of aware players ρ. In such a case, the model would be logically inconsistent. If the forgetful players were to condition their choice on the possibility that their opponents may be aware of renegotiation, in fact, they could not be unaware of renegotiation in the first place.

Finally, we should say that the description of the players' beliefs presented in Definition 1 is to be taken as a primitive. The state-space and the subjective priors are

[11] A formal definition of asymptotic stability may be found in Weibull (1995, Chapter 6).

only introduced to generate a representation of the population equilibrium play, and they are not supposed to be subject to epistemic reasoning by the players. If that were not the case, upon knowing that the state space includes the possibility that the opponent is aware of renegotiation, forgetful players should infer that renegotiation is possible after all.

4. Stability results

Proposition 2 characterizes the unique stable description of the society with infinitesimal forgetfulness. The stable fraction of aware players $\bar{\rho}$ is independent of the probability of forgetting renegotiation ε (as long as $0 < \varepsilon < 1$), and is determined as follows. Suppose that the population play is in equilibrium. Whenever too many players are aware of renegotiation, aware principals choose to monitor. Their opponents do not observe renegotiation, and thus they may forget it. When the proportion of forgetful players is large enough, aware principals offer an incentive contract, and then renegotiate it. The fraction of aware players is thus stable only if it makes aware principals indifferent between implementing the optimal monitoring scheme M^*, and offering an incentive contract that they subsequently renegotiate.

We will show that in the unique stable description of the society, the most profitable incentive contract C does not yield any informational rent to the agents. Despite this, all agents accept contract C, if they are offered it. Aware agents then play L, and forgetful ones play H. At the renegotiation stage, aware principals assign probability of $(1-\bar{\rho})p_H + \bar{\rho}p_L$ to the realization of output h. They offer the contract R^* that maximizes their expected profit

$$\Pi(R, \bar{\rho}) = \left[1 - \bar{\rho}p_L - (1-\bar{\rho})p_H\right](l - r_l) + \left[(1-\bar{\rho})p_H + \bar{\rho}p_L\right](h - r_h),$$

subject to the condition that $V(R, H) = 0$, so that the renegotiated contract R^* is accepted both by aware and forgetful agents. We denote by CR^* the strategy of initially offering a zero-rent incentive contract C, and then renegotiating it with R^*.

The intuition for these results is along the following observations. First note that it cannot be the case that in a stable society, aware principals renegotiate an initial contract and offer a contract accepted only by the agents who shirked. Following the logic of Proposition 1, an aware principal offers such a contract only if she expects her opponent to exert high effort with probability no larger than $x[C^*]$. But if this is the case, by construction, aware principals strictly prefer to monitor, and this cannot be a stable description of the society. Second, in a stable society aware agents must shirk, because they expect in equilibrium that aware principals renegotiate the initial contract so as to make better off the agents who shirked. Third, since in any stable society forgetful agents work hard when offered zero-rent incentive contracts, and aware agents shirk regardless of the initial contract, it is a waste of resources to initially propose a contract that yields positive rents to the agents.

An immediate consequence of the above results is that the stable fraction of aware players $\bar{\rho}$ is strictly smaller than $1 - x[C^*]$. Otherwise, aware principal would be able to elicit high effort only with probability smaller than $x[C^*]$. But then, aware principals strictly prefer to monitor, and this cannot be a stable description of the society. For $\bar{\rho}$ to be stable, moreover, it must be a stationary state, and thus there must be as many aware agents forgetting renegotiation, as there are forgetful agents observing a renegotiation offer on path of play. When ε is small enough, it is thus required that forgetful agents learn the possibility of renegotiation with infinitesimal probability, and thus that aware principals play CR^* with small probability. Forgetful principals implement the second-best contract without renegotiation (we denote this outcome by C^*), that thus occur with frequency $1-\bar{\rho}$.[12]

Proposition 2. *For any $\varepsilon > 0$ small enough, the unique stable pair $(\bar{\rho}, \bar{\sigma})$ is such that $\bar{\rho}$ and R^* solve the system*

$$\bar{\rho} = \frac{\Pi(R^*, H) - \Pi'(M^*, H)}{\Pi(R^*, H) - \Pi(R^*, L)},$$
$$\frac{(1-p_H)U(r_l^*)}{p_H U(r_h^*)} = \frac{(1-\bar{\rho})(1-p_H) + \bar{\rho}(1-p_L)}{(1-\bar{\rho})p_H + \bar{\rho}p_L}, \qquad V(R^*, H) = 0,$$

and such that $\bar{\sigma}_{P_F}(C^*) = 1$, $\bar{\sigma}_{P_A}(M^*) = 1 - \bar{\sigma}_{P_A}(CR^*)$, $\bar{\sigma}_{P_A}(CR^*) = \varepsilon/(1-(1-\varepsilon)\bar{\rho})$. The asymptotically stable fraction of aware players $\bar{\rho}$ is strictly smaller than $1 - x[C^*]$.

Before proving Proposition 2, we present our final results in terms of the aggregate principals' distribution of play in the stable society, calculated by compounding the type distribution with the principals' type strategies. Formally, given any equilibrium population play σ, and fraction of aware players ρ, the principals' aggregate distribution of play is $f = \rho\sigma_{P_A} + (1-\rho)\sigma_{P_F}$.

When the probability of forgetting renegotiation is strictly positive but infinitesimal, the stable frequency of renegotiation will be negligible, and both non-renegotiated, second-best incentive contracts, and monitoring contracts will be observed with non-negligible, stable frequency. Specifically, when taking limits as $\varepsilon \to 0$, the aggregate principals' distribution of play in the stable society is:

$$\bar{f}(C^*) = 1 - \bar{\rho} = \frac{\Pi'(M^*, H) - \Pi(R^*, L)}{\Pi(R^*, H) - \Pi(R^*, L)} > 0, \qquad \bar{f}(CR^*) = \bar{\rho}\bar{\sigma}_{P_A}(CR^*) \to 0,$$
$$\bar{f}(M^*) = \bar{\rho}\big(1 - \bar{\sigma}_{P_A}(CR^*)\big) \quad \to \quad \frac{\Pi(R^*, H) - \Pi'(M^*, H)}{\Pi(R^*, H) - \Pi(R^*, L)} < 1 - x[C^*].$$

We conclude the presentation of our final results with the following remark. Proposition 2 shows that in equilibrium, forgetful agents expect that their opponents will offer them a second-best incentive contract. Thus a forgetful player may suspect that her model of the game is incorrect when her opponent makes use of a monitoring device. In such a case, she may rationalize her opponent behavior in many possible ways. For instance, she may believe that the principal's choice is the result of an idiosyncratic tremble. Alternatively, she may think that her opponent has privately discovered a cheaper monitoring technology,

[12] Also note that the stable frequency of second-best contracts is *increasing* in the profit of the monitoring option. When monitoring is more valuable, in fact, aware principals will choose to monitor for lower proportions of forgetful players in the population.

that makes her more willing to monitor, rather than to offer an incentive contract. However, in the case that she rationalizes the principal's behavior by inferring that the principal has the option to renegotiate incentive contracts, our results change. In such a case, in fact, one can show that in the unique stable society all players are aware of renegotiation, and hence the analysis of the second section applies. Since the play is myopic and monitored agents' behavior does not depend on awareness of renegotiation, in fact, it is still the case that, for any given awareness population state, aware principal either choose monitoring or propose an incentive contract that they subsequently renegotiate. As a result, all their opponents are immediately made aware of renegotiation, and thus any awareness state with a strictly positive fraction of forgetful players must be unstable.[13]

Proof of Proposition 2. The analysis is conducted by first pinning down the population play in equilibrium σ as a function of the fraction of aware players ρ, by then deriving the law of motion of ρ induced by the associated equilibrium play, and finally by determining the asymptotically stable states of the resulting law of motion.

Lemma 3, proven in Appendix A, characterizes the principal's equilibrium play in the game $AG(\rho)$ for any ρ. For brevity, the statement of the lemma does not report the play off path, or the agents' play, which are derived in the proof.

Lemma 3. *In equilibrium, the forgetful principal plays C^* and does not renegotiate it. For any ρ, the aware principal either plays M^*, or initially proposes a contract C such that $V(C,L) \leqslant V(C,H) = 0$, and then renegotiates it with the contract R_ρ, defined as the contract R that solves*

$$\frac{(1-\rho)p_H + \rho p_L}{(1-\rho)(1-p_H) + \rho(1-p_L)} = \frac{p_H U(r_h)}{(1-p_H)U(r_l)}$$

together with $V(R,H) = 0$ (we denote that course of actions by CR_ρ). When $\rho > \bar{\rho}$ she plays M^*, when $\rho < \bar{\rho}$ she plays CR_ρ, when $\rho = \bar{\rho}$ she is indifferent. The threshold $\bar{\rho} \in (0, 1 - x[C^*])$.

For brevity, we henceforth denote the aware principal's strategies $\bar{\sigma}_{P_A}(M^*)$ and $\bar{\sigma}_{P_A}(CR_\rho)$ associated with the fraction of aware players ρ, by $\sigma_M(\rho_t)$ and $\sigma_{CR}(\rho_t)$, respectively, and we omit star superscripts. The aggregate principals' distribution of play at any period t is therefore:

$$f_t(M) = \rho_t \sigma_M(\rho_t), \qquad f_t(CR_\rho) = \rho_t \sigma_{CR}(\rho_t), \qquad f_t(C) = 1 - \rho_t.$$

Players are matched and assigned roles in such a way that a fraction ρ_t of the agents are aware and a fraction $1 - \rho_t$ are forgetful. The principals play either M, CR_ρ, or C, in proportions $f_t(M)$, $f_t(CR_\rho)$, and $f_t(C)$, respectively. This gives six types of pairs of players. In two of them (when an aware agent meets a principal playing M or C),

[13] Forgetful players are also surprised when offered a zero-rent incentive contract that does not coincide with the second-best contract C^*. Again, they may rationalize such an offer by inferring that the principal has the option to renegotiate incentive contracts. In such a case, our conclusions still hold, if we select the equilibria where aware principals offer contract C^* in Proposition 2, instead of any arbitrary zero-rent incentive contract.

a fraction ε of the agents (hence $\varepsilon/2$ of the players in such pairs) forget. In one of them (when a forgetful agent meets a principal who plays CR_ρ), the agents (one half of the involved players) become aware. Given these transition probabilities, by Alos-Ferrer (1999, Theorem 6.4),[14] we can approximate the stochastic evolution of a large population of players with the following difference equation in ρ_t:

$$\rho_{t+1} = \rho_t + \frac{1}{2}\big\{(1-\rho_t) f_t(CR_\rho) - \varepsilon\rho_t\big[f_t(C) + f_t(M)\big]\big\}.$$

For expositional purposes, we analyze the problem in continuous-time, rather than discrete-time, dynamics. As is customary, see Hale (1969), this is done by assuming that only a fraction h of the population is called to play during the time $[t, t+h)$, and for the remaining $1-h$ fraction of the population nothing changes. This is a mere convex combination of the above equation and full inertia, i.e.,

$$\rho_{t+h} = (1-h)\rho_t + h\left(\rho_t + \frac{1}{2}\big\{(1-\rho_t) f_t(CR_\rho) - \varepsilon\rho_t\big[f_t(C) + f_t(M)\big]\big\}\right),$$

and taking limits when $h \to 0$ one obtains the following differential equation:

$$\dot\rho_t = \frac{1}{2}\big\{(1-\rho_t) f_t(CR_\rho) - \varepsilon\rho_t\big[f_t(C) + f_t(M)\big]\big\}. \tag{1}$$

The evolution of Eq. (1) at time t depends on f_t which in turn depends on the equilibrium σ_t characterized in Lemma 3. Since Lemma 3 yields one equilibrium for $\rho < \bar\rho$, and a different one for $\rho > \bar\rho$, the equation is discontinuous. However, as the system is piecewise continuous, one can apply standard techniques to the segments $\rho > \bar\rho$ and $\rho < \bar\rho$, and then complete the analysis considering the discontinuity point $\bar\rho$.

For $\rho \in [0, \bar\rho)$, we obtain:

$$2\dot\rho = \rho(1-\varepsilon)(1-\rho),$$

and for $\rho \in (\bar\rho, 1]$, we obtain:

$$2\dot\rho = -\varepsilon\rho.$$

Since $\dot\rho > 0$ for any $\rho \in (0, \bar\rho)$, and $\dot\rho < 0$ for any $\rho \in (\bar\rho, 1]$, the equation admits only the stationary state $p = 0$ on the segment $[0, \bar\rho) \cup (\bar\rho, 1]$. But clearly, the state $\rho = 0$ is unstable because $\dot\rho > 0$ on $(0, \bar\rho)$.

So, the only candidate asymptotically stable state left is $\bar\rho$. By Lemma 3, any value $\sigma_{CR}(\bar\rho) \in [0, 1]$ is possible in equilibrium. In order to have a well-defined differential equation, one needs to select a unique $\sigma_{CR}(\bar\rho)$. For the equation to have a stable state, we need $\dot\rho = 0$ at $\bar\rho$. Solving out, we obtain:

$$\sigma_{CR}(\bar\rho) = \frac{\varepsilon}{1-(1-\varepsilon)\bar\rho}.$$

[14] Alos-Ferrer (1999) constructs matching schemes under which a continuous population stochastic evolution may be approximated with a dynamic system. Boylan (1992), proposes a similar result for countably infinite populations, with an argument often referred to as a "Law of Large Numbers" in evolutionary games.

When selecting that value, the above analysis shows that the state $\bar{\rho}$ is stationary, asymptotically stable and a global attractor. Also since $\lim_{\rho\uparrow\bar{\rho}}\dot{\rho} > 0$, and $\lim_{\rho\downarrow\bar{\rho}}\dot{\rho} < 0$, the state $\bar{\rho}$ is reached in finite time for any initial state $\rho \in (0, 1]$.

We conclude with a remark concerning the case when $\sigma_{CR}(\bar{\rho}) \neq \varepsilon/(1-(1-\varepsilon)\bar{\rho})$. In this case, the dynamics converge in finite time to $\bar{\rho}$; but, each time that this state is reached, the dynamics discontinuously jump away. The state $\bar{\rho}$ is not stationary, hence not asymptotically stable, but it is still the case that the average frequency of ρ_t over time is close to $\bar{\rho}$. Thus, the average frequency of second-best non-renegotiated $f_t(C)$ (equal to $1-\rho_t$) over time is close to $1-\bar{\rho}$.

The frequency of monitoring and renegotiation at $\bar{\rho}$ depend on the particular $\sigma_{CR}(\bar{\rho})$ selected to complete Eq. (1). Nevertheless, for $\sigma_{CR}(\bar{\rho}) > \varepsilon/(1-(1-\varepsilon)\bar{\rho})$, the dynamics jump into the region $\rho \in (\bar{\rho}, 1)$, where $\sigma_{CR}(\rho) = 0$. Thus the average frequency of renegotiation over time is approximately zero. Since

$$\lim_{\varepsilon\to 0}\frac{\varepsilon}{1-(1-\varepsilon)\bar{\rho}} = 0,$$

we conclude that with infinitesimal ε, under Eq. (1), for almost any selection $\sigma_{CR}(\bar{\rho})$, the average frequency of renegotiation over time is negligible. □

Acknowledgments

I wrote this paper while a PhD student in the Department of Economics at Northwestern University. I thank the Associate Editor, two anonymous referees, Eddie Dekel, Jeff Ely, Steve Matthews, John Panzar, Lars Stole, Juuso Valimaki, and John Woodruff. The usual caveat applies.

Appendix A. Omitted proofs

Proof of Proposition 1. The proof consists of 3 different steps.

Step 1. The principal's equilibrium renegotiation offer, given any initial contract C, and any agent's effort strategy x.

At the moment in which she is offered renegotiation, the agent knows the realization of her mixed strategy x; for brevity, we denote her as the a-agent if she played action a. For any contract C' we denote by $T(C', a) = p_H c'_h + (1-p_H)c'_l$, the expected transfer of the principal to the a-agent.

First suppose that the initial contract C is such that $c_h > c_l$. In such a case, since $p_H > p_L$, and $U'' < 0$, it follows that:

$$\begin{aligned} &V(R,L) > V(C,L) \qquad \text{for any } R \text{ such that } V(R,H) = V(C,H) \text{ and } T(R,H) < T(C,H),\\ &V(R,H) < V(C,H) \qquad \text{for any } R \text{ such that } V(R,L) = V(C,L) \text{ and } T(R,L) < T(C,L). \end{aligned} \tag{A.1}$$

The set of conditions (A.1), together with the fact that $U'' < 0$, implies that the contract $R^L[C] = (U^{-1}(V(C,L)), U^{-1}(V(C,L)))$ maximizes the principal's profit

$$\Pi^L(R,x) := x\Pi(C,H) + (1-x)\big[(1-p_L)(l-r_l) + p_L(h-r_h)\big],$$

among the offers R accepted in equilibrium only by the L-agent.

Secondly, the set of conditions (A.1) implies that the optimal offer among the renegotiation contracts accepted in equilibrium only by the H-agent is the initial contract C itself.

Finally, the set of conditions (A.1) allows us to pin down the contract $R[C,x]$ that maximizes the principal's profit

$$\Pi(R,x) := \big[x(1-p_H)+(1-x)(1-p_L)\big](l-r_l)+\big[xp_H+(1-x)p_L\big](h-r_h).$$

among the offers R accepted in equilibrium by both the H-agent and the L-agent.

Consider the principal's rate of substitution of r_h for r_l,

$$crs(x)=\frac{xp_H+(1-x)p_L}{x(1-p_H)+(1-x)(1-p_L)},$$

and the H-agent's marginal rate of substitution at the contract C,

$$MRS=\frac{p_H U(c_l)}{(1-p_H)U(c_h)}.$$

As long as $crs(x) \geqslant MRS$, the offer $R[C,x]$ coincides with the contract R that solves:

$$\frac{p_H U(r_l)}{(1-p_H)U(r_h)}=\frac{xp_H+(1-x)p_L}{x(1-p_H)+(1-x)(1-p_L)}, \qquad V(R,H)=V(C,H). \tag{A.2}$$

When $crs(x) < MRS$, the optimal offer $R[C,x]$ coincides with the initial contract C.

In order to determine the optimal renegotiation offer as a function of x and C, first notice that $\Pi^L(R^L[C],x) > \Pi^L(C,x)$ as long as $x<1$. Since $\Pi^L(C,x)=\Pi(C,x)$, it follows that $R^L[C]$ dominates C (the optimal offer accepted only by the H-agent), and that $R^L[C]$ dominates $R[C,x]$ whenever $crs(x) \leqslant MRS$.

Secondly, when $x=1$, it is the case that $\Pi^L(R^L[C],x)=\Pi(C,x)$, that $crs(x)=p_H/(1-p_H)$, and hence that $R[C,x]=(U^{-1}(V(C,H)+e), U^{-1}(V(C,H)+e))$. Since $U''<0$, it follows that $\Pi(R[C,x],x) > \Pi(C,x)$. Thus $R[C,x]$ strictly dominates $R^L[C]$ for $x=1$.

The function $\Pi(R[C,x],x)$ is continuous, strictly increasing in x, and (as U is strictly concave) it is a strictly convex function of x. The function $\Pi^L(R^L[C],x)$ is linear in x. Thus, denoting as $x'[C]$ the unique $x \in (0,1)$ that solves:

$$\Pi\big(R[C,x],x\big)=\Pi^L\big(R^L[C],x\big),$$

it follows that in equilibrium the principal plays $R^L[C]$ for any $x<x'[C]$, she plays $R[C,x]$ for any $x>x'[C]$, and is indifferent for $x=x'[C]$.

When the initial contract C is such that $c_h \leqslant c_l$, the problem is symmetric to the case in which $c_h > c_l$. It follows that, depending on x, the principal either plays $R^H[C]=(U^{-1}(V(C,H)+e), U^{-1}(V(C,H)+e))$ or she plays $R[C,x]$, and in either case, $r_h \leqslant r_l$. Finally, for any C such that $c_h=c_l$, the equilibrium renegotiation offer coincides with the initial contract C, for any effort strategy x.

Step 2. The optimal choice of $x[C]$ given the initial contract C and equilibrium renegotiation choices.

When the initial contract C is such that $c_h \leqslant c_l$, Step 1 shows that in equilibrium $r_h \leqslant r_l$, thus the agent sets $x[C]=0$.

Suppose the initial contract C is such that $c_h > c_l$. Denote by $V^*[C,x,a]$ the expected utility of the a-agent (for $a \in \{H,L\}$) given the initial contract C, the effort strategy x, and the equilibrium renegotiation choices. The results from Step 1 yield:

$$\begin{aligned} &V^*[C,x,L] > V(C,L) \quad \text{for any } x>x'[C], \qquad V^*[C,x,L]=V(C,L) \quad \text{for any } x<x'[C],\\ &V^*[C,x,L] \geqslant V(C,L) \quad \text{for } x=x'[C], \qquad V^*[C,x,H]=V(C,H) \quad \text{for any } x\in[0,1]. \end{aligned} \tag{A.3}$$

The set of conditions (A.3) implies that if the contract C satisfies $V(C,L) > V(C,H)$, then $x[C]=0$.

If the contract C satisfies $V(C,H)=V(C,L)$, the set of conditions (A.3) implies that $x[C] \leqslant x'[C]$; but since we are assuming that in the subgame following C, the players play a Pareto-undominated equilibrium, it follows that the agent must play $x'[C]$.

Finally, if the contract C satisfies $V(C,H) > V(C,L)$, the set of conditions (A.3) implies that $x[C] \in [x'[C],1)$. Thus $x[C]$ must coincide with the strategy $x \in [x'[C],1)$ satisfying $V(R[C,x],L) = V(R[C,x],H)$.

Step 3. The optimal choice of the initial contract C given the equilibrium effort strategies $x[C]$, and equilibrium renegotiation choices.

Denote by $\Pi^*[C]$ the principal's expected equilibrium profit after the renegotiation stage, when the initial contract is C.

Suppose that C satisfies $V(C,L) > V(C,H)$. Steps 1 and 2 then imply that $\Pi^*[C] = \Pi(C,L)$. In equilibrium, the agent accepts C only if $V(C,L) \geqslant 0$. Thus the contract C is dominated by the second-best contract C^*, which elicits equilibrium agent's effort $x[C^*] > 0$, and is renegotiated in equilibrium with $R^L[C^*]$ (which coincides with the contract 0).

If C satisfies $V(C,L) \leqslant V(C,H)$, Steps 1 and 2 imply that $\Pi^*[C] = \Pi(R[C,x[C]],x[C])$. Pick an arbitrary $\overline{V} \geqslant 0$, introduce the set $\mathcal{C} = \{C\colon\ V(C,L) < V(C,H) = \overline{V}\}$, and let C' be the contract satisfying $V(C',L) = V(C',H) = \overline{V}$. We will show that $\Pi^*[C'] > \Pi^*[C]$ for any $C \in \mathcal{C}$.

Step 2 shows that for any $x \geqslant x[C']$, it is the case $V(R[C',x],L) > V(C',L) = \overline{V}$. Condition (A.2) implies that $R[C,x] = R[C',x]$ and $V(C,H) = V(R[C,x],H)$ for any x and $C \in \mathcal{C}$. Thus

$$V\big(R[C,x],L\big) > \overline{V} = V(C,H) = V\big(R[C,x],H\big) \quad \text{for any } x \geqslant x[C'] \text{ and any } C \in \mathcal{C}.$$

For any contract $C \in \mathcal{C}$, the above result (together with the last result of Step 2) implies that $x[C] < x[C']$, and thus that $\Pi^*[C] < \Pi^*[C']$, because the profit $\Pi(R[C,x],x)$ is strictly increasing in x.

So we are left to compare $\Pi^*[C]$ across contracts C such that $V(C,L) = V(C,H) = \overline{V}$. When $-U''/U'$ is non-decreasing, since U is concave, for any fixed x, the H-agent's marginal rate of substitutions $crs(x)$ does not increase in $\overline{V}$. It follows that the equilibrium threshold $x[C]$ is also non-increasing in the rent $\overline{V}$. Since $\Pi^*[C]$ is strictly decreasing in $\overline{V}$, and strictly increasing in x, it follows that the optimal initial contract must yield rent $\overline{V} = 0$, and thus coincides with the second-best contract C^* in equilibrium.[15]

The result that $\overline{\Pi} < \Pi(C^*,H)$ follows from the fact that $x[\overline{C}] < 1$, and that $\Pi(C^*,H) > \Pi(0,L)$. □

Proof of Lemma 3. The results of the second section yield the behavior of forgetful types, and imply that we can rule out any monitoring contract other than M^* and say that both types of agents play H after M^*.

Step 1. Equilibrium choices when the aware type of principal does not offer M^*.

First of all note that for any initial contract, the effort equilibrium choice of the agent of type A_A is determined by the equilibrium choices of the principal of type P_A. In fact, upon being initially offered a contract $C \neq C^*$, the agent of type A_A infers that her opponent is of type P_A. If she is offered the contract C^*, her opponent may be of type P_F. But in such a case, she knows that her expected payoff is $0 = V(C^*,H) = V(C^*,L)$ regardless of the action she takes.

If P_A initially proposes an (individually rational) contract C such that $V(C,H) < V(C,L)$, both types of agents play L. Thus the principal is strictly better off initially offering the second-best contract C^*, and then renegotiating with the contract 0. So we henceforth consider initial contracts C such that $V(C,H) \geqslant V(C,L)$.

Consider the equilibrium choice of P_A at the renegotiation stage, given that the initial contract is C, and agent of type A_A plays H with probability σ_H. The principal of type P_A knows her opponent has played H with probability $\mu = 1 - \rho + \rho\sigma_H$. Substituting μ for x in the proof of Proposition 1 (Step 1), we obtain that P_A offers $R[C,\mu]$ whenever $\mu > \mu'[C]$, and $R^L[C]$ if $\mu < \mu'[C]$.

For any contract C, consider the game G_1 equilibrium strategy $x[C]$ determined in Proposition 1. Also, for any fraction ρ, let $\sigma_H[C,\rho]$ be the equilibrium effort strategy of A_A, introduce $\mu[C,\rho] := 1 - \rho + \rho\sigma_H[C,\rho]$, and let $\Pi^*(C,\rho)$ be the principal's equilibrium expected profit.

[15] If $-U''/U'$ is strictly decreasing, then $x[C]$ is strictly decreasing in $\overline{V}$. Thus for $(p_H - p_L)(h - l)$ large enough, the optimal contract $\overline{C}$ yields the agent $\overline{V} > 0$.

Case 1. $\rho < 1 - x[C]$.

For any initial contract C, it follows that $\mu[C, p] \geqslant 1 - \rho > x[C]$ and P_A renegotiates C offering $R[C, \mu[C, \rho]]$. Since $\mu[C, \rho] > x[C]$, from the proof of Proposition 1 (Step 2), we know that

$$V\big(R\big[C, \mu[C, \rho]\big], L\big) > V(C, L) = V(C, H) = V\big(R\big[C, \mu[C, \rho]\big], H\big),$$

so that $\sigma_H[C, \rho] = 0$. It thus follows that $\mu[C, \rho] = 1 - \rho$ and $\Pi^*(C, \rho) = \Pi(R[C, 1 - \rho], 1 - \rho)$.

For any arbitrary $\overline{V}$, let $\mathcal{C}(\overline{V}) := \{C\colon V(C, L) \leqslant V(C, H) = \overline{V}\}$. From the proof of Proposition 1 (Step 3) we know that $V(R[C, x], H) = V(C, H) = \overline{V}$ for any x and $C \in \mathcal{C}(\overline{V})$. Since ρ is independent of $\overline{V}$, it follows that $\Pi(R[C, 1 - \rho], 1 - \rho)$ is decreasing in $\overline{V}$. Thus we conclude that $\Pi^*(C^*, \rho) > \Pi^*(C, \rho)$ for any $C \notin \mathcal{C}(0)$, and that $\Pi^*(C^*, \rho) = \Pi^*(C, \rho)$ for any $C \in \mathcal{C}(0)$.

Case 2. $\rho \geqslant 1 - x[C]$.

First, notice that there cannot exist any equilibrium in which, following any contract C, it is the case that $\mu[C, \rho] > x[C]$. Otherwise, P_A would offer $R[C, \mu[C, \rho]]$. From the proof of Proposition 1 (Step 2), we know that $x[C]$ coincides with the strategy x that satisfies $V(R[C, x], L) = V(R[C, x], H)$, and that, since $\mu[C, \rho] > x[C]$, it is the case that $V(R[C, \mu[C, \rho]], L) > V(R[C, \mu[C, \rho]], H)$. Thus A_A would be better off playing L, and so it would be the case that $\mu[C, \rho] \leqslant x[C]$.

Secondly, since we are restricting attention to equilibria in which the play in the subform following any contract C is a Pareto-undominated equilibrium, we select the equilibria such that $\mu[C, \rho] = x[C]$. We thus conclude that $\Pi^*(C, \rho) = \Pi^*(C, 1)$ for any $\rho \in [1 - x(C), 1]$.

For the case when $\rho = 1 - x(C)$, we have shown that $\sigma_H[C, \rho] = 0$, and thus $\Pi^*(C, \rho) = \Pi(R[C, 1 - \rho], 1 - \rho)$. As in Case 1 it follows that $\Pi^*(C^*, \rho) > \Pi^*(C, \rho)$ if $C \notin \mathcal{C}(0)$, and that $\Pi^*(C^*, \rho) = \Pi^*(C, \rho)$ if $C \in \mathcal{C}(0)$.

Wrapping together Cases 1 and 2 allows us to determine the optimal initial contract.

Suppose that $\overline{C} \neq C^*$. Since $x[\overline{C}] > x[C^*]$, there is a unique threshold ρ' such that

$$\Pi\big(R[C^*, 1 - \rho'], 1 - \rho'\big) = \Pi^*(\overline{C}, 1),$$

and $\rho' \in (1 - x[\overline{C}], 1 - x[C^*])$. The principal of type P_A chooses $\overline{C}$ if $\rho > \rho'$, and any contract $C \in \mathcal{C}(0)$ if $\rho < \rho'$, and is indifferent between these two alternatives for $\rho = \rho'$.

If $\overline{C} = C^*$, instead, P_A initially offers contract C^* if $1 - \rho < x[C^*]$, and any contract $C \in \mathcal{C}(0)$ if $1 - \rho \geqslant x[C^*]$.

Step 2. The aware principal's choice between a contract C, and the monitoring option M^*.

By construction, it is the case that

$$\Pi'(M^*, H) > \Pi^*[\overline{C}, 1] = \Pi^*[\overline{C}, \rho] \quad \text{for any } \rho \in \big[1 - x[\overline{C}], 1\big],$$

and (hence) that $\Pi'(M^*, H) > \Pi^*[C^*, \rho]$ for any $\rho \in [1 - x[C^*], 1]$.

For $\rho = 0$, Step 1 shows that for any $C \in \mathcal{C}(0)$, it is the case that

$$\Pi^*[C, \rho] = \Pi^*[C^*, \rho] > \Pi(C^*, H) > \Pi'(M^*, H),$$

where the last inequality is by construction.

Since $\Pi^*[C^*, \rho]$ is strictly decreasing in ρ on the interval $[0, 1 - x[C^*]]$, it follows that there exists a unique value $\bar{\rho}$ satisfying

$$\Pi^*[C^*, \bar{\rho}] = \Pi'(M^*, H),$$

and that $\bar{\rho} \in [0, 1 - x[C^*]]$. When $\rho > \bar{\rho}$, the principal of type P_A chooses M^*, when $\rho < \bar{\rho}$ she initially offers a contract $C \in \mathcal{C}(0)$ and then she renegotiates it with $R[C^*, 1 - \rho]$, and finally she is indifferent between these two courses of action when $\rho = \bar{\rho}$.

The expression for $\bar{\rho}$ and R^* in the statement of the lemma follows from the definitions of $\Pi^*[C^*, \bar{\rho}]$, and $R[C^*, 1 - \bar{\rho}]$. □

References

Alos-Ferrer, C., 1999. Dynamical systems with a continuum of randomly matched agents. J. Econ. Theory 86, 245–267.
Boylan, R., 1992. Laws of large numbers for dynamical systems with randomly matched individuals. J. Econ. Theory 57, 473–504.
Border, K., Sobel, J., 1987. Samurai accountant: a theory of auditing and plunder. Rev. Econ. Stud. 54, 525–540.
Fudenberg, D., Levine, D.K., 1998. The Theory of Learning in Games. MIT Press, Cambridge, MA.
Fudenberg, D., Tirole, J., 1990. Moral hazard and renegotiation in contracts. Econometrica 58, 1279–1319.
Hale, J., 1969. Ordinary Differential Equations. Wiley, New York.
Harsanyi, J., 1967. Games with incomplete information played by 'Bayesian' players. Parts I–III. Manage. Sci. 14, 159–182, 320–334, 486–502.
Huberman, G., Kahn, C., 1988. Limited contract enforcement and strategic renegotiation. Amer. Econ. Rev. 78, 471–484.
Kandori, M., 1992. Social norms and community enforcement. Rev. Econ. Stud. 59, 63–80.
Kreps, D., Wilson, R., 1982. Reputation and imperfect information. J. Econ. Theory 27, 253–279.
Milgrom, P., Roberts, J., 1982. Limit pricing and entry under incomplete information. Econometrica 53, 889–904.
Mirrlees, J., 1976. The optimal structure of incentives and authority within an organization. Bell J. Econ. 7, 105–131.
Nachbar, J., 1990. Evolutionary selection dynamics in games: convergence and limit properties. Int. J. Game Theory 19, 59–89.
Weibull, J.W., 1995. Evolutionary Game Theory. MIT Press, Cambridge, MA.