Economics Bulletin

# Volume 36, Issue 3

# Short-run immiseration in repeated moral hazard

Pedro Hemsley
*FCE/UERJ*

## Abstract

In standard models of infinitely-repeated moral hazard with risk aversion, the agent´s continuation value from the relationship with the principal decreases non-monotonically over time and diverges to minus infinity almost surely, meaning that the agent gets most of the compensation in the early stages of the relationship and is left with little to receive in the future. In the short run, however, this pattern is less clear as continuation values may drift up as a response to good outcomes. I provide conditions under which the drift upwards is small and the probability of increased continuation values is very low even for short periods. Hence the immiseration result, an asymptotic property, is a good approximation in the short run.

**Citation:** Pedro Hemsley, (2016) "Short-run immiseration in repeated moral hazard", *Economics Bulletin*, Volume 36, Issue 3, pages 1479-1485
**Contact:** Pedro Hemsley - pedrohemsley@gmail.com.
**Submitted:** February 03, 2016. **Published:** August 03, 2016.

## 1. Introduction

Repeated moral hazard models have been analyzed formally since the pioneering works of Lambert [1983] and Rogerson [1985], the first to establish a major result in the literature: it is generally optimal to make the contract of a risk-averse agent history-dependent by tailoring current payments to past performances - that is, the optimal contract displays *memory*. Intuition is straight-forward: by doing so, the principal spreads the risk of a one-period shock over the whole relationship, and thus provides a natural source of insurance to the agent without decreasing his incentive to exert effort. When the relationship is infinitely-repeated[1], history-dependence has a stark implication: the agent's continuation value decreases on average over time and diverges to minus infinity with probability one. Intuitively, the principal drives the agent's continuation value towards a region where it is cheaper to offer incentives: exactly when the agent has low average consumption, and hence has a high marginal utility from consumption, since utility is concave.

This extreme asymptotic property is known as the immiseration result[2]. Although it is a powerful long-term property, the short-term behavior of continuation values may be different: Spear and Srivastava [1987] shows that if the observable outcome is good enough, then the principal raises the agent's continuation value. Specifically, there is a cutoff value such that for every outcome above it, the continuation value increases. Although continuation values decrease on average, they may increase in any given period - hence, they may drift upwards in the short-run.

This short-term pattern may vary in different settings. I provide conditions under which the upward drift is small and the probability of observing an increased continuation value is negligible even over short periods. Specifically, I consider a dynamic moral hazard model with two possible outcomes (success or failure) and two effort levels (high or low). I add two assumptions to this framework: (i) the probability of success is equal to zero if effort is low; and (ii) the principal's payoff is zero in case of failure. There is no baseline output and no chance of success in the absence of effort.

I show that under these assumptions, the short-term behavior of the optimal contract nearly repeats the long-term pattern if the agent is sufficiently risk-averse, which is typically the case in principal-agent relationships as the principal (e.g., a pharmaceutical company investing in a new drug) has a diversified capital structure and access to credit and insurance markets, while the agent (e.g., a researcher) has a more limited ability to cope with risk. Hence the immiseration result, an asymptotic result, is a good approximation for relatively short periods. Intuitively, the pay is front-loaded even in the short-run: most of the compensation to the agent is given in the beginning of the relationship as the cost of postponing rewards is too high. Continuation values only increase in any given sequence of periods if the agent is successful every time; in all other contingencies, it decreases.

The rest of the paper is organized as follows. Section 2 describes the model. Section 3 presents the results, and section 4 briefly concludes.

## 2. Model

Consider a simple version of the infinitely-repeated moral hazard model (see Laffont and Martimort [2002], chapter 8). A risk-neutral principal wants to hire a risk-averse agent to perform unobservedly a task $e$. The agent has a separable utility function $U(c,e) = u(c) - \psi(e)$, in which $c$ is consumption and $e$ is effort such that $u(0) = 0$, $u' > 0$ and $u" < 0$. In every period the agent chooses privately either high effort ($e_h$), with disutility $\psi(e_h) = K > 0$, or low effort ($e_l$), with no disutility: $\psi(e_l) = 0$. I assume that $K$ is not as high as to render an equilibrium with strictly positive effort infeasible. There are two possible outcomes: success, with a high output $Y_s$, or failure, with low output $Y_f$ such that $Y_s > Y_f$. Let $p_h = Prob(Y_s/e_h)$ and $p_l = Prob(Y_s/e_l)$, $p_h > p_l$. For tractability, in proposition 2 (and related lemmas) I will use $u(x) = \frac{1}{r} ln(1 + rx)$, in which $r \geq 0$ captures the level of risk aversion.

Both the principal and the agent discount the future with the same rate $\beta$. The principal is subject to an incentive compatibility constraint if he wants to induce high effort. If they do not sign a contract, the agent gets a lifetime utility denoted by $u_{res}$ and the principal makes zero profit in a competitive market.

[1] See Laffont and Martimort [2002], chapter 8.

[2] It also arises in infinitely-repeated adverse selection models: see Thomas and Worrall [1990].

This is formally a repeated game with imperfect public monitoring. A recursive structure applies to this class of games, as shown in Abreu et al. [1990], and there is a stationary value function $V(.)$ that represents the principal's expected payoff from the relationship. A recursive contract is characterized by the corresponding policy functions: a present payment and a continuation value following each possible realization of output, plus a choice of effort: $\{(c_f, w_f), (c_s, w_s), e\}$. History is summarized in the state variable $w$, which is the agent's beginning-of-period promised value from the relationship. When $w$ is constant over time, the contract is said to be stationary, or to present *amnesia*, as it disposes of history; hence studying the time pattern of the contract amounts to establishing how $w$ evolves over time. Assume the principal achieves strictly positive profits under high effort for every history $w$. The initial value $w_0 \geq u_{res} > 0$ is given.

In order to induce high effort after history $w$, the principal solves the following recursive program:

$$V(w/e_h) = \underset{c_s, c_f, w_s, w_f}{Max} \quad p_h [Y_s - c_s + \beta V(w_s)] + (1 - p_h) [Y_f - c_f + \beta V(w_f)]$$

subject to the following constraints:

$$(IC) \quad p_h [u(c_s) + \beta w_s] + (1 - p_h) [u(c_f) + \beta w_f] - K \geq p_l [u(c_s) + \beta w_s] + (1 - p_l) [u(c_f) + \beta w_f]$$

$$(PK) \quad p_h [u(c_s) + \beta w_s] + (1 - p_h) [u(c_f) + \beta w_f] - K - w = 0$$

The first restriction is the usual incentive compatibility constraint (IC) for the agent: his expected utility when making effort must be higher than if he shirks. The second is the promise-keeping restriction that guarantees the principal will deliver the promised value $w$. Defining $p^R \equiv \frac{p_h}{1-p_h}$, this problem may be simplified as follows.

$$V(w/e_h) = \underset{c_s, c_f, w_s, w_f}{Max} \quad p_h [Y_s - c_s + \beta V(w_s)] + (1 - p_h) [Y_f - c_f + \beta V(w_f)]$$

subject to the following constraints:

$$(IC) \quad [u(c_s) + \beta w_s] - [u(c_f) + \beta w_f] - \frac{K}{p_h - p_l} \geq 0$$

$$(PK') \quad p^R [u(c_s) + \beta w_s] + [u(c_f) + \beta w_f] - \frac{K}{1 - p_h} - \frac{w}{1 - p_h} = 0$$

I follow Laffont and Martimort [2002] (chapter 8) and substitute (PK) for (PK'), which is defined with a weak inequality. Strict inequality cannot hold in equilibrium, as it would affect the incentive compatibility constraint in the previous period, compromising the recursive structure of the problem; hence it must be checked whether it is binding in equilibrium, which is the case if the multiplier of (PK') is strictly positive.

The technology available to the agent has the following property.

**Definition 1.** Effort is *essential*:

i)$Y_f = 0$
ii) $p_l = 0$

In words, this assumption states that output is equal to zero in the absence of effort: the probability of success is zero, and there is no output to split in this case. This may be interpreted as a depiction of any principal-agent relationship in which effort is essential: there is no baseline output that can be improved upon if the agent works harder. Notice that the usual Monotone Likelihood Ratio Condition is trivially satisfied in this setup: $\frac{p_h}{1-p_h} > \frac{p_l}{1-p_l} = 0$. Lastly, I will assume for the sake of simplicity that $p_h = 1 - p_h$, so that $p_h = \frac{1}{2}$. This acts as a lower bound on the probability of success in a given period; hence I preclude the possibility that decreased continuation values in the short-run might stem from artificially low probabilities of success. As such, it allows one to trace the behavior of the optimal contract to other primitives of the model, as established in the next section.

## 3. Optimal Contract

Assume that both the principal and the agent are bound to the contract they sign. In this case, there is no need for a participation constraint for either of them, since it is assumed that $w_0 \geq u_{res}$ and the principal has positive profit after every history. In order to show that the program above effectively characterizes the problem posed, low effort must be ruled away. Notice that choosing $e = 0$ in a given period is the same as destroying output. This hurts the principal, but also hurts the agent and in theory could be used to relax the incentive compatibility constraint: the principal trades this benefit off against the forgone output. In the present model, however, it is optimal for the principal to induce high effort ($e = e_h$) in every period due to the technology available to the agent: the cost in terms of output is too high.

The principal induces high effort in every period for two reasons. First, inducing low effort is a stationary solution, meaning that he will choose low effort again in every period afterward. Second, low effort gives him strictly negative profits, which cannot be optimal as it was assumed that he makes strictly positive profit under high effort following any history - in particular, high effort yields positive profit in the optimal static (or stationary) contract. This captures what "essential effort" means: the level of effort induced by a constant payment gives the principal a strictly lower payoff than he can get in the optimal static contract[3].

To see that the value function is strictly concave and differentiable, substitute $h \equiv u^{-1}$ into the problem and notice that the constraint set defined by (IC) and (PK') is strictly convex since $h$ is convex[4].

It follows that the solution to the program is fully characterized by the Kuhn-Tucker conditions plus the envelope theorem (I will consider an interior solution, which is the case for example if $u'(0) = \infty$). Let $\gamma$ and $\mu$ be the multipliers of IC and PK, respectively. First-order conditions are:

FOC($c_h$):$-p_h + \gamma u'(c_s) + \mu p^R u'(c_s) = 0$
FOC($c_l$):$-(1 - p_h) - \gamma u'(c_f) + \mu u'(c_f) = 0$
FOC($w_h$): $p_h V'(w_s) + \gamma + \mu p^R = 0$
FOC($w_l$): $(1 - p_h) V'(w_f) - \gamma + \mu = 0$

The envelope theorem states that $V'(w) = \frac{-\mu}{1-p_h}$.

Moreover, the argument in Spear and Srivastava [1987] applies without change: $V(w)$ is the Pareto frontier, and hence must be non-increasing - otherwise the principal could simply offer the agent more and be better off. In fact, it is strictly decreasing: from FOC($c_l$), one has $\mu = \frac{1-p_h}{u'(c_f)} + \gamma$, which is strictly positive due to $u' > 0$. Since this result does not depend on $w$, it follows that (PK') holds with equality after every history.

The optimal contract is front-loaded: the agent's average payment is higher in the beginning of the relationship and decreases as it evolves. In other words, the longer the time horizon is, the lower the expected payment per period becomes. Intuitively, the principal wants to provide incentives at the lowest cost. To do so, he will spread shocks over time - i.e., make the contract history-dependent, so that the state variable $w$ should change over time. As shown in Spear and Srivastava [1987], the principal rewards a success both in the present ($c_s > c_f$) and in the future ($w_s > w > w_f$) in such a way that the state variable decreases on average: when consumption goes down, the marginal utility of consumption increases, allowing the principal to achieve a given spread in payoffs after success and failure with a smaller difference

---

[3] To see that it is indeed optimal to induce $e = e_h$ in every period, assume the principal wants to induce low effort when the promised value is $\tilde{w}$. Then the probability of success is zero and he must choose only one present payment $c$ and one continuation value $w'$. The value function becomes $V(\tilde{w}) = \underset{c,w'}{Max} \ [-c + \beta V(w')]$ subject to $(PK) \quad [u(c) + \beta w'] - \tilde{w} \geq 0$. In this case the weak inequality poses no problem since there is no (IC). This program is strictly concave and differentiable. Let $\mu$ be the multiplier of the promise-keeping constraint; first-order conditions with respect to $c$ and $w$ are:$-1 + \mu u'(c) = 0$ and $V'(w') + \mu = 0$. The envelope theorem implies that $V'(\tilde{w}) = -\mu$.

The first line implies that $\mu > 0$ since $u' > 0$. Then $V'(w') = V'(\tilde{w})$. Since $V$ is strictly concave, it follows that $w' = \tilde{w}$, so that the principal will be facing the same problem in the next period and will choose low effort again, so that it is optimal to choose low effort in every period afterward. But then his profit will be less than zero as he must give the agent strictly positive consumption levels (due to $u_{res} > 0$ and the fact that $u(0) = 0$). But this cannot be optimal as it was assumed the principal makes a strictly positive profit under high effort after any history - in particular, he would be strictly better off by offering the optimal static contract.

[4] For details, see Laffont and Martimort [2002] (chapter 8) and Stokey et al. [1989].

in consumption. Notice that under such a front-loaded contract, the agent would like to save part of his gains; if allowed to do so, this would affect incentives in the second period. I assume the the agent consumes immediately the transfer he receives from the principal[5].

The following proposition describes the features of the optimal contract in the present context.

**Proposition 1.** *Spear and Srivastava [1987]: Continuation values to the agent decrease on average and diverge to minus infinity.*

*Proof.* FOC($c_s$) and FOC($w_s$) imply that $V'(w_s)\,u'(c_s) = -1$, while FOC($c_f$) and FOC($w_f$) imply $V'(w_f)\,u'(c_f) = -1$. Then:
$\frac{u'(c_f)}{u'(c_s)} = \frac{V'(w_s)}{V'(w_f)}$.

If $w_f \geq w_s$, it follows that $c_f \geq c_s$ (since $V$ is strictly decreasing and strictly concave), which is not incentive compatible - hence one must have $w_s > w_f$.

In order to see that $w_s > w > w_f$, solve for $\gamma$ in FOC($w_f$) and FOC($w_s$) to get $(1-p_h)\,V'(w_f) + \mu = -p_h V'(w_s) - \mu p^R$. Rearranging, one has:

$$(1-p_h)\,V'(w_f) + p_h V'(w_s) = -\mu p^R - \mu$$

The left-hand side is simply the expected value of the derivative of the value function $V$; label it $E[V'(\tilde{w})]$, in which $\tilde{w}$ is the next-period continuation value (this is a random variable as it depends on the realized output). Using the envelope theorem ($V'(w) = \frac{-\mu}{1-p_h}$), the previous line becomes:

$$E[V'(\tilde{w})] = V'(w)$$

This is simply the martingale property for $V'$ first identified in Spear and Srivastava [1987]. Together with $w_s > w_f$, it implies that $V'(w_s) < V'(w) < V'(w_f)$. The concavity of $V$ then implies $w_s > w > w_f$.

Lastly, $V' < 0$ implies that $V'$ is a non-positive martingale, which converges almost surely. Since $w_f$ is unbounded from below, the same argument made in Ljungqvist and Sargent (2004, p. 670) applies: if $V'$ converges to a finite limit, continuation values also converge - the principal does not spread continuation values even when the cost of doing so is very low. Hence $w$ converges almost surely to minus infinity. □

Let $w_N$ denote the continuation value $w$ after $N$ periods. Proposition 1 implies that $w_N$ goes to minus infinity as $N \to \infty$ with probability one, but for $N = 1$ there is a strictly positive probability that it will increase. This characterization is silent about the short-term behavior of continuation values, i.e., $1 < N < \infty$. To establish the result, I will consider initially the value of $N$ for which the probability of observing an increase is highest: $N = 2$.

For the next results, I consider the logarithmic utility presented at the beginning of the previous section. Under this specification, the inverse of the agent's marginal utility of consumption at $x$ is equal to $1 + rx$, which is linear in $x$. Notice that the martingale property $E[V'(\tilde{w})] = V'(w)$ identified above, which follows form the first-order conditions and is cast in terms of continuation values, is a generalization of a basic result found in finite-horizon repeated moral hazard problems: namely, the agent's inverse marginal utility of consumption is also a martingale, which reads $1 + rx = E[1 + r\tilde{x}]$ in the present setting for some level of current consumption $x$ and a random level of consumption in the following period $\tilde{x}$. This reduces to $x = E[\tilde{x}]$, implying that consumption is a martingale. The concavity of $u$ then implies that $u(x)$ is a super-martingale[6], and decreases on average towards minus infinity, while consumption converges to its lower bound $\frac{-1}{r}$. Also notice that, under this specification for the utility function, $r$ is positively correlated with the agent's absolute risk aversion[7]. In what follows, I will interpret a higher $r$ as a higher level of risk aversion[8].

[5] In other words, the agent has no access to financial markets or to any type of saving technology.

[6] See Lemma 5.1 in Varadhan [1999].

[7] When the level of consumption is zero, the agent's absolute risk aversion is equal to $r$: $\frac{-u"(0)}{u'(0)} = r$

[8] The results based on the logarithmic utility function are not specific to this note. For a detailed discussion, including a complete derivation of time patterns for consumption and utility, see Laffont and Martimort [2002] (chapter 8, pages 336-341).

The next lemma computes the expected value of the agent's continuation value after two periods and its drifts after a success or a failure. This computation is based on an explicit construction of the principal's value function, which is done through a guess-and-verify approach that starts with a function of the form $V(w) = \alpha - \frac{1}{\beta} exp(\beta w + \gamma)$ for some real parameters $\alpha$, $\beta$ and $\gamma$ that may be found using the first-order conditions[9].

**Lemma 1.** *The optimal contract has the following features.*

(1) The upward jump in the agent's continuation value in case of success is $K - \frac{1}{r(1-\beta)} ln\left(\frac{e^X + e^{-X}}{2}\right)$.

(2) The downward jump in the agent's continuation value in case of failure is $-K - \frac{1}{r(1-\beta)} ln\left(\frac{e^X + e^{-X}}{2}\right)$.

(3) The expected value of $w$ after two periods is $w_2^E = w - \frac{\beta(1+\beta)}{2r(1-\beta)} ln\left(\frac{e^X + e^{-X}}{2}\right)$.

in which $X = r(1-\beta)K$

*Proof.* Straightforward application of Laffont and Martimort [2002] (chapter 8, section 8.2.6). □

One may see immediately that $w_2^E < w$, as would be expected from Proposition 1.

The additional structure in the current setup implies that the continuation value decreases in all but one contingency provided that the agent is sufficiently risk averse. Namely, $w$ only increases after a sequence of two successes; if there is at least one failure, either in the first or in the second period, $w$ decreases.

**Lemma 2.** *There is a cutoff value $\overline{r}$ such that for all $r > \overline{r}$, the continuation value $w_2$ is higher than $w$ if and only if output is high in both periods.*

*Proof.* The "if" part is a consequence of Proposition 1. To prove the "else if" part, start by defining $F(r) \doteq \beta\left(K - \frac{1}{r(1-\beta)} ln\left(\frac{e^X + e^{-X}}{2}\right)\right) + \beta^2\left(-K - \frac{1}{r(1-\beta)} ln\left(\frac{e^X + e^{-X}}{2}\right)\right)$: according to Lemma 1, this is the agent's future gain, in terms of continuation values for two periods, that comes from a current success followed by a failure. This is higher than the gain when success and failure are reversed because $\beta < 1$; hence it is sufficient to show that $F(r) < 0$, which in turn implies that $\beta\left(-K - \frac{1}{r(1-\beta)} ln\left(\frac{e^X + e^{-X}}{2}\right)\right) + \beta^2\left(K - \frac{1}{r(1-\beta)} ln\left(\frac{e^X + e^{-X}}{2}\right)\right) < 0$.

Define also $f(r) \doteq \frac{1}{r(1-\beta)} ln\left(\frac{e^X + e^{-X}}{2}\right)$; hence $F(r) = \beta[K - f(r)] + \beta^2[-K - f(r)]$. L'Hospital's rule implies $\lim_{r\to 0} f(r) = 0$ and $\lim_{r\to\infty} f(r) = K$. It follows that $\lim_{r\to 0} F(r) = \beta K - \beta^2 K = K\beta(1-\beta) > 0$ and $\lim_{r\to\infty} F(r) = -\beta^2 K < 0$. Since $F(r)$ is continuous, one may apply the intermediate value theorem to establish that there is a threshold such that $F(r) = 0$. An analogous argument shows that $f'$ becomes strictly positive and bounded away from zero for $r$ high enough; hence, $F'$ becomes strictly negative and bounded away from zero.

This implies the highest value $r$ such that $F(r) = 0$ is finite and $F$ does not become positive above it. If there is only one value of $r$ such that $F(r) = 0$, take this as $\overline{r}$. If there are multiple values of $r$ that solve $F(r) = 0$, then take the highest one as $\overline{r}$, concluding the argument. □

It follows that for any sequence of two periods, $w$ will decrease except in one only case, which limits severely the probability that it will increase in the short-run in spite of proposition 1. This is summarized in proposition 2.

**Proposition 2.** *If $r$ is high enough, then in every sequence of two periods, the continuation value $w$ only increases if the agent is successful in both of them, and decreases in all other contingencies.*

In the current setup, the probability that this will occur is $p_h^2 = \frac{1}{4}$. Even if this is non-negligible, the expected value $w_2$ is still significantly lower than $w$ due to the restriction identified in proposition 2. This argument may be indefinitely extended for $N > 2$ as long as one chooses $r$ accordingly.

[9] Again, the reader may refer to Laffont and Martimort [2002], pages 336-339, for a detailed derivation of these parameters.

**Corollary 1.** *If $r$ is high enough, then in every sequence of $N$ periods, the continuation value $w$ increases if and only if the agent is successful in all periods, and decreases in all other contingencies.*

The probability that this will happen is $p^N$, which converges to zero as $N$ increases. In the current setting, $p_h = 0.5$ and $p^N$ becomes 3% after 5 periods and only 0.1% after 10 periods. Intuitively, the optimal contract is very front-loaded: the principal rewards the agent for a success at the very beginning of the relationship and then drives him quickly into a region with low utility and, therefore, high marginal utility, decreasing the cost of providing incentive to work hard. When effort is essential and the agent's risk aversion is high, the optimal contract is more tilted towards current rewards. For practical terms, the immiseration result is a good approximation even in the short-run in this setting.

To grasp the intuition behind Proposition 2 and Corollary 1, consider as a benchmark the equilibrium when the agent is risk-neutral. In this case, the optimal contract may be made stationary without loss of generality, and the continuation value is constant and equal to $w_0$: it is sufficient for the principal to use current payments to provide rewards and punishments[10]. In other words, current payments and continuation values are perfect substitutes at a rate determined by the discount factor $\beta$. This happens because there is no tradeoff between incentives and insurance.

As the agent's level of risk aversion increases, this tradeoff develops, and it is optimal for the principal to offer a history-dependent contract so as to make the provision of incentives cheaper: the repetition of the game allows the principal to spread the agent's rewards and punishments over time, so that the agent bears only a fraction of the risk associated with his effort in any given period. As discussed above, this implies that continuation values decrease on average and drift towards minus infinity, so that the contract becomes front-loaded.

This has an implication for the design of the optimal contract even in the short run, i.e., for a number of periods $N < \infty$. For any given range of $N$ periods, the expected value of the continuation value decreases when the agent becomes more risk averse[11]. Since it is optimal to implement high effort in every period, the probabilities of success and failure are exogenous when it comes to the decision about current and future payments, implying that the principal needs to decrease the expected continuation value by actually reducing this value in at least some contingencies. In principle, he might achieve that with a reduction for some (or all) contingencies without affecting the sign of $E(w_N - w_0)$ after any sequence successes and failures: that is, the principal might implement a lower expected continuation value without affecting the direction of the drift (whether upward or downward) in any contingency.

If this were the case, then the immiseration result would not be a reasonable approximation in the short-run. The property identified in Proposition 2, however, shows that the principal does not limit himself to such changes: he also alters the direction of the drift for some sequences, and does so to maximum extent when the agent's level of risk aversion is high enough: then the continuation value decreases for all but one sequence of outcomes - namely, if the agents hits only successes in a sequence of $N$ periods. As a result, the immiseration result, an asymptotic property, becomes a good approximation for the short-run behavior of the optimal contract when the agent is very risk-averse.

## 4. Final Remarks

This paper provides conditions under which the long-term property of immiseration nearly holds in the short-run. In practical terms, this means that contracts are very front-loaded: the reward for success is provided as early as possible by the principal. This happens whenever the agent is sufficiently risk averse, the probability of success is equal to zero in the absence of effort, and output in case of failure is equal to zero.

Future research should aim at establishing the robustness of the results in this note. Specifically, one should consider more general utility functions and technologies to establish an upper bound on the probability of observing an increased continuation value in the short-run.

[10] For a detailed discussion of stationarity under risk neutrality in the context of relational contracts, see Levin [2003].

[11] Otherwise, the expected continuation value would have to increase for periods $n > N$, but this cannot be the case as it would violate Proposition 1.

## References


Abreu, D., Pearce, D., and Stacchetti, E. (1990) "Toward a theory of discounted repeated games with imperfect monitoring" Econometrica, 58:1041-1063.

Laffont, J.J. and Martimort, D. (2002) *Theory of Incentives: The Principal-Agent Model* Princeton University Press.

Lambert, R. (1983) "Long-term contracts and moral hazard" Bell Journal of Economics, 14:441-452.

Levin, J. (2003) "Relational incentive contracts" The American Economic Review, 93:835-857.

Rogerson, W. (1985) "Repeated moral hazard" Econometrica, 53:69-76.

Spear, S. and Srivastava, S. (1987) "On repeated moral hazard with discounting" The Review of Economic Studies, 54:599-617.

Stokey, N., Lucas, R. and Prescott, E. (1989) *Recursive Methods in Economic Dynamics* Harvard University Press.

Thomas, J. and Worrall, T. (1990) "Income Fluctuation and asymmetric information: An example of a repeated principal-agent problem" Journal of Economic Theory, 51:367-390.

Varadhan, S.R.S. (1999). *Probability theory* Courant Institute for Mathematical Sciences, New York University.