Working Paper Series

Repeated Moral Hazard with Effort Persistence

WP 08-04

Arantxa Jarque
Federal Reserve Bank of Richmond

THE FEDERAL RESERVE BANK OF RICHMOND
RICHMOND ▪ BALTIMORE ▪ CHARLOTTE

Repeated Moral Hazard with Effort Persistence*

Arantxa Jarque†
FRB– Richmond
and U. Carlos III de Madrid

July 2008
Federal Reserve Bank of Richmond Working Paper 08-04

Abstract

I study a problem of repeated moral hazard in which the effect of effort is persistent over time: each period's outcome distribution is a function of a geometrically distributed lag of past efforts. I show that when the utility of the agent is linear in effort, a simple rearrangement of terms in his lifetime utility translates this problem into a related standard repeated moral hazard. The solutions for consumption in the two problems are observationally equivalent, implying that the main properties of the optimal contract remain unchanged with persistence. To illustrate, I present the computed solution of an example.

Journal of Economic Literature Classification Numbers: D30, D31, D80, D82.

Key Words: mechanism design; repeated agency.

*I would like to thank Harold Cole, Huberto Ennis, Juan Carlos Hatchondo, Nicola Pavoni and especially Hugo Hopenhayn and Iván Werning for very useful suggestions and discussions. I would also like to thank an anonymous referee for feedback on an earlier version of the article. The paper benefited from comments by seminar audiences at CEMFI, the University of Iowa, the Federal Reserve Bank of Richmond, and the Meetings of the SED 2005. I thank the Ministry of Education, Culture and Sports of Spain for its support through the Research Project SEJ2004 - 08011ECON. This article is based on part of my dissertation; I gratefully acknowledge financial support from the Banco de España during my Ph.D.

†Email: Arantxa.Jarque@rich.frb.org.

1 Introduction

The moral hazard literature has pointed out repeatedly the importance of generalizing the current models of asymmetric information to setups in which either the hidden endowment of the agent or the effect of the agent's effort are correlated in time, or "persistent." The difficulty of this generalization lies on the fact that, with persistence, incentives for deviation in a given period depend on the history of private information of the agent (what is sometimes referred to as the "joint deviations" problem). As with the case in which the agent has access to hidden savings (Abraham and Pavoni, 2006, and Werning, 2001), persistence implies that there is no common knowledge of preferences at the beginning of each period. Hence, the standard recursive formulation with continuation utility as a state variable (Spear and Srivastava, 1987) is not valid in the presence of persistence.

In this paper, I characterize a class of agency problems of repeated moral hazard (hidden effort) with persistence for which a simple solution exists. I model persistence by assuming that the effort of the agent affects not only the current output distribution but also the distribution of output in every future period. The class of problems that I characterize satisfies two key assumptions: the agent is risk averse with linear disutility of effort, and output depends on the sum of depreciated past efforts. I show that the (constrained) optimal contract can be found by solving an auxiliary problem – a related repeated moral hazard problem without persistence. The two problems are observationally equivalent. Hence, the intertemporal properties of consumption under persistence do not differ from those of the optimal contract in problems without persistence. The (unobservable) effort sequences do typically exhibit different properties when effort is persistent. I present a numerical example that illustrates the main conclusions.

1.1 Related Literature

The existing literature on repeated moral hazard with persistence has not yet generally characterized the effect of persistence on the properties of effort and consumption in the optimal contract. There are, however, some interesting examples for which a solution can be found. Fernandes and Phelan (2000) provide the first recursive treatment of agency problems with effort persistence. In their paper, the current effort of the agent affects output in the same period and in the following one. Their setup is characterized by three parameters: the number of periods for which the effect of effort lasts, the number of possible effort levels and the number of possible outcome realizations. All three parameters are set to two, and this makes their formulation and their computational approach feasible. The curse of dimensionality applies whenever any of the three parameters is increased. Moreover, no results are given in their paper on how the properties of the optimal contract differ from the case without persistence. In my formulation, I allow for a continuum of efforts, infinite persistence and multiple outcomes, but under my assumptions the recursive formulation of the problem turns out to be particularly simple. Mukoyama and Sahin (2005) and Kwon (2006) show in a discrete two–effort model that, if persistence is high, it may be optimal for the principal to perfectly insure the agent in the initial periods. The results that I present here apply to a different class of problems in which past efforts influence output less than current effort; in this context, perfect insurance is never optimal.

2 Model

I study a T–period relationship between a risk neutral principal and a risk averse agent, where T may be infinite. I consider the repeated moral hazard (RMH) problem that arises because the effort of the agent is unobservable. I generalize the standard RMH problem by considering the case of "persistent" effort : effort carried out by the agent each period affects current as well as future output distributions.

I assume that both parties commit to staying in the contract and that the principal has perfect control over the savings of the agent. They both discount the future at a rate β. I assume that the agent has additively separable utility that is linear in effort.

A1 The agent's utility is given by $U(c_t, e_t) = u(c_t) - e_t$, where u is twice continuously differentiable and strictly concave and c_t and e_t denote consumption and effort at time t, respectively.

There is a finite set of possible outcomes each period, $Y = \{y_i\}_{i=1}^N$, and the set of histories of outcome realizations up to time t is denoted by Y^t, with typical element $y^t = (y_1, ..., y_t)$. Histories of outcomes are assumed to be common knowledge. I assume both consumption and effort lie in a compact set: $c_t \in [0, y_t]$ and $e_t \in E = [0, e_{\max}]$ for all t.

To capture the persistence of effort, I model the probability distribution of current output as a function of the whole history of past efforts, and I denote it by $\pi\left(y_t|e^t\right)$, where e^t is the history of effort choices up to time t. I assume the distribution has full support: in every period t, $\pi\left(y_t|e^t\right) > 0$ for all $y_t \in Y$ and for all $e^t \in E^t$.

The strategy of the principal consists of a sequence of consumption transfers to the agent contingent on the history of outcomes, $\mathbf{c} = \left\{c\left(y^t\right)\right\}$, to which he commits when offering the contract at the beginning of time. The agent's strategy is a sequence of period best–response effort choices that maximize his expected utility from t onward, given the past history of output: $\mathbf{e} = \left\{e\left(y^t\right)\right\}$. At the end of each period, output y_t is realized according to the distribution determined by the effort choices up to time t, and the corresponding amount of consumption $c\left(y^t\right)$ is given to the agent.

An optimal contract is a pair of contingent sequences $\mathbf{c}^* = \left\{c^*\left(y^t\right)\right\}$ and $\mathbf{e}^* = \left\{e^*\left(y^t\right)\right\}$ that maximize the expected discounted difference between output and the promised contingent payments, subject to two constraints: the Participation Constraint (PC), which states that the initial expected utility of the agent in the contract should be at least as large as his outside utility, w_0 , and the Incentive Constraint (IC), which states that the sequence $\mathbf{e}^*$ should be a solution to the maximization problem faced by the agent, given the contingent consumption transfers established by $\mathbf{c}^*$:

$$\max_{\substack{\mathbf{e} \in E^T, \\ \mathbf{c}:\, c\left(y^t\right) \in [0, y_t]\ \forall t}} \sum_{t=1}^{T} \sum_{y^t} \beta^{t-1} \left[y_t - c\left(y^t\right)\right] \prod_{\tau=1}^{t} \pi\left(y_\tau \middle| \left\{e_j\left(y^{j-1}\right)\right\}_{j=1}^{\tau}\right)$$

$$s.t.$$

$$\sum_{t=1}^{\infty}\sum_{y^t}\beta^{t-1}\left[u\left(c\left(y^t\right)\right)-e_t\left(y^{t-1}\right)\right]\prod_{\tau=1}^{t}\pi\left(y_\tau\middle|\left\{e_j\left(y^{j-1}\right)\right\}_{j=1}^{\tau}\right)\geq w_0 \tag{PC}$$

$$\mathbf{e}\in\arg\max_{\widehat{\mathbf{e}}\in E^T}\sum_{t=1}^{T}\sum_{y^t}\beta^{t-1}\left[u\left(c\left(y^t\right)\right)-\widehat{e}_t\left(y^{t-1}\right)\right]\prod_{\tau=1}^{t}\pi\left(y_\tau\middle|\left\{e_j\left(y^{j-1}\right)\right\}_{j=1}^{\tau}\right) \tag{IC}$$

Following the influential work of Spear and Srivastava (1987), the usual procedure in the literature to solve for the optimal contract in the standard RMH problem without persistence is to write the problem of the principal recursively using the continuation utility of the agent as a state variable. In this recursive formulation, the IC is a per–period constraint on the level of effort. If the problem of the agent stated in the IC is concave in his effort choice, the corresponding first order condition is both necessary and sufficient. In such case, the IC can be replaced by this first order condition and a Lagrangian can be constructed for the problem of the principal.[1]

In general, it is not possible to follow a similar approach to find the solution to the problem with persistence . When effort is persistent, incentives for deviation may also depend on the particular sequence of past and future efforts chosen by the agent.[2] Therefore, one needs to check for the possibility of joint deviations involving effort choices in more than one period; this implies that the standard recursive formulation is no longer valid, complicating the computation of the optimal contract.[3]

In what follows, I characterize a class of problems where the specification of persistence is such that the optimal contract can be found in a simple way. Within this class, the problem above can be translated into a standard RMH problem, where the usual recursive tools can be used to derive the optimal contract. I make the following simplifying assumption about how the effect of effort persists in time:

A2 The distribution of output depends on a "productive state," denoted s, determined by the history of effort choices in the following way:

$$s_t=\sum_{\tau=1}^{t}\rho^{t-\tau}e_\tau, \tag{1}$$

were $\rho\in(0,1)$ measures the persistence of effort through its effect on future productive states, and $s_0=0$.

Under A2, s_t is a sufficient statistic for the history of effort up to time t, and the distribution of outcomes at time t can now be simply written as $\pi\left(y_t|s_t\right)$.[4] The limit case of $\rho=0$ corresponds

[1] This solution procedure is often called the First Order Approach. See Rogerson (1985b) and Jewitt (1988).

[2] For a given continuation contract, the first order condition of the IC problem with respect to e_t includes the term $\frac{\partial\pi\left(y_{t+j}|\{e_\tau\}_{\tau=t}^{t+j}\right)}{\partial e_t}$ for all j periods between t and T. Persistence implies that this derivative depends on $\{e_k\}_{k=1}^{t-1}$ and that it is non–zero for some or all future j's.

[3] A modified recursive formulation including three state variables may be possible. See Fernandes and Phelan (2000) for an example of this approach.

[4] Under A2 but without A1, any recursive formulation would in general need three state variables (Jarque, 2002). It is the combination of the two assumptions that further simplifies the problem.

to the standard RMH problem in which the probability distribution of current outcomes depends only on current effort. The set of feasible productive states at time t, S_t, is derived recursively from the set of feasible efforts and equation (1):

$$S_t = [\rho s_{t-1}, \rho s_{t-1} + e_{\max}] . \tag{2}$$

I can express the strategy of the agent using s as a function of the history of outputs, by substituting in (1) the corresponding $e\left(y^t\right)$ for each t. This allows me to write the expected utility of the agent in terms of the sequence $\mathbf{s} = \left\{s\left(y^t\right)\right\}_{t=1}^{T}$:

$$\sum_{t=1}^{T}\sum_{y^t}\beta^{t-1}\left[u\left(c\left(y^t\right)\right) - \left(s\left(y^{t-1}\right) - \rho s\left(y^{t-2}\right)\right)\right]\prod_{\tau=1}^{t}\pi\left(y_\tau | s\left(y^{\tau-1}\right)\right).$$

After a simple rearrangement of terms in the above expression, the problem of the principal can be written as follows:

$$\max_{\substack{\mathbf{s}\in S^T,\\ \mathbf{c}:\, c\left(y^t\right)\in[0,y_t]\ \forall t}} \sum_{t=1}^{T}\sum_{y^t}\beta^{t-1}\left[y_t - c\left(y^t\right)\right]\prod_{\tau=1}^{t}\pi\left(y_\tau | s\left(y^{\tau-1}\right)\right), \tag{\textbf{OP}}$$

$$s.t.$$

$$w_0 = \sum_{t=1}^{T}\sum_{y^t}\beta^{t-1}\left[u\left(c\left(y^t\right)\right) - (1-\beta\rho)\, s_t\left(y^{t-1}\right)\right]\prod_{\tau=1}^{t}\pi\left(y_\tau | s\left(y^{\tau-1}\right)\right)$$

$$\mathbf{s} \in \arg\max_{\widehat{\mathbf{s}}\in S^T}\sum_{t=1}^{T}\sum_{y^t}\beta^{t-1}\left[u\left(c\left(y^t\right)\right) - (1-\beta\rho)\, \widehat{s}_t\left(y^{t-1}\right)\right]\prod_{\tau=1}^{t}\pi\left(y_\tau | \widehat{s}\left(y^{\tau-1}\right)\right).$$

Simple inspection of this formulation of the problem shows that the solution for the optimal contract, $\mathbf{c}^*$, and the sequence of productive states that it implements, $\mathbf{s}^*$, are also a solution for the optimal consumption and *effort* sequences in an RMH problem without persistence in which the utility of the agent is given by

$$U\left(\widetilde{c}_t, \widetilde{e}_t\right) = u\left(\widetilde{c}_t\right) - (1-\beta\rho)\,\widetilde{e}_t,$$

and the distribution function is

$$\widetilde{\pi}\left(y_t|\widetilde{e}_t\right) = \pi\left(y_t|\widetilde{e}_t\right),$$

where $\widetilde{e}_t$ and $\widetilde{c}_t$ denote the choice variables in this new problem. I refer to this related moral hazard problem as the auxiliary problem (**AP**), with solution $\widetilde{\mathbf{c}}^*$ and $\widetilde{\mathbf{e}}^*$. The solution for effort in the original problem (**OP**) can be recovered from $\widetilde{\mathbf{e}}^*$ according to

$$\begin{aligned} e_1^* &= s_1^* = \tilde{e}_1^*, \\ e_t^* &= s_t^* - \rho s_{t-1}^* = \tilde{e}_t^* - \rho\tilde{e}_{t-1}^* \quad t>1. \end{aligned} \tag{3}$$

Some additional technical conditions are needed for the solutions of problems **OP** and **AP** to coincide:

A3 The expected utility of the agent for a given consumption scheme $\mathbf{c}$ is concave in each s_t.

A4 The distribution function $\pi(y_t|s_t)$ satisfies

$$\lim_{s\to 0}\frac{\partial E[y]}{\partial s}=+\infty \quad \text{and} \quad \lim_{s\to e_{\max}}\frac{\partial E[y]}{\partial s}=0.$$

These two assumptions ensure that, in problem **AP,** the solution for effort is always interior and the expected utility of the agent is a concave function of his effort choice. Hence, the IC can be substituted by the implied first order conditions and problem **AP** can be solved with the standard recursive techniques.[5]

In any RMH problem with linear disutility of effort, effort choices must lie in a closed set; this ensures that the domain of the problem is not empty by putting a limit on the utility that the agent can get by choosing a very low effort.[6] In problem **OP,** I assume a closed set $E=[0,e_{\max}]$, which translates into a closed S_t. Ideally, in problem **AP** we would like to impose $\widetilde{e}_t\in S_t$. However, the set S_t is endogenously determined according to Eq. (2), *given* the effort solution for problem **OP**. In practice, the auxiliary problem must be solved using an exogenously determined domain for effort, such as

$$\widetilde{E_t}=[0,e_{\max}]\ \forall t.$$

At any t after the first period, the true set S_t corresponds to the set $\widetilde{E_t}$ shifted to the right by ρs_t. To establish that $\widetilde{E_t}$ is a good alternative set, one needs to check two things: first, that the solution $\widetilde{\mathbf{e}}^*$ belongs to $\widetilde{E_t}\cap S_t=[\rho s_t,e_{\max}]$, so that $\mathbf{s}^*$ is in fact a feasible sequence in problem **AP**; and second, that the different domain does not make the solutions differ across the two problems. I now argue that, for the proposed set $\widetilde{E_t}$, both conditions are satisfied under a restriction on the values of ρ.

First, consider the feasibility of $\mathbf{s}^*$. That $\widetilde{e}_t<e_{\max}$ for all t is implied by the second condition in A4. That $\widetilde{e}_t>\rho s_{t-1}$ for all t can be checked using the solution $\widetilde{\mathbf{e}}^*$ to problem **AP** to determine if ρ satisfies the following condition:[7]

$$s^*\left(y^t,y_{t+1}\right)\geq \rho s^*\left(y^t\right) \qquad \forall\left(y^t,y_{t+1}\right)\text{ and }\forall t.$$

Note that this condition rules out combinations of parameters that would imply negative effort recommendations, which would be necessary if the drop in the required s from one period to the next was too big to be achieved just by letting the current s depreciate. In particular, I denote as $\rho^*_{\max}$ the maximum persistence that guarantees that effort is positive at all times:

$$\rho^*_{\max}=\min_{y^t\in Y^T,y_i\in Y}\left(\frac{\widetilde{e}^*\left(y^t,y_i\right)}{\widetilde{e}^*\left(y^t\right)}\right),$$

[5] Kocherlakota (2004) studies a related repeated moral hazard problem in which the agent has access to hidden savings. He demonstrates that the problem of the agent may fail to be globaly concave in effort and savings when the disutility of the agent is linear. Effort persistence could, in general, also imply that the problem in the IC of problem **OP** fails to be globally concave in effort across periods. Here, however, I only require the IC of the *auxiliary* problem **AP** to be concave in the current period's effort; since this problem is a standard repeated moral hazard, the usual conditions on the probability function suffice to guarrantee A4 (see Rogerson (1985b) and Jewitt, 1988.)

[6] I thank an anonymous referee for raising this issue.

[7] Numerical simulations suggest that this condition is fulfilled by a wide range of parameter values.

and study only problems with $\rho < \rho_{\max}$.

Second, consider the effect of the difference in the domains in the solution. Under the concavity and interiority implied by A3 and A4, the deviations added and those eliminated when using domain $\widetilde{E}_t$ are always dominated by deviations involving effort choices closer to the optimum, which is included both in $\widetilde{E}_t$ and S_t; hence, the solutions to the two problems are in fact the same when one uses $\widetilde{E}_t$ as the domain for $\widetilde{e}_t$.

I summarize the above discussion in the following proposition:

Proposition 1 *Consider a repeated moral hazard problem between a risk neutral principal and a risk averse agent with utility $u(c) - e$ in which the distribution of output is given by $\pi(y_t|s_t)$, where $s_t = \sum_{\tau=1}^{t} \rho^{\tau-1} e_\tau$ and $e_t \in [0, e_{\max}]$ $\forall t$. If the expected utility of the agent for a given contract is concave in s_t for all t and the optimal choice for s_t is always interior, there is an auxiliary repeated moral hazard problem where the agent's utility is given by $u(\widetilde{c}) - (1-\beta\rho)\widetilde{e}$ and the distribution of output is given by $\widetilde{\pi}(y_t|\widetilde{e}_t) = \pi(y_t|\widetilde{e}_t)$ which can be solved recursively. Furthermore, if $\rho \le \rho^*_{\max}$, we have that **(i)** the optimal consumption coincides in both problems and **(ii)** the optimal sequence of effort in the original problem, e^*, can be obtained from the solution to the auxiliary problem, $\widetilde{e}^*$, using the following system of equations:*

$$
\begin{aligned}
e_1^* &= s_1^* = \tilde{e}_1^*, \\
e_t^* &= s_t^* - \rho s_{t-1}^* = \tilde{e}_t^* - \rho \tilde{e}_{t-1}^* \quad t > 1.
\end{aligned}
$$

In the presence of persistence, when the agent increases effort today – incurring some disutility today – he can achieve the same productive state tomorrow with less effort, therefore saving some disutility tomorrow. This intertemporal trade–off implies, in general, that persistence introduces history dependence. The assumption of the productive state being a geometrically discounted sum of past efforts, combined with linearity in the disutility function, make the marginal cost and the marginal benefit of effort depend only on s, simplifying the history dependence.[8]

The intuition for the equivalence between problem **OP** and **AP** relies on the fact that with linear disutility, the actual period in which the agent experiences the cost of effort is not important. This means that the principal can solve the problem by choosing directly the optimal level of s in every period, modifying accordingly the utility function: persistence of effort can be understood as a lower marginal cost of acumulating s.

An important implication of the relationship I established between the problem with persistence and the auxiliary problem is that the sequence of contingent consumption will be exactly the same in both problems. This makes the two problems observationally equivalent. The results found in the moral hazard literature on the long run distribution of utilities and the individual consumption paths will also hold in the environment studied here with persistence and linear disutility.[9]

[8] Farhi (2006) independently proposes a similar set of simplifying assumptions in a capital taxation problem.

[9] See Rogerson (1985a) for a seminal contribution to the study of the properties of consumption dynamics in the optimal contract. See Spear and Srivastava (1987), Thomas and Worral (1990), Phelan and Townsend (1991), Phelan (1994, 1995), and Atkeson and Lucas (1995) for long run results and applications.

3 An example

Consider a problem in which there are only two possible outcomes each period, y_H and y_L. The utility of the agent is

$$U(c_t, e_t) = \frac{c^{1-\sigma}}{1-\sigma} - e.$$

The probability of the high outcome depends on s in the following way:

$$\pi(y_H | s_t) = 1 - \exp\{-r\sqrt{s_t}\} \ \forall t.$$

Under these specifications, A1–A4 are satisfied.

I compute the numerical solution to the optimal contract for $\rho = 0.4$, and compare it to the solution under the same parameters but with no persistence of effort, i.e. $\rho = 0$.[10] As seen in Figure 1, for a given w, the levels of consumption and continuation utilities are lower with persistence. Also, the set of feasible and incentive compatible values of w is bigger with persistence (higher values are included). As seen in the last plot of Fig. 1, for a given w, the value to the principal is higher with persistence. Also, higher levels of $s(w)$ are implemented (Fig. 2, bottom): consistently with what we expect from our knowledge of RMH problems, a lower marginal disutility means the solution $\widetilde{e}(w)$ in the **AP** problem is higher, which translates into higher $s(w)$ in problem **OP**.

The solution for effort is plotted in Fig. 2 (top). When $\rho = 0$, we have $e(w) = s(w) \ \forall w$. When $\rho = 0.4$, however, effort depends on both the current and the last period's promised utility (see Eq. 3). In the two–output case studied in this numerical example, there are (at most) two values of last period's w that are compatible with a given w today, each corresponding to a different output realization yesterday. This allows me to label the effort solution as $e_L(w)$ and $e_H(w)$ and to plot them as a function of w only.[11] The figure shows that, for a given w, effort is generally slightly lower with persistence than without. Note, however, that in the first period $e_1 = s_1(w_0)$ since $s_0 = 0$, which creates a front–loading effect on effort. Fig. 3 plots the average time path of both s and e over 7,000 realizations. The persistence of effort makes it efficient to build up the optimal level of s from the first period. If the disutility of effort was strictly convex, I conjecture that this front–loading force would still be present, although the build up of s would presumably take place over several periods. With convex disutility, however, the contract may exhibit quite different properties because of the joint deviations problem. The case characterized here can be a useful benchmark for future research in this subject.

[10] The rest of the parameters are: $\sigma = 1/2$, $y_H = 25$, $y_L = 8$, $r = 0.8$, $\beta = 0.85$, $T = \infty$. Note that the assumption $c_t \leq y_t$ puts a lower bound on s for high values of continuation utility, and hence the condition $\rho \leq \rho_{\max}$ is satisfied in the parameterization of the example.

[11] Generally, effort is a function of two states, that is, $e(w, w_{-1})$, where w_{-1} denotes the previous period continuation utility. Hence, $e_i(w)$, for $i = L, H$, is the effort recommendation today if the current promised utility is w, and it is the case that $w = w_i(w_{-1})$, where $w_i(\cdot)$ is the continuation value that the contract prescribes for the previous period contingent on output i. Some values of w are never chosen as a continuation utility (corresponding to the discontinuities in $e_i(w)$) because the numerical solution uses a discretized state space. Fitting a polynomial through the solutions, for example, produces a smooth and continuous effort function.

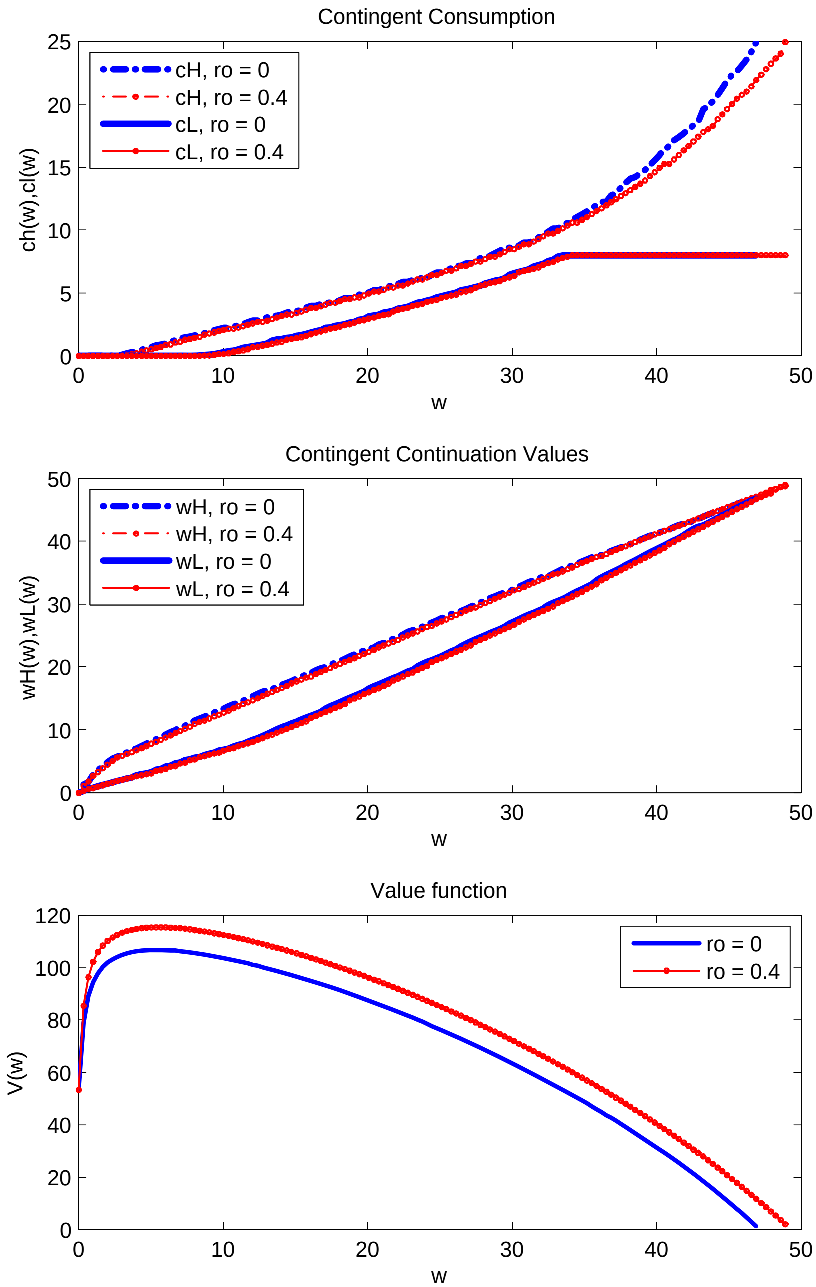

Figure 1: Computed solution for consumption, continuation utility and value function.

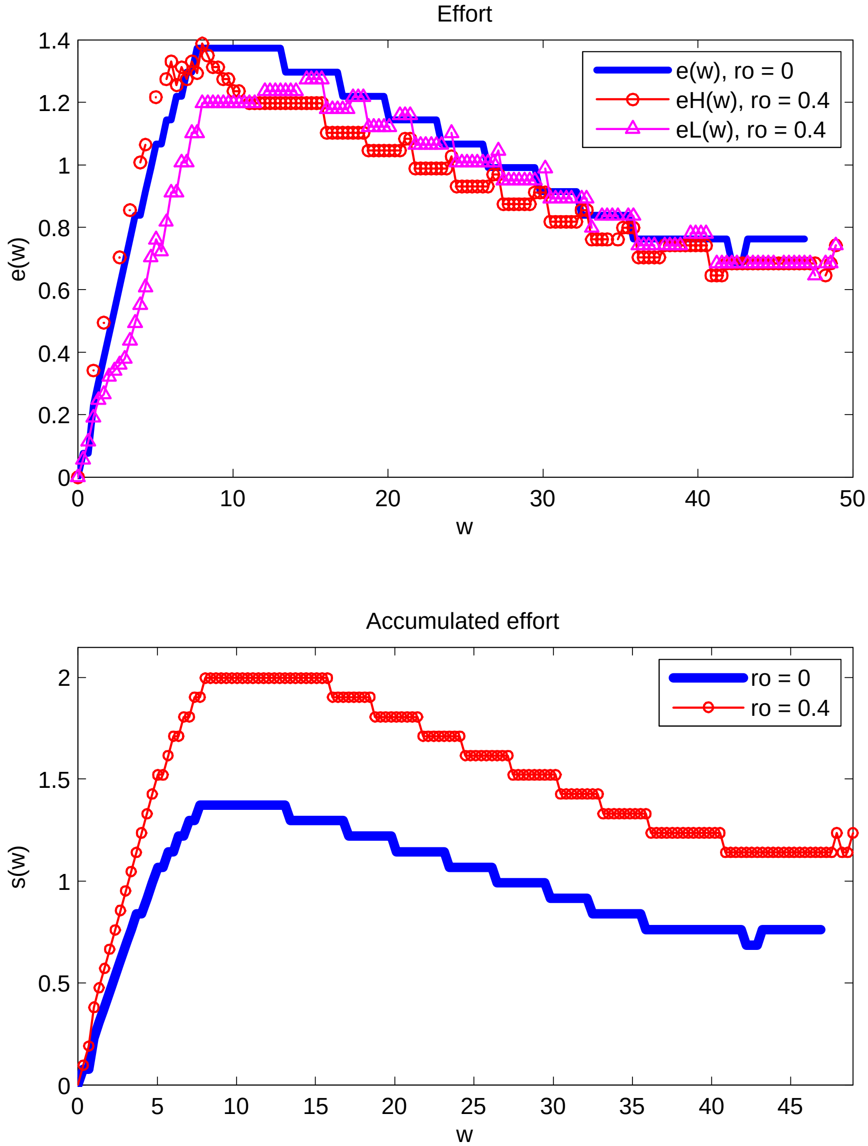

Figure 2: Computed solution for effort and accumulated effort.

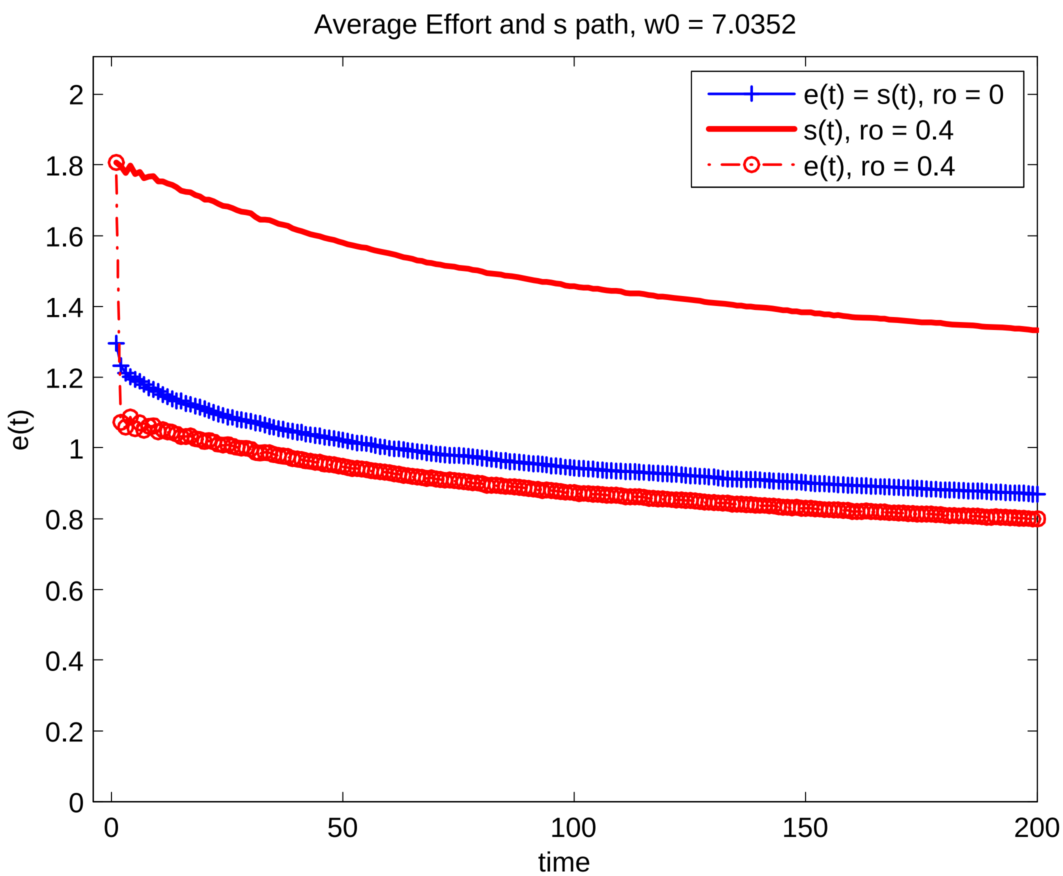

Figure 3: Example of an average 200 period long path for effort and accumulated effort, with 7,000 different path draws.

References

[1] Ábrahám, Á. and N. Pavoni, "Efficient Allocations with Moral Hazard and Hidden Borrowing and Lending," *mimeo*, University College London (2006).

[2] Atkeson, A. and R. E. Lucas, Jr, "Efficiency and Equality in a Simple Model of Efficient Unemployment Insurance," *Journal of Economic Theory,* **66** (1995), 64-88.

[3] Farhi, E., "Capital Taxation and Ownership when Markets are Incomplete," mimeo, Harvard University (2006).

[4] Fernandes, A. and C. Phelan, "A Recursive Formulation for Repeated Agency with History Dependence," *Journal of Economic Theory,* **91** (2000): 223-247.

[5] Jewitt, I., "Justifying the First–Order Approach to Principal–Agent Problems," *Econometrica*, **56**(5), 1177-1190 (1988).

[6] Kwon, I. "Incentives, Wages, and Promotions: Theory and Evidence," *Rand Journal of Economics*, **37** (1), 100-120 (2006).

[7] Mukoyama, T. and A. Sahin, "Repeated Moral Hazard with Persistence," *Economic Theory,* vol. **25**(4), pages 831-854, 06 (2005).

[8] Phelan, C., "Repeated Moral Hazard and One–Sided Commitment," *Journal of Economic Theory,* **66** (1995), 468-506.

[9] Phelan, C., "Incentives, Insurance, and the Variability of Consumption and Leisure," *Journal of Economic Dynamics and Control* **18**, Issues 3-4, (1994), Pages 581-599.

[10] Phelan, C. and R. M. Townsend, "Computing Multi–Period, Information–Constrained Optima," *Rev. Econ. Studies* (1991) **58**, 853-881.

[11] Rogerson, W. P. "Repeated Moral Hazard," *Econometrica*, **53**(1), pp. 69-76. (1985a).

[12] Rogerson, W. P. "The First–Order Approach to Principal–Agent Problems," *Econometrica*, **53**(6), pp. pp. 1357-1367 (1985b).

[13] Spear, S. E. and S. Srivastava, "On Repeated Moral Hazard with Discounting," *Rev. Econ. Studies* **54** (1987), 599-617.

[14] Thomas, J. and T. Worral, "Income Fluctuations and Asymetric Information: An Example of Repeated Principal–Agent Problem," *Journal of Economic Theory,* **51** (1990), 367:390.

[15] Werning, I., "Moral Hazard with Unobserved Endowments: A Recursive Approach," *mimeo*, University of Chicago (2001).