

Repeated Moral Hazard and Recursive Lagrangeans

Antonio Mele*

Universitat Pompeu Fabra, Barcelona (Spain)

February 2008

PRELIMINARY DRAFT

Abstract

I solve a repeated moral hazard model with a fast and flexible numerical algorithm. Instead of applying the traditional Abreu, Pierce and Stacchetti (1990), I extend the Lagrangean techniques developed in Marcet and Marimon (1998) to the principal-agent framework. A numerical procedure is proposed, that is much faster than the traditional algorithms based on the promised utilities approach, and that can easily deal with large state spaces. Given the computational speed, the algorithm is especially suitable for applications with many state variables and for calibration purposes.

1 Introduction

There has been a huge amount of research, in last two decades, on the repeated principal agent model. Recent contributions on dynamic agency theory have focused both on the case in which agent's consumption is observable (see for example Rogerson (1985a), Spear and Srivastava (1987), Thomas and Worrall (1990), Phelan and Townsend (1991), Fernandes and Phelan (2000)) and more recently on the case in which agents can secretly save and borrow (Werning (2001), Abraham and Pavoni (2006, 2008)). The main difficulty in the repeated moral hazard framework is that the optimal contract is history-dependent, i.e. it is not recursive over its natural state space, and therefore it is not possible to directly apply dynamic programming algorithms to solve it. The traditional way of dealing with this problem is based on the *expected promised utilities approach*: a new state space is defined that includes all the natural state variables and the agent's continuation value. By a standard argument due to Spear and Srivastava (1987), it can be shown that

***Acknowledgements:** I am grateful to Albert Marcet for smart suggestions and long fruitful discussions on the topic. I also owe a special thank to Luigi Balletta and Sevi Rodriguez-Mora for advices at a very early stage of the work, and to Davide Debortoli and Ricardo Nunes for their generosity in discussing infinitely many numerical aspects of the paper. This paper has benefitted from comments by Klaus Adam, Sofia Bauducco, Francesco Caprioli, Martina Cecioni, Josè Dorich, Harald Fadinger, Eva Luethi, Angelo Mele, Matthias Messner, Krisztina Molnar, Nicola Pavoni, Josep Pijoan-Mas, Michael Reiter and Gilles Saint-Paul. A special thank to participants at UPF MACRO BREAK, and to students participating at UPF MACRO DISCUSSION GROUP. All mistakes are mine. **email:** antonio.mele@upf.edu

the problem is recursive on the new state space, and therefore can be solved numerically by value function iteration. However, there is an additional complication: promised utilities are constrained to belong to a feasible set, which has to be characterized by solving a fixed point problem over a set operator.

This approach has a limitation: as the number of natural state variables grows, it becomes unmanageable, since the dimensionality of the feasible set of promised utilities increases when the dimensionality of the natural state space increases. It is easy to characterize this set if there is just one exogenous shock; but if the problem includes capital accumulation or other endogenous states (like in the literature of dynamic agency with hidden states, for example), then the set is multidimensional, which implies that the fixed point of the feasible set is extremely difficult to find, if not computationally impossible (see for example the discussion in Abraham and Pavoni (2006)).

This paper provides two contributions to the literature. First, I adapt the techniques of recursive Lagrangeans developed in Marcet and Marimon (1999) (MM from now on) to the dynamic agency model, and I show that this idea translates in a very simple and intuitive way of characterizing the optimal contract. Second, I provide an algorithm based on the recursive Lagrangean which is much faster¹ than the usual dynamic programming techniques and does not suffer from the same dimensionality issues as the promised utilities approach.

Here are the steps of the procedure: I solve the agent's problem, *given the principal's strategy*, by taking first-order conditions; then, I use agent's first-order conditions as constraints in order to write the Lagrangean of the principal². Using arguments similar to MM, I show that this problem is recursive in an enlarged state space, which includes the natural states and an endogenously evolving Pareto-Negishi weight. Given this result, I can obtain a solution from Lagrangean first order conditions.

This methodology is much simpler to implement and less mathematically and computationally demanding than Abreu, Pierce and Stachetti (1990) (APS in the following). The numerical algorithm just amounts to find the roots of a nonlinear system of equations (the Lagrangean first order conditions). The main gain is the possibility to deal with large state spaces and solve complicated models in a reasonable amount of time, which is useful in macroeconomic problems with heterogeneous agents (especially when one would like to calibrate the model to real data).

The constrained-efficient allocation is characterized by an *endogenously evolving Pareto-Negishi weight*³: this variable summarizes the principal's promises, by rewarding the agent with more consumption if a "good" realization of the state of nature is observed, and by punishing him with less consumption if a "bad" outcome happens. This Pareto-Negishi weight not only has the above mentioned intuitive interpretation, but also allows to characterize many standard results of the dynamic agency literature in a straightforward way. For

¹The numerical solution is obtained in around 2 seconds in a state-of-the-art personal computer.

²In order to be sure that agent's first-order conditions are sufficient to get the optimal solution for the problem of the principal, I assume that Rogerson (1985b) conditions of monotone likelihood ratio and convex distribution function are satisfied.

³In the case with hidden assets, the optimal contract is characterized by two co-state variables: the endogenous Pareto-Negishi weight and the Lagrange multiplier of the agent's Euler equation.

example, Rogerson (1985a) shows that optimal contracts satisfy an *"inverted Euler equation"* in which the inverse of marginal utility of consumption evolves as a martingale; the same result is obtained here by taking expectations of the law of motion of the Pareto-Negishi weight (see Proposition 2).

The Lagrangean approach has been criticized by Messner and Pavoni (2004): they show that the functional equation associated with the problem can yield unfeasible or suboptimal solutions, unless the problem is strictly concave. However, this issue is irrelevant to my approach, since I use the Lagrangean first order conditions to characterize the solution, therefore feasibility is imposed directly (i.e., feasibility constraints are among the equations that the numerical algorithm has to solve for). The functional equation is used in this paper only to show that policy functions are recursive, and this fact is exploited in the numerical computations: hence, I solve for Markovian feasible policy functions. Moreover, my numerical method yields a decreasing and strictly concave Pareto frontier, a case that also Messner and Pavoni recognize as not problematic for the Lagrangean method.

The paper is organized as follows: Section 2 introduces the basic framework, Section 3 gives a simple explanation of the methodology applied. Section 4 applies the techniques informally described in Section 3 to a standard problem of repeated moral hazard. Section 5 shows how to extend the basic formulation to models with endogenous state variables and provides the Lagrangean solution of a model with hidden effort and hidden assets. Section 6 gives some numerical examples of the application of the technique. Section 7 concludes.

2 The basic model

The economy is inhabited by a risk neutral principal and a risk averse agent. Time is discrete, and the state of the world is determined by an observable Markov state process $\{s_t\}_{t=0}^{\infty}$, where $s_t \in S$, and $\#S = I$. I assume s_0 is known, and the realizations of the process are public information. I will denote with subscripts the single realizations, and with superscripts the histories:

$$s^t \equiv \{s_0, \dots, s_t\} \in S^{t+1}$$

At each period, the agent gets a state-contingent income flow $y(s_t)$, enjoys consumption $c_t(s^t)$, receives a transfer $\tau_t(s^t)$ from the principal, and exerts a costly unobservable action $a_t(s^t) \in A \subseteq \mathbb{R}_+$, $\{0\} \in A$. I will refer to $a_t(s^t)$ as action or effort. The costly action affects the probability distribution of the state of the world tomorrow. Let $\pi(s_{t+1} = s_i | s_t, a_t(s^t))$ be the probability that state tomorrow is $s_i \in S$ conditional on past state and effort exerted by the agent at the beginning of the period⁴, with $\pi(s_0) = 1$. I assume $\pi(\cdot)$ is twice continuously differentiable in $a_t(s^t)$, and there is *full support*: $\pi(s_{t+1} = s_i | s_t, a) > 0 \forall i, \forall a$. Let $\Pi(s^{t+1} | s_0, a^t(s^t)) = \prod_{j=0}^t \pi(s_{j+1} | s_j, a_j(s^j))$ be the cumulated probability of an history s^{t+1} given the whole history of unobserved actions $a^t(s^t) \equiv (a_0(s^0), a_1(s^1), \dots, a_t(s^t))$.

⁴Notice that I allow for persistence; in the numerical examples, I focus on i.i.d. shocks, but it should be clear that persistence does not create particular problems neither theoretically nor numerically.

The instantaneous utility of the agents is

$$u(c_t(s^t)) - v(a_t(s^t))$$

with $u(\cdot)$ strictly increasing, strictly concave and satisfying Inada conditions, while $v(\cdot)$ is strictly increasing and strictly convex; both are twice continuously differentiable. Agents do not accumulate assets autonomously: the only source of insurance is the principal. Then, the budget constraint of an agent will be simply:

$$c_t(s^t) = y(s_t) + \tau_t(s^t) \quad \forall s^t, t \geq 0$$

For simplicity, let $\hat{s}_i$, $i = 1, 2, \dots, I$ be the possible realizations of $\{s_t\}$ and let them be ordered such that $y(s_t = \hat{s}_1) < y(s_t = \hat{s}_2) < \dots < y(s_t = \hat{s}_I)$. I assume that the action $a_t(s^t)$ cannot be observed by the principal. The principal only observes the realization of the state of nature s_t . Both principal and agents are fully committed once they sign the contract at time zero.

A *contract* (or *allocation*) in this framework is a plan $(a^\infty, c^\infty, \tau^\infty) \equiv \{a_t(s^t), c_t(s^t), \tau_t(s^t)\}_{t=0}^\infty$ that belongs to the following set:

$$\begin{aligned} \Gamma \equiv \{ & (a^\infty, c^\infty, \tau^\infty) : a_t(s^t) \in A, \quad c_t(s^t) \geq 0, \\ & \tau_t(s^t) = c_t(s^t) - y(s_t) \quad \forall s^t \in S^{t+1}, t \geq 0 \} \end{aligned}$$

3 Methodological overview

The recent literature on repeated moral hazard has focused on a now widely known technique, developed during late '80s by APS and Spear and Srivastava (1987), which helps to overcome the history-dependence that characterizes optimal contracts with repeated moral hazard (see e.g., Rogerson (1985a)). APS method introduces a new endogenous state variable in the model, which corresponds to expected discounted utility⁵. In the typical application, the model at hand is proved to be recursive in an enlarged state space that includes expected discounted utility and natural states of the economy (i.e., endogenous or exogenous state variables, like capital stock or technology shocks). The main limitation of the APS approach is the need to calculate the feasible set for expected discounted utilities, the dimensionality of which depends on the natural states. When the cardinality of the state space is large, it becomes very burdensome (if not impossible) to characterize this set.

This is true, for example, in the case of dynamic moral hazard with hidden assets. Abraham and Pavoni (2006) (AP from here on) analyze a model with hidden effort and hidden assets: the agent can borrow or lend without being monitored by the principal. The natural states of the model are this income and hidden assets. Unfortunately, this problem generates a continuum of incentive constraints (for each realizations

⁵For an introduction to APS techniques, see Ljunqvist and Sargent (2004)

of the income, there is a continuum of asset positions), and the feasible set of continuation values has infinite dimension. Hence, the APS technique is problematic. In order to avoid this complication, they characterize the optimal contract by defining an auxiliary problem, where agent's first-order conditions over effort and bonds are used as constraints for the principal's problem. They show that the marginal utility of consumption becomes a state variable. Their auxiliary problem is therefore characterized by three states: income, expected discounted utility and consumption marginal utility, and can be solved recursively by value function iteration. However, in models with hidden assets, the use of agent's first-order conditions can be problematic, because those constraints do not always generate a convex set, hence convex optimization techniques cannot be applied. AP therefore implement an ex-post verification algorithm, that checks if the allocation obtained is globally optimal.

In this paper, I use a first-order approach in the same line of AP, by substituting incentive constraints with agent first order conditions. However, instead of applying APS techniques, I follow MM methodology of recursive Lagrangeans. This methodology allows for the presence of many endogenous states without big computational issues.

In a standard repeated moral hazard problem with no asset accumulation, Rogerson (1985b) shows that we can guarantee that first-order approach yields the optimal solution by assuming that the probability functions satisfy two conditions: *monotonicity of likelihood ratio* and *convexity of the distribution function*. (see details below). Given these assumptions, it is possible to solve the planner problem by using agent's first-order condition with respect to effort, instead of the (much more complicated) incentive compatibility constraint. In order to solve the principal's problem, I write down the corresponding Lagrangean and I manipulate it algebraically, getting an endogenously evolving Pareto-Negishi weight, which is the new co-state variable of the problem. By using an argument similar to MM, the problem is shown to be recursive in the state space that includes income and the Pareto-Negishi weight. This means that we can restrict ourselves to look for a Markovian solution, i.e. policy functions that only depend on natural states and the new costate.

For the case with hidden assets, there is a possibility that the constraint set is not convex even if the two Rogerson (1985b) conditions are satisfied. To the best of my knowledge, there is not any restrictive assumption that makes sure the problem is convex, but there are numerical procedures that can be used to verify the global optimality of a solution obtained from the Lagrangean (the ex-post verification algorithm in Abraham and Pavoni (2006) is an excellent candidate). Therefore, the methodology has general applicability, provided that a verification of the optimality is implemented.

The numerical algorithm solves Lagrangean first order conditions by collocation method, parameterizing the policy functions and the conditional expectations (i.e. approximating them as functions of the co-state variable and the income shock). This numerical procedure is very fast compared to APS, since it skips both the set operator fixed point problem and the optimization step: with the current version of the algorithm, the

solution of a standard repeated moral hazard problem requires only few seconds of machine time⁶. Moreover, additional states can be added to the basic framework at the cost of small complications, even if we cannot guarantee that Rogerson's (1985b) conditions are sufficient. In this case, the ex-post verification procedure proposed by AP can be used to check for global optimality.

4 Repeated moral hazard

Assume, for simplicity, that the discount factor of the agent and the principal is the same. The principal evaluates allocations according to the following

$$\begin{aligned} W(s_0; a^\infty, c^\infty, \tau^\infty) &= - \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \tau_t(s^t) \Pi(s^t | s_0, a^{t-1}(s^{t-1})) \\ &= \sum_{t=0}^{\infty} \sum_{s^t} \beta^t [y(s^t) - c_t(s^t)] \Pi(s^t | s_0, a^{t-1}(s^{t-1})) \end{aligned} \quad (1)$$

therefore efficient contracts can be characterized by solving the principal problem of maximizing (1), subject to incentive compatibility and to the requirement of providing at least a minimum level of ex-ante utility V^{out} to the agent:

$$\begin{aligned} W^*(s_0) &= \max_{\{a_t(s^t), c_t(s^t)\}_{t=0}^{\infty} \in \Gamma} \sum_{t=0}^{\infty} \sum_{s^t} \beta^t [y(s^t) - c_t(s^t)] \Pi(s^t | s_0, a^{t-1}(s^{t-1})) \\ s.t. \quad a^\infty &\in \arg \max_{\{a_t(s^t)\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \sum_{s^t} \beta^t [u(c_t(s^t)) - v(a_t(s^t))] \Pi(s^t | s_0, a^{t-1}(s^{t-1})) \\ &\sum_{t=0}^{\infty} \sum_{s^t} \beta^t [u(c_t(s^t)) - v(a_t(s^t))] \Pi(s^t | s_0, a^{t-1}(s^{t-1})) \geq V^{out} \end{aligned} \quad (2)$$

We will call this *the original problem*. The last constraint is called the *promise-keeping constraint* (PKC in the following) in the literature. Notice that (2) is a very complicated object. In this work, I use the first order conditions of the agent's problem as a substitute for the constraint (2). In order to be sure that this substitution will lead to the actual solution of the original problem, we assume the following two conditions hold (see Rogerson (1985b)):

Condition 1 (Monotone Likelihood-Ratio Condition (MLRC)) $\hat{a} \leq \hat{\hat{a}} \implies \frac{\pi(s_{t+1}=s_i | s_t, \hat{a})}{\pi(s_{t+1}=s_i | s_t, \hat{\hat{a}})}$ is nonincreasing in i .

The above property can be restated in a simpler way: if $\pi(\cdot)$ is differentiable, then MLRC is equivalent to $\frac{\pi_a(s_{t+1}=s_i | s_t, a)}{\pi(s_{t+1}=s_i | s_t, a)}$ being nondecreasing in i for any a , where $\pi_a(s_{t+1}=s_i | s_t, a)$ is the derivative of $\pi(\cdot)$ with respect to a . An important consequence of the MLRC is the following: let $F(\cdot)$ be the cumulative distribution function of $\pi(\cdot)$; then MLRC implies $F'(s_{t+1}=s_i | s_t, a)$ is nonpositive for any i and every

⁶With a personal computer with processor Intel Core 2 6420 2.13GHz and 2 Gb RAM, the Matlab code needs about 2 seconds to converge.

a. Therefore, more effort implies a first order stochastic dominance shift of the distribution (see Rogerson (1985b)).

Condition 2 (Convexity of the Distribution Function Condition (CDFC)) $F''(s_{t+1} = s_i | s_t, a)$ is nonnegative for any i and every a .

We now define the problem of the agent and we derive his first order conditions with respect to effort, which will be our substitute for the incentive compatibility constraint.

The problem of the agent, given the principal's strategy profile $\tau^\infty \equiv \{\tau_t(s^t)\}_{t=0}^\infty$, is:

$$V(s_0; \tau^\infty) = \max_{\{c_t(s^t), a_t(s^t)\}_{t=0}^\infty \in \Gamma} \left\{ \sum_{t=0}^\infty \sum_{s^t} \beta^t [u(c_t(s^t)) - v(a_t(s^t))] \Pi(s^t | s_0, a^{t-1}(s^{t-1})) \right\}$$

The first order condition for effort is:

$$\begin{aligned} v'(a_t(s^t)) &= \sum_{j=1}^\infty \beta^j \sum_{s^{t+j}|s^t} \pi_a(s_{t+1} | s_t, a_t(s^t)) \times \\ &\times [u(c_{t+j}(s^{t+j})) - v(a_{t+j}(s^{t+j}))] \Pi(s^{t+j} | s^{t+1}, a^{t+j}(s^{t+j} | s^{t+1})) \end{aligned} \quad (3)$$

Intuitively, the marginal cost of effort today (LHS) has to be equal to future expected benefits (RHS). In particular, we can notice that effort today affects all future expected utilities through its marginal effect on tomorrow's outcome distribution ($\pi_a(s_{t+1} | s_t, a_t(s^t))$).

In the following, for simplicity I will refer to (3) as the *incentive-compatibility constraint* (ICC). We can write the problem of the principal as:

$$\begin{aligned} W^*(s_0) &= \max_{\{\tau_t(s^t), a_t(s^t), c_t(s^t)\}_{t=0}^\infty \in \Gamma} \sum_{t=0}^\infty \sum_{s^t} \beta^t [y(s_t) - c_t(s^t)] \Pi(s^t | s_0, a^{t-1}(s^{t-1})) \\ s.t. \quad v'(a_t(s^t)) &= \sum_{j=1}^\infty \beta^j \sum_{s^{t+j}|s^t} \frac{\pi_a(s_{t+1} | s_t, a_t(s^t))}{\pi(s_{t+1} | s_t, a_t(s^t))} \times \\ &\times [u(c_{t+j}(s^{t+j})) - v(a_{t+j}(s^{t+j}))] \Pi(s^{t+j} | s^t, a^{t+j-1}(s^{t+j-1} | s^t)) \\ &\quad \forall s^t, t \geq 0 \\ &\sum_{t=0}^\infty \sum_{s^t} \beta^t [u(c_t(s^t)) - v(a_t(s^t))] \Pi(s^t | s_0, a^{t-1}(s^{t-1})) \geq V^{out} \end{aligned} \quad ((PP \ 1))$$

It should be obvious to the reader that (PP 1) is not recursive: *ICC* implies that today's effort depends on tomorrow's effort and consumption. In other words, the problem of the principal is affected by time-inconsistency.

4.1 The Lagrangean approach

We will write down the Lagrangean and transform the problem in a recursive one with the techniques developed in MM. First of all, notice that the PKC must be binding in the optimum: otherwise, the

principal can increase her expected discounted utility by asking the agent to increase effort in period 0 by $\delta > 0$, provided that δ is small enough. PKC therefore will be associated with a strictly positive Lagrange multiplier (say, γ), which will be a function of V^{out} , i.e. for every V^{out} , there will be a unique γ associated with the PKC. Therefore, this Lagrange multiplier can be seen as a Pareto-Negishi weight on the agent's utility. Since each γ implies a unique V^{out} , by solving the problem for different values of γ between zero and infinity, we can fully characterize the Pareto frontier of this economy. Hence, in the following, I am going to consider γ as a parameter, that represents the PKC. Moreover, notice that by fixing γ , V^{out} will appear in the Lagrangean only in the constant term γV^{out} , thus it will be irrelevant for the optimal allocation. I can therefore write the Lagrangean as:

$$\begin{aligned} L(s_0, \gamma, c^\infty, a^\infty, \lambda^\infty) = & \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \{y(s_t) - c_t(s^t) + \gamma [u(c_t(s^t)) - v(a_t(s^t))]\} \Pi(s^t | s_0, a^{t-1}(s^{t-1})) + \\ & - \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \lambda_t(s^t) \left\{ v'(a_t(s^t)) - \sum_{j=1}^{\infty} \beta^j \sum_{s^{t+j}|s^t} \frac{\pi_a(s_{t+1} | s_t, a_t(s^t))}{\pi(s_{t+1} | s_t, a_t(s^t))} \times \right. \\ & \times [u(c_{t+j}(s^{t+j})) - v(a_{t+j}(s^{t+j}))] \Pi(s^{t+j} | s^t, a^{t+j-1}(s^{t+j-1} | s^t)) \} \Pi(s^t | s_0, a^{t-1}(s^{t-1})) \end{aligned}$$

where $\lambda_t(s^t)$ is the Lagrange multiplier of ICC. MM show that the Lagrangean can be manipulated with simple algebra to get the following expression:

$$\begin{aligned} L(s_0, \gamma, c^\infty, a^\infty, \lambda^\infty) = & \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \{y(s_t) - c_t(s^t) + \phi_t(s^t) [u(c_t(s^t)) - v(a_t(s^t))]\} + \\ & - \lambda_t(s^t) v'(a_t(s^t)) \} \Pi(s^t | s_0, a^{t-1}(s^{t-1})) \end{aligned}$$

where

$$\phi_{t+1}(s^t, \hat{s}) = \phi_t(s^t) + \lambda_t(s^t) \frac{\pi_a(s_{t+1} = \hat{s} | s_t, a_t(s^t))}{\pi(s_{t+1} = \hat{s} | s_t, a_t(s^t))} \quad \forall \hat{s} \in S \quad (4)$$

$$\phi_0(s^0) = \gamma$$

The multiplier $\lambda_t(s^t)$ is the *planner's shadow cost to enforce an incentive-compatible allocation*. The essence of MM is the (new co)state variable $\phi_t(s^t)$: by definition, it corresponds to the sum of all λ 's multiplied by the likelihood ratio, from period 0 till period $t-1$ plus the constant γ

$$\begin{aligned} \phi_t(s^{t-1}, s_t) &= \phi_{t-1}(s^{t-1}) + \lambda_{t-1}(s^{t-1}) \frac{\pi_a(s_t | s_{t-1}, a_{t-1}(s^{t-1}))}{\pi(s_t | s_{t-1}, a_{t-1}(s^{t-1}))} = \\ &= \gamma + \sum_{i=0}^{t-1} \lambda_i(s^i) \frac{\pi_a(s_{i+1} | s_i, a_i(s^i))}{\pi(s_{i+1} | s_i, a_i(s^i))} \end{aligned}$$

The intuition is simple. For any s^t , the expression $\lambda_t(s^t) \frac{\pi_a(s_{t+1} | s_t, a_t(s^t))}{\pi(s_{t+1} | s_t, a_t(s^t))}$ is a *planner's promise* about how much she will care about agent's welfare in the future, depending on which realization of state s_{t+1} is observed. By keeping track of all λ 's and $\frac{\pi_a}{\pi}$'s realized in the past, $\phi_t(s^t)$ summarizes all the promises made by the planner in previous periods.

In this framework, there is straightforward interpretation of $\phi_t(s^t)$: it is the *Pareto-Negishi weight* of the agent's utility, that evolves *endogenously* in order to take into account the effort of the agent. To simplify things, let us assume there are only two possible realizations of the state of nature: $s_t \in \{s_L, s_H\}$. At time 0, the weight is equal to γ . In period 1, given our assumption on the likelihood ratio, the Pareto-Negishi weight is higher than γ if the principal observes s_H , while it is lower than γ if she observes s_L (see Lemma 1 below). Therefore the agent is induced to exert effort by the principal's promise to care more about agent's welfare in the next periods.

There are two technical problems in this setup. First, the endogenous evolution of the Pareto-Negishi weight is a deviation from the MM framework, since in their paper the costate law of motion only depends on $\lambda_t(s^t)$, while here also depends on $a_t(s^t)$. Second, the probability distribution of the future states is endogenous. These facts may potentially complicate the proof that the problem is recursive in the costate variable $\phi_t(s^t)$, but I show in Proposition 1 that the argument works also here with some minor modifications. Therefore, the next step is to make sure that our problem is recursive in the costate variable $\phi_t(s^t)$.

4.2 Recursive formulation

By the duality theory (see for example Luenberger (1969)), we know that a solution of the original problem corresponds to a saddle point of the Lagrangean, i.e. the contract $\{c_t^*(s^t), a_t^*(s^t), y(s_t) - c_t^*(s^t)\}_{t=0}^\infty$ is a solution for the principal's problem if $(c^{\infty*}, a^{\infty*}, \lambda^{\infty*}) = \{c_t^*(s^t), a_t^*(s^t), \lambda_t^*(s^t)\}_{t=0}^\infty$ satisfy:

$$L(s_0, \gamma, c^{\infty}, a^{\infty}, \lambda^{\infty*}) \leq L(s_0, \gamma, c^{\infty*}, a^{\infty*}, \lambda^{\infty*}) \leq L(s_0, \gamma, c^{\infty*}, a^{\infty*}, \lambda^{\infty})$$

MM show that we can associate a saddle-point functional equation to the Lagrangean. Let

$$r(a, c, \lambda, s, \phi) = y(s) - c + \phi[u(c) - v(a)] - \lambda v'(a)$$

Then, the saddle-point functional equation associated with our problem is:

$$\begin{aligned} J(s, \phi) &= \min_{\lambda} \max_{a, c} \left\{ r(a, c, \lambda, s, \phi) + \beta \sum_{s'} \pi(s' | s, a) J(s', \phi'(s')) \right\} \\ \text{s.t. } \phi'(s') &= \phi + \lambda \frac{\pi_a(s' | s, a)}{\pi(s' | s, a)} \quad \forall s' \end{aligned} \tag{5}$$

As usual, in order to show that the problem is recursive, it is sufficient to prove that the operator associated with the functional equation is a contraction.

Proposition 1 *Fix an arbitrary constant $K > 0$ and let $K_\phi = \max\{K, K\|\phi\|\}$. The operator*

$$\begin{aligned} (T_K f)(s, \phi) &\equiv \min_{\{\lambda > 0: \|\lambda\| \leq K_\phi\}} \max_{a, c} \left\{ r(a, c, \lambda, s, \phi) + \beta \sum_{s'} \pi(s' | s, a) f(s', \phi'(s')) \right\} \\ \text{s.t. } \phi'(s') &= \phi + \lambda \frac{\pi_a(s' | s, a)}{\pi(s' | s, a)} \quad \forall s' \end{aligned}$$

is a contraction.

Proof. *In the Proof's appendix. ■*

Proposition 1 is the crucial result of the paper: it states that, if we bound the costate variable, our problem is recursive in the state space (s, ϕ) . All the other theorems in MM apply directly to our framework. This means that the result of Proposition 1 is valid for any $K > 0$. Therefore, a recursive solution of the Problem (5) is a solution of the Lagrangean, and more importantly it is a solution of the original problem. As a consequence, we can restrict the search of optimal contracts to the set of Markovian policy functions in the space (s, ϕ) .

Notice that, since in the Lagrangean formulation we eliminated the constant γV^{out} , the value of the original problem is:

$$W(s_0; \tau^{\infty*}, a^{\infty*}, c^{\infty*}) = W^*(s_0) = J(s_0, \gamma) - \gamma V^{out}$$

There are two numerical strategies to get the optimal contract. The first is value function iteration of the saddle-point functional equation. But as well known, this is a slow procedure, since it involves the optimization step at each iteration. Therefore, we adopt a second strategy: we solve for the Lagrangean first order conditions.

4.3 Characterization of the optimal contract

We can now take the first order conditions of the Lagrangean:

$$c_t(s^t) : \quad 0 = -1 + \phi_t(s^t) u_c(c_t(s^t)) \quad (6)$$

$$\begin{aligned} a_t(s^t) : \quad & 0 = -\lambda_t(s^t) v''(a_t(s^t)) - \phi_t(s^t) v'(a_t(s^t)) + \\ & + \sum_{j=1}^{\infty} \beta^j \sum_{s^{t+j}|s^t} \frac{\pi_a(s_{t+1}|s_t, a_t(s^t))}{\pi(s_{t+1}|s_t, a_t(s^t))} \{y(s_t) - c_t(s^t) - \lambda_{t+j}(s_{t+j}) v'(a_{t+j}(s^{t+j})) - \\ & + \phi_{t+j}(s^{t+j}) [u(c_{t+j}(s^{t+j})) - v(a_{t+j}(s^{t+j}))]\} \Pi(s^{t+j} | s^t, a^{t+j-1}(s^{t+j-1})) + \\ & + \beta \lambda_t(s^t) \sum_{s^{t+1}|s^t} \frac{\partial \left(\frac{\pi_a(s_{t+1}|s_t, a_t(s^t))}{\pi(s_{t+1}|s_t, a_t(s^t))} \right)}{\partial a} [u(c_{t+1}(s^{t+1})) - v(a_{t+1}(s^{t+1}))] \pi(s_{t+1} | s_t, a_t(s^t)) \end{aligned}$$

and

$$\begin{aligned} \lambda_t(s^t) : \quad & 0 = -v'(a_t(s^t)) + \sum_{j=1}^{\infty} \sum_{s^{t+j}|s^t} \frac{\pi_a(s_{t+1}|s_t, a_t(s^t))}{\pi(s_{t+1}|s_t, a_t(s^t))} \times \\ & \times [\beta^j [u(c_{t+j}(s^{t+j})) - v(a_{t+j}(s^{t+j}))] \Pi(s^{t+j} | s^t, a^{t+j-1}(s^{t+j-1} | s^t))] \end{aligned} \quad (7)$$

We can use these equations to show some properties of the optimal contract. Lemma 1 makes clear how $\phi_t(s^t)$ incorporates the promises of the principal. From (6) we can see that $c_{t+1}(s^{t+1}) = u_c^{-1}\left(\frac{1}{\phi_{t+1}(s^{t+1})}\right)$, then $c_{t+1}(s^{t+1})$ is increasing in $\phi_{t+1}(s^{t+1})$. Lemma 1 says that, tomorrow, the principal will reward a high income realization with higher consumption than today, and a low income realization with lower consumption than today.

Lemma 1 *In the optimal contract, $\phi_{t+1}(s^t, \hat{s}_1) < \phi_t(s^t) < \phi_{t+1}(s^t, \hat{s}_I)$ for any t .*

Proof. *In the Proof's appendix. ■*

The following Proposition characterizes the time series properties of the Pareto Negishi weight.

Proposition 2 *$\phi_t(s^t)$ is a martingale that converges to zero almost surely.*

Proof. *In the Proof's appendix. ■*

Proposition 2 is the well known result that $\frac{1}{u_c(c_t(s^t))}$ evolves as a martingale (see Rogerson (1985a)). The a.s.-convergence to zero is the so called *immiseration property* that implies zero consumption almost surely as $t \rightarrow \infty$, which is a standard result in models with asymmetric information (see Thomas and Worrall (1990), for example). In this framework, the immiseration property has an intuitive interpretation: in order to keep strong incentives for the agent, the planner must ensure that the Pareto-Negishi weight goes to zero as $t \rightarrow \infty$ for any possible sequence of realizations of the income shock.

The result in Proposition 2 is obtained by using the law of motion of $\phi_t(s^t)$ and (6), which yields

$$E_t^a \left[\frac{1}{u_c(c_{t+1}(s^{t+1}))} \right] = \frac{1}{u_c(c_t(s^t))}$$

We can use Jensen's inequality and the strict concavity of $u(\cdot)$ to get that $E_t^a [u_c(c_{t+1}(s^{t+1}))] > u_c(c_t(s^t))$: the profile of expected consumption is decreasing across time.

5 Adding hidden assets to the problem

As anticipated, the main difficulties of APS technique arise in models with many endogenous state variables. In this section, I show how the Lagrangean approach can easily deal with this complication, providing an example of repeated moral hazard with hidden assets.

Let $\{b_t(s^t)\}_{t=-1}^\infty$, b_{-1} given, be a sequence of one-period bond that the agent pays 1 today, getting R tomorrow. Assume that the principal cannot monitor the bond market, so that the asset accumulation is unobservable to her. Then agent's budget constraint becomes:

$$c_t(s^t) + b_t(s^t) = y(s_t) + \tau_t(s^t) + Rb_{t-1}(s^{t-1})$$

Accordingly, agent's first order conditions are (3) and the following Euler equation:

$$u'(c_t(s^t)) = \beta R \sum_{s_{t+1}} u'(c_t(s^t, s_{t+1})) \pi(s_{t+1} | s_t, a_t(s^t)) \quad (8)$$

Assume $\beta R = 1$ to simplify algebra. The presence of hidden assets requires (8) to be included in the set of constraints for the principal's problem. Notice that Conditions 1 and 2 are not sufficient to ensure that

first-order approach is correct (see, on this point, the discussion in AP). Still, we can deal numerically with this problem by using the ex-post verification method proposed in AP.

Let $\eta_t(s^t)$ be the Lagrange multiplier for (8). The Lagrangean can be manipulated to get:

$$L(s_0, \gamma, c^\infty, a^\infty, \lambda^\infty, \eta^\infty) = \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \{ y(s_t) - c_t(s^t) + \phi_t(s^t) [u(c_t(s^t)) - v(a_t(s^t))] + \\ - \lambda_t(s^t) v'(a_t(s^t)) + [\eta_t(s^t) - \zeta_t(s^t)] u_c(c_t(s^t)) \} \Pi(s^t | s_0, a^{t-1}(s^{t-1}))$$

where

$$\phi_{t+1}(s^t, \hat{s}) = \phi_t(s^t) + \lambda_t(s^t) \frac{\pi_a(s_{t+1} = \hat{s} | s_t, a_t(s^t))}{\pi(s_{t+1} = \hat{s} | s_t, a_t(s^t))} \quad \forall \hat{s} \in S \quad \text{and} \quad \phi_0(s^0) = \gamma$$

$$\zeta_{t+1}(s^t, \hat{s}) = \eta_t(s^t) \quad \forall \hat{s} \in S \quad \text{and} \quad \zeta_0(s^0) = 0$$

Define

$$\hat{r}(a, c, \lambda, \eta, s, \phi, \zeta) = y(s) - c + \phi[u(c) - v(a)] - \lambda v'(a) + (\eta - \zeta) u_c(c)$$

Using the same arguments of Proposition 1, it is possible to show that the saddle-point functional equation associated with our problem:

$$\hat{J}(s, \phi, \zeta) = \min_{\lambda, \eta} \max_{a, c} \left\{ \hat{r}(a, c, \lambda, \eta, s, \phi, \zeta) + \beta \sum_{s'} \pi(s' | s, a) \hat{J}(s', \phi'(s'), \zeta') \right\} \\ \text{s.t.} \quad \phi'(s') = \phi + \lambda \frac{\pi_a(s' | s, a)}{\pi(s' | s, a)} \quad \forall s' \\ \zeta' = \eta$$

is a contraction in the state space $(s, \phi, \zeta)^7$.

The important thing to notice is that solving the Lagrangean problem with hidden savings is much simpler than solving the correspondin APS problem. While APS requires to first find the feasible expected utilities space, and then to run a dynamic programming algorithm to get a solution, in my approach I only have to solve for the Lagrangean first order conditions, taking into account that the problem is recursive in (s, ϕ, ζ) . Under both approaches, one needs to verify numerically that the solution is actually optimal for the original problem, but it should be obvious at this point that this verification is much faster under the Lagrangean approach than under APS. Indeed, the ex-post verification algorithm proposed by AP amounts to remaximize the agent's lifetime utility by choosing consumption, asset holdings and effort, taking as given the transfer policy function obtained by solving the Lagrangean problem. Given the transfer scheme and under the assumptions I made, the agent's problem is strictly concave, therefore first order conditions are necessary and sufficient to characterize the optimal solution.

⁷Since the proof is the same as for Proposition 1, it is omitted.

6 Numerical simulations: a new algorithm

In this section, I present some numerical examples of the repeated moral hazard model solved with the Lagrangean approach. The algorithm that I propose is much faster than the traditional dynamic programming tools, since it only has to solve the Lagrangean first order conditions; the advantage with respect to the traditional dynamic programming is due to the elimination of the maximization step. Moreover, compared to APS method, there is no need to calculate the expected utility space in which a solution is looked for, and this allows the researcher to solve models with a large number of state variables.

6.1 Two realizations, i.i.d. income, no hidden assets

I simplify the notation by writing a generic variable as x_t instead of $x_t(s^t)$. I assume that the income process has two possible realizations (y^L and y^H). I also assume there is no persistence across time (i.e., the state is i.i.d.), and I use the simpler notation $\pi(a_t) = \pi(y_{t+1} = y^L \mid a_t)$.

The Lagrangean becomes:

$$L = E_0^a \sum_{t=0}^{\infty} \beta^t \{ (y_t - c_t) + \phi_t [u(c_t) - v(a_t)] - \lambda_t v'(a_t) \}$$

with

$$\phi_{t+1}^H = \phi_t + \lambda_t \frac{\pi_a(a_t)}{\pi(a_t)}$$

$$\phi_{t+1}^L = \phi_t - \lambda_t \frac{\pi_a(a_t)}{1 - \pi(a_t)}$$

$$\phi_0(s^0) = \gamma$$

where E_t^a is the expectation operator over histories induced by the probability distribution $\pi(a_t)$. The first order conditions can be rewritten as

$$c_t : \quad u'(c_t) = \frac{1}{\phi_t} \tag{9}$$

$$\begin{aligned} a_t : \quad 0 = & -\lambda_t v''(a_t) - \phi_t v'(a_t) + \\ & + \pi_a(a_t) \beta E_{t+1}^a \left\{ \sum_{j=1}^{\infty} \beta^{j-1} \{ (y_{t+j} - c_{t+j}) - \lambda_{t+j} v'(a_{t+j}) + \phi_{t+j} [u(c_{t+j}) - v(a_{t+j})] \} \mid y_{t+1} = y^H \right\} + \\ & - \pi_a(a_t) \beta E_{t+1}^a \left\{ \sum_{j=1}^{\infty} \beta^{j-1} \{ (y_{t+j} - c_{t+j}) - \lambda_{t+j} v'(a_{t+j}) + \phi_{t+j} [u(c_{t+j}) - v(a_{t+j})] \} \mid y_{t+1} = y^L \right\} + \\ & + \beta \lambda_t \pi(a_t) \frac{\partial \left(\frac{\pi_a(a_t)}{\pi(a_t)} \right)}{\partial a_t} [u(c_{t+1}) - v(a_{t+1}) \mid y_{t+1} = y^H] + \\ & + \beta \lambda_t (1 - \pi(a_t)) \frac{\partial \left(\frac{-\pi_a(a_t)}{1 - \pi(a_t)} \right)}{\partial a_t} [u(c_{t+1}) - v(a_{t+1}) \mid y_{t+1} = y^L] \end{aligned} \tag{10}$$

and

$$\begin{aligned} \lambda_t : \quad 0 = & -v'(a_t) + \pi_a(a_t) \beta E_{t+1}^a \left\{ \sum_{j=1}^{\infty} \beta^{j-1} [u(c_{t+j}) - v(a_{t+j}) | y_{t+1} = y^H] \right\} + \\ & -\pi_a(a_t) \beta E_{t+1}^a \left\{ \sum_{j=1}^{\infty} \beta^{j-1} [u(c_{t+j}) - v(a_{t+j}) | y_{t+1} = y^L] \right\} \end{aligned} \quad (11)$$

6.2 The Algorithm

The numerical procedure is an application of the collocation method (see Judd (1998)) to solve the Lagrangean first order conditions. I use the Miranda-Fackler toolbox *Compecon* for function approximations.

I fix γ and I choose a discrete grid for ϕ_t that contains γ ; let a generic collocation grid node be indicated by (y, ϕ) . I approximate with cubic splines $a \equiv a(y, \phi)$ and $\lambda \equiv \lambda(y, \phi)$. I get consumption directly from ϕ by using (9): $c \equiv c(y, \phi) = u'^{-1}(\phi^{-1})$. Let $\phi^H = \phi + \lambda \frac{\pi_a(a)}{\pi(a)}$ and $\phi^L = \phi - \lambda \frac{\pi_a(a)}{1-\pi(a)}$. Notice that the expectations that enter in the FOCs are also representable as functions of the state variables. Therefore, (10) and (11) can be rewritten⁸ as:

$$\begin{aligned} 0 = & -\lambda v''(a) - \phi v'(a) + \\ & + \pi_a(a) \beta F(y^H, \phi^H) - \pi_a(a) \beta F(y^L, \phi^L) + \\ & + \beta \lambda \pi(a) \frac{\partial \left(\frac{\pi_a(a)}{\pi(a)} \right)}{\partial a} \left[u(c(y^H, \phi^H)) - v(a(y^H, \phi^H)) \right] + \\ & + \beta \lambda (1 - \pi(a)) \frac{\partial \left(\frac{-\pi_a(a)}{1-\pi(a)} \right)}{\partial a} \left[u(c(y^L, \phi^L)) - v(a(y^L, \phi^L)) \right] \end{aligned}$$

and

$$0 = -v'(a) + \pi_a(a) \beta D(y^H, \phi^H) - \pi_a(a) \beta D(y^L, \phi^L)$$

where

$$D(y, \phi) = u(c) - v(a) + \beta \left[\pi(a) D(y^L, \phi^L) + (1 - \pi(a)) D(y^H, \phi^H) \right]$$

$$F(y, \phi) = y - c - \lambda v'(a) + \phi [u(c) - v(a)] + \beta \left[\pi(a) F(y^L, \phi^L) + (1 - \pi(a)) F(y^H, \phi^H) \right]$$

I solve these four equations with a nonlinear equation solver.

⁸Notice that $D(\cdot)$ is the expected discounted utility of the agent, while $F(\cdot)$ is the expected discounted utility of the principal.

6.2.1 Results

I choose the following functional forms:

$$\begin{aligned} u(c) &= \frac{c^{1-\sigma}}{1-\sigma} \\ v(a) &= \alpha a^\varepsilon \\ \pi(a) &= 1 - a^\nu, \quad a \in (0, 1) \end{aligned}$$

The baseline parameters are summarized in the table:

α	ε	ν	σ	y^L	y^H	β	γ
0.5	2	0.5	2	0	1	0.95	0.5955

The algorithm delivers a set of parameterized policy functions. Figure 1 shows consumption, effort, the next period Pareto weights and the ICC Lagrange multiplier as functions of the current state ϕ . As we already said, consumption is increasing in ϕ , while effort is decreasing in the Pareto weight. Notice also that the policy functions for the Pareto weights satisfy Lemma 1. The Lagrange multiplier, interestingly, is an increasing function of the current state: as long as ϕ increases (i.e., as long as the realizations of high income is preponderant), the shadow cost of enforcing an incentive compatible allocation decreases.

Figure 2 represents the parameterized policy functions for transfers and expectations. Transfers are increasing in ϕ which is not surprising. Looking at expectations, notice that $F(y, \phi)$ corresponds to the planner expected value⁹, while $D(y, \phi)$ corresponds to the agent expected value. Under this interpretation, it is easy to see that the first has to be decreasing in ϕ while the second has to be increasing: if the planner cares more about the agent, her part of the surplus has to decrease. It is also interesting to notice that marginal disutility of effort decreases with current state, implying a lower marginal cost for the agent to exert more effort as long as the Pareto weight increases.

Figure 3 and 4 show the average allocations across 50 thousands independent simulations for 200 periods, with $y_0 = y^H$. Notice from Figure 3 that average consumption has a downward shape, as we were expecting from the discussion in Section 4.3. Effort, on the other hand, increases on average. As in Thomas and Worrall (1990), the average path for agent's lifetime utility is decreasing, while the Lagrange multiplier λ is reduced on average along the optimal path. Figure 4 has some interesting features to illustrate: the average ϕ does not show a monotone pattern; average transfers jump after period zero and then slightly decrease across time. The average principal's lifetime utility shows the time-inconsistency problem of the principal: after period zero, the value has a jump, showing that the principal would allow to reoptimize at time 1. To understand the last plot of Figure 4, let us notice that it is possible to reinterpret the value of the principal

⁹Notice from the definitions that $F(y(s), \phi) = J(s, \phi)$

$F(y_t, \phi_t)$ as indebttness of the agent. To see this, notice that we can define bond holdings recursively as:

$$\begin{aligned}
b_t(y_t, \phi_t) &= -E_t^a \sum_{j=0}^{\infty} \beta^j (y_{t+j} - c_{t+j}) = \\
&= - \left[E_t^a \sum_{j=0}^{\infty} \beta^j \{ (y_{t+j} - c_{t+j}) + \phi_{t+j} [u(c_{t+j}) - v(a_{t+j})] - \lambda_{t+j} v'(a_{t+j}) \} - \gamma V^{out} \right] \\
&= - [F(y_t, \phi_t) - \gamma V^{out}]
\end{aligned}$$

where the second line is due to the optimality of the contract and the definition of the Lagrangean, and the third to the definition of $F(y_t, \phi_t)$. It is easy to show (see Abraham and Pavoni (2006)) that this bond holdings sequence is only one of a continuum of possible equilibria of the model with moral hazard and monitorable asset markets, since we can always find a combination of transfers and bond holdings that deliver the optimal allocations of consumption and effort defined by the constrained efficient contract. Notice that from Figure 4 it is clear that the agent will like to decumulate assets.

Finally, Figure 5 shows the Pareto frontier of the optimal contract. It is clearly decreasing and strictly concave, implying that the critique by Messner and Pavoni (2004) is not relevant for the proposed algorithm.

6.3 An example with many income realizations

TBA

6.4 An example with hidden assets

TBA

7 Conclusions

I have presented a Lagrangean approach to repeated moral hazard problems, and an algorithm which is much faster than the traditional APS techniques. My methodology allows the researcher to deal with models with many states, and to calibrate the simulated series to real data in a reasonable amount of time.

This method has many possible applications. Given the speed, the algorithm can be useful (as a time-saving technique) also for those models that do not involve any endogenous state variable, like models of unemployment insurance *à la* Hopenhayn and Nicolini (1997). However, the main gain of the Lagrangean method can be seen in more complicated models. For example Zhao (2007) solves a model of risk sharing between two agents, where both have hidden effort. The next step would be to solve the same model with hidden effort and hidden savings, but this would be very complicated with APS. At the contrary, with the

Lagrangean approach, it will be much easier to characterize the optimal contract and get a numerical solution to illustrate the main properties.

There is a huge set of possible applications of the technique I developed: models with repeated moral hazard and hidden savings as in Abraham and Pavoni (2006), or models with repeated moral hazard and heterogeneous agents where the heterogeneity comes from some state variable, like wealth in models of income distribution (Cagetti and De Nardi (2006)) or productive capital (as in Meh and Quadrini (2006)). The price to pay is either to restrict the class of models we can deal with (by imposing assumptions on some primitives like Condition 1 and 2 in the present work), or to use some numerical tools to verify the optimality of the solution obtained with the first order Lagrangean approach. In any case, these are small costs compared to the benefit of being able to solve models that are at present unmanageable.

References

- [1] Abraham, A. and N. Pavoni (2006), "Efficient Allocations with Moral Hazard and Hidden Borrowing and Lending", mimeo, UCL
- [2] Abraham, A. and N. Pavoni (2008), "Principal-Agent Relationships with Hidden Borrowing and Lending: The First-Order Approach in Two Periods", mimeo, UCL
- [3] Abreu, D., Pearce, D. and E. Stacchetti (1990) "Toward a Theory of Discounted Repeated Games With Imperfect Monitoring," *Econometrica*, vol. 58(5), pp. 1041-1063.
- [4] Albanesi, S. and C. Sleet (2006), "Dynamic Optimal Taxation with Private Information", *Review of Economic Studies* 73: 1-30
- [5] Cagetti, M. and M. De Nardi (2006), "Estate Taxation, Entrepreneurship, and Wealth", mimeo
- [6] Cole, H. and N. Kocherlakota (1998), "Efficient allocations with hidden income and hidden storage", *Federal Reserve Bank of Minneapolis Staff Report* 238
- [7] Fernandes, A. and C. Phelan (2000), "A Recursive Formulation for Repeated Agency with History Dependence", *Journal of Economic Theory* 91(2): 223-247
- [8] Hopenhayn, H. A. and Nicolini, J. P. (1997), "Optimal Unemployment Insurance", *Journal of Political Economy* 105(2): 412-438
- [9] Laffont, J.-J. and D. Martimort (2002), "The Theory of Incentives", Princeton University Press
- [10] Ljungqvist, L. and T. Sargent (2004), "Recursive Macroeconomic Theory", MIT Press
- [11] Marcet, A. and R. Marimon (1999), "Recursive contracts", mimeo, Universitat Pompeu Fabra

- [12] Meh, C. and V. Quadrini (2006), "Endogenous market incompleteness with investment risks", *Journal of Economic Dynamics & Control*, 30:2143–2165
- [13] Messner, M. and N. Pavoni (2004). "On the Recursive Saddle Point Method: A Note", IGIER Working Paper n. 255
- [14] Phelan, C. and R. M. Townsend (1991), "Computing Multi-Period, Information Constrained Equilibria", *Review of Economic Studies* 58(5): 853-881
- [15] Rogerson, W. (1985a), "Repeated Moral Hazard", *Econometrica*, 53: 69-76
- [16] Rogerson, W. (1985b), "The First-Order Approach to Principal-Agent Problems", *Econometrica*, 53 (6): 1357-1368
- [17] Spear, S. and S. Srivastava (1987), "On Repeated Moral Hazard with Discounting", *Review of Economic Studies* 54(4): 599-617
- [18] Sleet, C. and S. Yeltekin (2005), "Social credibility, social patience and long run inequality", mimeo, Carnegie Mellon University
- [19] Thomas, J. and T. Worrall (1990) "Income fluctuations and asymmetric information: An example of a repeated principal-agent problem", *Journal of Economic Theory* 51: 367 390
- [20] Werning, I. (2001), "Repeated Moral-Hazard with Unmonitored Wealth: A Recursive First-Order Approach", mimeo, MIT
- [21] Zhao, R. (2007), "Dynamic risk-sharing with two-sided moral hazard", *Journal of Economic Theory* 136: 601-640.

Proofs

Proposition 1 Fix an arbitrary constant $K > 0$ and let $K_\phi = \max \{K, K \|\phi\|\}$. The operator

$$(T_K f)(s, \phi) \equiv \min_{\{\lambda > 0: \|\lambda\| \leq K_\phi\}} \max_{a, c} \left\{ r(a, c, \lambda, s, \phi) + \beta \sum_{s'} \pi(s' | s, a) f(s', \phi'(s')) \right\}$$

$$s.t. \quad \phi'(s') = \phi + \lambda \frac{\pi_a(s' | s, a)}{\pi(s' | s, a)} \quad \forall s'$$

is a contraction..

Proof. I have to prove that the above operator is a contraction in some complete metric space. Let $\varsigma \equiv (a, c)$. Following Marcet and Marimon (1998), it is possible to decompose the objective function of the SPFE as:

$$\begin{aligned} r(\varsigma, \lambda, s, \phi) &\equiv r_0(\varsigma, s) + \lambda r_1(\varsigma, s) + \phi r_2(\varsigma, s) \\ &= y(s) - c - \lambda v'(a) + \phi [u(c) - v(a)] \end{aligned}$$

and consequently, the value of the SPFE in two parts:

$$f(s, \phi) \equiv f_0(s, \phi) + f_1(s, \phi)$$

where $f_0(s, \cdot)$ is homogeneous of degree zero and $f_1(s, \cdot)$ is homogeneous of degree 1. Therefore, the space

$$M = \{f : S \times \mathbb{R}_+ \quad s.t.$$

$$\begin{aligned} a) \quad & f(\cdot, \cdot) = f_0(\cdot, \cdot) + f_1(\cdot, \cdot) \quad \text{and} \quad \forall \alpha > 0 \\ & f_0(\cdot, \alpha \phi) = f_0(\cdot, \phi) \quad \text{and} \quad f_1(\cdot, \alpha \phi) = \alpha f_1(\cdot, \phi) \\ b) \quad & f_j(s, \cdot) \text{ is continuous and bounded,} \quad j = 0, 1 \end{aligned}$$

will be our candidate, with norm

$$\|f\| = \sup \{|f(s, \phi)| : \|\phi\| \leq 1, s \in S\}$$

Marcet and Marimon (1998), Lemma 1 shows that M is a nonempty complete metric space. Now, fix a positive constant K and let $K_\phi = \max \{K, K \|\phi\|\}$. Define the auxiliary operator

$$(T_K f)(s, \phi) = \min_{\{\lambda \geq 0: \|\lambda\| \leq K_\phi\}} \max_{\varsigma} \left\{ r(\varsigma, \lambda, s, \phi) + \beta \sum_{s'} \pi(s' | s, a) f(s', \phi'(s')) \right\}$$

$$s.t. \quad \phi'(s') = \phi + \lambda \frac{\pi_a(s' | s, a)}{\pi(s' | s, a)}$$

I have to show that $T_K : M \longrightarrow M$. Notice that

$$\begin{aligned} (T_K f)(s, \phi) &= r_0(\varsigma^*, s) + \beta \sum_{s'} \pi(s' | s, a^*) f_0(s', \phi^{*'}(s')) + \\ &\quad + \lambda^* r_1(\varsigma^*, s) + \phi r_2(\varsigma^*, s) + \beta \sum_{s'} \pi(s' | s, a^*) f_1(s', \phi^{*'}(s')) \end{aligned}$$

hence by Schwartz's inequality

$$\begin{aligned} \|(T_K f)(s, \phi)\| &\leq \|r_0(\varsigma^*, s)\| + \beta \left\| f_0\left(s', \phi^{*'}(s')\right) \right\| + \\ &\quad + \max\{K, K\|\phi\|\} \|r_1(\varsigma^*, s)\| + \|\phi\| \|r_2(\varsigma^*, s)\| + \\ &\quad + \beta \left(\max\{K, K\|\phi\|\} \left\| \frac{\pi_a(s' | s, a^*)}{\pi(s' | s, a^*)} \right\| + \|\phi\| \right) \left\| f_1\left(s', \frac{\phi^{*'}(s')}{\|\phi^{*'}(s')\|} \right) \right\| \end{aligned}$$

and therefore $(T_K f)(s, \phi)$ is bounded. A generalized Maximum Principle argument gives continuity of $(T_K f)(s, \phi)$. To check for homogeneity properties, let $(\varsigma_\alpha^*, \lambda_\alpha^*)$ be a solution associated to $(s, \alpha\phi)$, and let $\tilde{\lambda} = \frac{\lambda_\alpha^*}{\alpha}$ and $\phi'_\alpha(s') = \alpha\phi + \lambda_\alpha^* \frac{\pi_a(s' | s, a_\alpha^*)}{\pi(s' | s, a_\alpha^*)}$. By the properties of homogeneous function, if

$$(\varsigma_\alpha^*, \lambda_\alpha^*) \in \arg \min_{\{\lambda \geq 0: \|\lambda\| \leq K_\phi\}} \max_{\varsigma} \left\{ r(\varsigma, \lambda, s, \alpha\phi) + \beta \sum_{s'} \pi(s' | s, a_\alpha^*) f(s', \phi'_\alpha(s')) \right\}$$

then

$$(\varsigma_\alpha^*, \tilde{\lambda}) \in \arg \min_{\{\lambda \geq 0: \|\lambda\| \leq K_\phi\}} \max_{\varsigma} \left\{ r(\varsigma, \lambda, s, \phi) + \beta \sum_{s'} \pi(s' | s, a_\alpha^*) f(s', \phi'(s')) \right\}$$

Hence

$$\begin{aligned} (Tf)(s, \alpha\phi) &= (Tf)_0(s, \alpha\phi) + (Tf)_1(s, \alpha\phi) = \\ &= r_0(\varsigma_\alpha^*, s) + \beta \sum_{s'} \pi(s' | s, a_\alpha^*) f_0\left(s', \alpha\phi + \lambda_\alpha^* \frac{\pi_a(s' | s, a_\alpha^*)}{\pi(s' | s, a_\alpha^*)}\right) + \\ &\quad + \lambda_\alpha^* r_1(\varsigma_\alpha^*, s) + \alpha\phi r_2(\varsigma_\alpha^*, s) + \beta \sum_{s'} \pi(s' | s, a_\alpha^*) f_1\left(s', \alpha\phi + \lambda_\alpha^* \frac{\pi_a(s' | s, a_\alpha^*)}{\pi(s' | s, a_\alpha^*)}\right) \\ &= r_0(\varsigma_\alpha^*, s) + \beta \sum_{s'} \pi(s' | s, a_\alpha^*) f_0\left(s', \phi + \tilde{\lambda} \frac{\pi_a(s' | s, a_\alpha^*)}{\pi(s' | s, a_\alpha^*)}\right) + \\ &\quad + \alpha \left[\tilde{\lambda} r_1(\varsigma_\alpha^*, s) + \phi r_2(\varsigma_\alpha^*, s) + \beta \sum_{s'} \pi(s' | s, a_\alpha^*) f_1\left(s', \phi + \tilde{\lambda} \frac{\pi_a(s' | s, a_\alpha^*)}{\pi(s' | s, a_\alpha^*)}\right) \right] \\ &= (Tf)_0(s, \phi) + \alpha (Tf)_1(s, \phi) \end{aligned}$$

and therefore the operator preserves the homogeneity properties. To see monotonicity, let $g, h \in M$ such that $g \leq h$. Therefore

$$\begin{aligned} &\max_{\varsigma} \left\{ r(\varsigma, \lambda, s, \phi) + \beta \sum_{s'} \pi(s' | s, a) g(s', \phi'(s')) \right\} \\ &\leq \max_{\varsigma} \left\{ r(\varsigma, \lambda, s, \phi) + \beta \sum_{s'} \pi(s' | s, a) h(s', \phi'(s')) \right\} \end{aligned}$$

and then

$$\begin{aligned} &\min_{\{\lambda \geq 0: \|\lambda\| \leq K_\phi\}} \max_{\varsigma} \left\{ r(\varsigma, \lambda, s, \phi) + \beta \sum_{s'} \pi(s' | s, a) g(s', \phi'(s')) \right\} \\ &\leq \min_{\{\lambda \geq 0: \|\lambda\| \leq K_\phi\}} \max_{\varsigma} \left\{ r(\varsigma, \lambda, s, \phi) + \beta \sum_{s'} \pi(s' | s, a) h(s', \phi'(s')) \right\} \end{aligned}$$

which implies $(T_K g)(s, \phi) \leq (T_K h)(s, \phi)$. To see discounting, let $k \in \mathbb{R}_+$, and define $f + k \in M$ as $(f + k)(s, \phi) = f(s, \phi) + k$. Therefore:

$$\begin{aligned} & \max_{\varsigma} \left\{ r(\varsigma, \lambda, s, \phi) + \beta \sum_{s'} \pi(s' | s, a) (g + k)(s', \phi'(s')) \right\} \\ &= \max_{\varsigma} \left\{ r(\varsigma, \lambda, s, \phi) + \beta \sum_{s'} \pi(s' | s, a) g(s', \phi'(s')) + \beta k \right\} \\ &= \max_{\varsigma} \left\{ r(\varsigma, \lambda, s, \phi) + \beta \sum_{s'} \pi(s' | s, a) g(s', \phi'(s')) \right\} + \beta k \end{aligned}$$

Hence we get

$$\begin{aligned} T_K(f + k)(s, \phi) &= \\ &= \min_{\{\lambda \geq 0: \|\lambda\| \leq K_\phi\}} \max_{\varsigma} \left\{ r(\varsigma, \lambda, s, \phi) + \beta \sum_{s'} \pi(s' | s, a) (f + k)(s', \phi'(s')) \right\} \\ &= \min_{\{\lambda \geq 0: \|\lambda\| \leq K_\phi\}} \max_{\varsigma} \left\{ r(\varsigma, \lambda, s, \phi) + \beta \sum_{s'} \pi(s' | s, a) g(s', \phi'(s')) \right\} + \beta k \\ &= (T_K f)(s, \phi) + \beta k \end{aligned}$$

and then $T_K(f + k) \leq T_K f + \beta k$. Now it is possible to use the above properties to show the contraction property for the operator T_K . In order to see this, let $f, g \in M$. By homogeneity, we get

$$\begin{aligned} f(s, \phi) &= g(s, \phi) + f(s, \phi) - g(s, \phi) \\ &\leq g(s, \phi) + |f(s, \phi) - g(s, \phi)| \end{aligned}$$

and then

$$f(s, \phi) \leq g(s, \phi) + \|f(s, \phi) - g(s, \phi)\|$$

Now applying the operator T_K and using monotonicity and discounting we get:

$$\begin{aligned} (T_K f)(s, \phi) &\leq T_K(g + \|f - g\|)(s, \phi) \\ &\leq (T_K g)(s, \phi) + \beta \|f - g\| \end{aligned}$$

which implies finally

$$\|T_K f - T_K g\| \leq \beta \|f - g\|$$

and given $\beta \in (0, 1)$ this concludes the proof that the operator T_K is a contraction. ■

Lemma 1 *In the optimal contract, $\phi_{t+1}(s^t, \hat{s}_1) < \phi_t(s^t) < \phi_{t+1}(s^t, \hat{s}_I)$ for any t .*

Proof. Notice first that, for any t , $\exists i, j : \pi_a(\hat{s}_i | s_t, a_t^*(s^t)) > 0$ and $\pi_a(\hat{s}_j | s_t, a_t(s^t)) < 0$. Suppose not: then the only possibility is that $\pi_a(\hat{s}_i | s_t, a_t(s^t)) = 0$ for any i (otherwise, $\sum_{\hat{s}_i} \pi_a(\hat{s}_i | s_t, a_t(s^t)) \neq 0$, which is impossible). This implies, by (7), $0 = v'(a_t(s^t))$ which is a contradiction since $v(\cdot)$ is strictly increasing.

Adding the full support assumption and the fact that $\lambda_t(s^t) > 0$, we get that $\exists i, j : \phi_{t+1}(s^t, \hat{s}_j) < \phi_t(s^t) < \phi_{t+1}(s^t, \hat{s}_i)$. By MLRC, $\phi_{t+1}(s^t, \hat{s}_1) \leq \phi_{t+1}(s^t, \hat{s}_j)$ for any j and $\phi_{t+1}(s^t, \hat{s}_i) \leq \phi_{t+1}(s^t, \hat{s}_I)$ for any i , which proves the statement. ■

Proposition 2 $\phi_t(s^t)$ is a martingale that converges to zero.

Proof. Use the law of motion of $\phi_t(s^t)$ and take expectations on both sides:

$$\sum_{s_{t+1}} \phi_{t+1}(s^t, s_{t+1}) \pi(s_{t+1} | s_t, a_t(s^t)) = \phi_t(s^t) + \lambda_t(s^t) \sum_{s_{t+1}} \frac{\pi_a(s_{t+1} | s_t, a_t(s^t))}{\pi(s_{t+1} | s_t, a_t(s^t))} \pi(s_{t+1} | s_t, a_t(s^t))$$

Notice that $\lambda_t(s^t) \sum_{s_{t+1}} \frac{\pi_a(s_{t+1} | s_t, a_t(s^t))}{\pi(s_{t+1} | s_t, a_t(s^t))} \pi(s_{t+1} | s_t, a_t(s^t)) = 0$, which implies

$$E_t^a[\phi_{t+1} | s^t] = \phi_t(s^t) \quad (12)$$

where $E_t^a[\cdot]$ is the expectation operator induced by $a_t(s^t)$. Therefore $\phi_t(s^t)$ is a martingale. To see that it converges to zero, rewrite (12) by using (6):

$$E_t^a \left[\frac{1}{u_c(c_{t+1}(s^{t+1}))} \right] = \frac{1}{u_c(c_t(s^t))}$$

By Inada conditions, $\frac{1}{u_c(c_t(s^t))}$ is bounded above zero and below infinity. Therefore $\phi_t(s^t)$ is a nonnegative martingale, and by Doob's theorem it converges almost surely to a random variable (call it X). To see that $X = 0$ almost surely, I follow the proof strategy of Thomas and Worrall (1990), to which I refer for details. Suppose not, and take a path $\{s^t\}_{t=0}^\infty$ such that $\lim_{t \rightarrow \infty} \phi_t(s^t) = \bar{\phi} > 0$ and state $\hat{s}_I$ happens infinitely many times. I claim that this sequence cannot exist. Take a subsequence $\{s^{t(k)}\}_{k=1}^\infty$ of $\{s^t\}_{t=0}^\infty$ such that $s_{t(k)} = \hat{s}_I \forall k$. This subsequence has to converge to some limit $\bar{\phi} > 0$, since at some point will be in a ϵ -neighborhood of $\bar{\phi}$ for some $\epsilon > 0$. Call $f(\phi_t(s^t), \hat{s}_i) = \phi_{t+1}(s^t, \hat{s}_i)$ and notice that $f(\cdot)$ is continuous, hence $\lim_{k \rightarrow \infty} f(\phi_{t(k)}(s^{t(k)}), \hat{s}_I) = f(\bar{\phi}, \hat{s}_I)$. By definition, $f(\phi_{t(k)}(s^{t(k)}), \hat{s}_I) = \phi_{t(k)+1}(s^t, \hat{s}_I)$, then $\lim_{k \rightarrow \infty} \phi_{t(k)+1}(s^{t(k)}, \hat{s}_I) = f(\bar{\phi}, \hat{s}_I)$. However, notice that it must be $\lim_{k \rightarrow \infty} \phi_{t(k)}(s^{t(k)}) = \bar{\phi}$ and $\lim_{k \rightarrow \infty} \phi_{t(k)+1}(s^{t(k)}, \hat{s}_I) = \bar{\phi}$. But by Lemma 1, $\phi_{t(k)}(s^{t(k)}) < \phi_{t(k)+1}(s^{t(k)}, \hat{s}_I)$ for any k . Therefore, this is a contradiction and this sequence cannot exist. Since paths where state $\hat{s}_I$ occurs only a finite number of times have probability zero, this implies that

$$\Pr \left\{ \lim_{t \rightarrow \infty} \phi_t(s^t) > 0 \right\} = 0$$

which implies $X = 0$ almost surely. ■

Figures




Figure 1: Policy functions: consumption, effort, Pareto weights and ICC Lagrange multiplier




Figure 2: Policy functions: transfers, expectations and marginal disutility of effort




Figure 3: Time series properties: Average allocations over 50000 independent simulations (consumption, effort, agent lifetime utility, Lagrange multiplier)




Figure 4: Time series properties. Average allocations over 50000 independent simulations (Pareto weight, transfer, planner value, asset holdings)




Figure 5: Pareto frontier