



Munich Personal RePEc Archive

Repeated moral hazard and recursive Lagrangeans

Antonio Mele

University of Oxford, Department of Economics, Nuffield College,
Oxford

29. March 2010

Online at <http://mpa.ub.uni-muenchen.de/21741/>

MPRA Paper No. 21741, posted 31. March 2010 05:52 UTC

Repeated Moral Hazard and Recursive Lagrangeans

Antonio Mele*

University of Oxford and Nuffield College

March 29, 2010

Abstract

This paper shows how to solve dynamic agency models by extending recursive Lagrangean techniques *à la* Marcet and Marimon (2009) to problems with hidden actions. The method has many advantages with respect to promised utilities approach (Abreu, Pearce and Stacchetti (1990)): it is a significant improvement in terms of simplicity, tractability and computational speed. Solutions can be easily computed for hidden actions models with several endogenous state variables and several agents, while the promised utilities approach becomes extremely difficult and computationally intensive even with just one state variable or two agents. Several numerical examples illustrate how this methodology outperforms the standard approach.

1 Introduction

This paper shows how to solve repeated moral hazard models using recursive Lagrangean techniques. In particular, this approach can be used in the analysis of dy-

***Acknowledgements:** I am grateful to Albert Marcet for his 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, to Davide Debortoli and Ricardo Nunes for their generosity in discussing infinitely many numerical aspects of the paper, and to Chris Sleet for pointing out a mistake in a previous version of the work. This paper has benefitted from comments by Klaus Adam, Sofia Bauducco, Toni Braun, Filippo Brutti, Andrea Caggese, Francesco Caprioli, Martina Cecioni, Federica Dispenza, Josè Dorich, Martin Ellison, Giuseppe Ferrero, Harald Fadinger, Tom Holden, Michal Horvath, Tom Krebs, Eva Luethi, Angelo Mele, Matthias Messner, Krisztina Molnar, Juan Pablo Nicolini, Nicola Pavoni, Josep Pijoan-Mas, Michael Reiter, Pontus Rendahl, Gilles Saint-Paul, Daniel Samano, Antonella Tutino and from participants at Macro Break and Macro Discussion Group at Universitat Pompeu Fabra, Macro Workshop at University of Oxford, SED Meeting 2008 in Cambridge (MA), Midwest Economic Theory Meeting 2008 in Urbana-Champaign, 63rd European Meeting of the Econometric Society 2008 in Milan, 14th CEF Conference 2008 in Paris, 7th Workshop on "Macroeconomic Dynamics: Theory and Applications" in Rome, North American Summer Meeting of the Econometric Society 2009 in Boston, and seminar audience at University of Mannheim, Paris School of Economics, Queen Mary - University of London, University of Oxford, Nuffield College and Federal Reserve Board. This paper was awarded the prize of the 2008 CEF Student Contest by the Society for Computational Economics. All mistakes are mine. **email:** antonio.mele@economics.ox.ac.uk

dynamic hidden-actions models with several endogenous state variables and many agents. While these models are extremely complicated to solve with commonly used solution strategies, my methodology is simpler and numerically faster than the alternatives.

The recent literature on dynamic principal-agent models is vast¹. Typically these models do not have closed form solution, therefore it is necessary to solve them numerically. The main technical difficulty is the history dependence of the optimal allocation: the principal must keep track of the whole history of shock realizations, use it to extract information about the agent's unobservable behavior, and reward or punish the agent accordingly. As a consequence, it is not possible to derive a standard recursive representation of the principal's intertemporal maximization problem. The traditional way of dealing with this complication is based on the *promised utilities approach*: the dynamic program is transformed into an auxiliary problem with the same solution, in which the principal chooses allocations and the agent's future continuation value, taking as given the continuation value chosen in the previous period. The latter (also called *promised utility*) incorporates the whole history of the game, and hence becomes a new endogenous state variable to be chosen optimally. By using a standard argument, due to Abreu, Pearce and Stacchetti (1990) (APS henceforth) among others, it can be shown that the auxiliary problem has a recursive representation in a new state space that includes the continuation value and the state variables of the original problem. However, there is an additional complication: in order for the auxiliary problem to be equivalent to the original one, promised utilities must belong to a particular set (call it the feasible set), which has to be characterized numerically before computation of the optimal allocation². It is trivial to characterize this set if there is just one exogenous shock, but it becomes complicated, if not computationally unfeasible, in models with several endogenous states or with many agents. Therefore, with this approach, there is a large class of models that we cannot analyze even with numerical methods.

This paper provides a way to overcome the limits of the promised utilities approach: under assumptions that justify the use of the first-order approach³, it extends the recur-

¹Many contributions have focused 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 (2008, 2009)); other works have explored what happens in presence of more than one agent (see e.g. Zhao (2007) and Friedman (1998)), while few researchers have extended the setup to production economies with capital (Clementi et al. (2008a,2008b)). Among applications, a non-exhaustive list includes unemployment insurance (Hopenhayn and Nicolini (1997), Shimer and Werning (forthcoming), Werning (2002), Pavoni (2007, forthcoming)), executive compensation (Clementi et al. (2008a,2008b), Clementi et al. (2006), Atkeson and Cole (2008)), entrepreneurship (Quadrini (2004), Paulson et al. (2006)), credit markets (Lehnert et al. (1999), and many more.

²The feasible set is the fixed point of a set-operator (see APS for details). The standard numerical algorithm proposed by APS starts with a large initial set, and iteratively converges to the fixed point. Sleet and Yeltekin (2003) and Judd, Conklin and Yeltekin (2003) provide two efficient ways of computing it.

³The first-order approach, consisting of the substitution of the incentive-compatibility constraint with the first-order conditions of the agent's maximization problem with respect to hidden actions, is

sive Lagrangean techniques developed in Marcet and Marimon (2009) (MM henceforth) to the dynamic agency model. These techniques are well understood and widely used for full information problems of optimal policy and enforcement frictions, but MM do not analyze their applicability to environments with private information. Sleet and Yeltekin (2008a) make a crucial contribution in applying recursive Lagrangean techniques to dynamic models with privately observed idiosyncratic preference shocks. This paper instead focuses on a particular class of dynamic models with hidden actions, i.e. models that admit the use of the first-order approach⁴.

The approach can be better illustrated in a dynamic principal-agent model such as the one in Spear and Srivastava (1987), where no endogenous state variables are present. The recursive Lagrangean formulation of this model has a straightforward interpretation: the optimal contract can be characterized by maximizing a weighted sum of the lifetime utilities of the principal and the agent (i.e., a utilitarian social welfare function), where in each period the social planner optimally updates the weight of the agent in order to enforce an incentive compatible allocation. These Pareto-Negishi weights⁵ become the new state variables that "recursify" the dynamic agency problem. In particular, this endogenously evolving weight summarizes the contract's promises according to which the agent is rewarded or punished. Imagine, for simplicity, that there are only two possible realizations for output, either "good" or "bad". The contract promises that, if tomorrow a "good" realization of the output is observed, the Pareto-Negishi weight will increase, therefore the principal will care more about the expected discounted utility of the agent from tomorrow on. Analogously, if a "bad" outcome happens, the Pareto-Negishi weight will decrease, hence the principal will care less about the expected discounted utility of the agent from tomorrow on. An optimal contract chooses the sequence of Pareto-Negishi weights in such a way that rewards and punishments are incentive compatible.

Under this interpretation, it is easy to understand why the recursive Lagrangean approach is simpler than APS: it does not require the additional step of characterizing

widely used in the solution of static models with moral hazard since the seminal work of Mirrlees (1975). Unfortunately, as Mirrlees pointed out, this approach is not justified in all setups. The literature has provided several sets of assumptions that guarantee its validity.

⁴This paper is different from Sleet and Yeltekin (2008a) in two aspects, besides the focus on a different type of private information. Firstly, the structure of the hidden shocks framework is such that Sleet and Yeltekin (2008a) can use recursive Lagrangeans directly on the original problem without need of a first-order approach. Secondly, they mainly focus on theoretical aspects of the method, while this paper also aims at providing an efficient way of characterizing the numerical solution. A third and minor difference is technical: they do not exploit the homogeneity of the value and policy functions, which is crucial in my proof strategy and in numerical applications. Their work is complementary to this paper in the analysis of dynamic models with asymmetric information. They also use their techniques in several applied papers, for example Sleet and Yeltekin (2008b) and Sleet and Yeltekin (2006).

⁵Lustig and Chien (2005) use the term "Pareto-Negishi weight" in a model of an endowment economy with limited enforcement, where agents face both aggregate and idiosyncratic shocks. In their work, the weight of each agent evolves stochastically in order to keep track of occasionally binding enforcement constraints. Sleet and Yeltekin, in their papers, use the same terminology.

a feasible set for the new state variables, as we did with APS for continuation values. In the recursive Lagrangean approach, the social welfare function maximization problem is well defined for any real-valued weight⁶.

This line of reasoning can be easily extended to more general problems of repeated moral hazard with many agents and many observable endogenous state variables. The dynamic optimization problem has a recursive formulation based on Pareto-Negishi weights and the endogenous state variables. These weights are updated in each period to enforce an incentive compatible allocation, while the endogenous states follow their own law of motion. Also in these more complicated environments there is no need for characterizing the feasible set of Pareto-Negishi weights. Given this, the main gain in using recursive Lagrangeans is in terms of tractability, since we eliminate the often intractable step of characterizing feasible values for the auxiliary problem, a crucial aspect of the APS approach.

Extending the recursive Lagrangean approach to models with endogenous unobservable state variables is more challenging. In particular, it is well known that the first-order approach is rarely justified in these cases, and we do not have sufficient conditions that guarantee its validity. However, we can follow a "solve-and-verify" approach along the lines of Abraham and Pavoni (2009): first solve the problem with recursive Lagrangeans, using the first-order approach⁷, and then verify that the agent does not have incentives to deviate from the choices implied by the optimal contract. The last verification step can be done with standard dynamic programming techniques, as Abraham and Pavoni suggested in their work.

This paper also propose an efficient way to compute the optimal contract based on the theoretical results. The idea is to find approximated policy functions by solving Lagrangean first-order conditions. The procedure is an application of the collocation method (see Judd (1998)). The algorithm is simple: firstly, approximate the policy functions for allocations, Lagrange multipliers, agents' and principal's continuation values over a set of grid nodes, with standard interpolation techniques, either splines or Chebychev polynomials depending on the particular application. Then look for the coefficients of these approximated policy functions that satisfy Lagrangean first-order conditions. The gain in terms of computational speed is large: as a benchmark, in a state-of-the-art laptop, the Fortran code provided by Abraham and Pavoni (2009) solves a model with hidden effort and hidden assets accumulation in 15 hours, while my Matlab code obtains an accurate solution in around 20 seconds. This large computational gain is obtained for two reasons. The first has already been mentioned: we do not need to find a feasible set for Pareto-Negishi weights. The second reason is that solving a system of nonlinear equations is much faster than value function iteration (the standard algorithm used for promised utility approach)⁸.

⁶This is also valid for the recursive Lagrangeans approach in dynamic optimization problems with full information. For a discussion of this issue, see Marcet and Marimon (2009).

⁷Notice that we need to use the agent's first-order conditions with respect to all unobservable choice variables.

⁸The proposed procedure is a local characterization of the saddle-point, and therefore second-order

The paper is organized as follows: section 2 provides an illustration of the recursive Lagrangean approach in a simple dynamic principal-agent model; section 3 contains a more general theorem for problems with several endogenous state variables and more than one agent; section 4 discusses how the recursive Lagrangean approach can be used in models with unobservable endogenous states and illustrates these ideas in a model with hidden effort and hidden asset accumulation; section 5 explains the details of the algorithm, and provides some numerical examples and a performance analysis of the algorithm in terms of accuracy and computational speed; section 6 concludes.

2 An illustration with a simple dynamic agency model

In order to illustrate the Lagrangean approach, it is easier to start with a dynamic agency problem without endogenous states, as in Spear and Srivastava (1987). This is helpful in understanding the differences between this approach and the promised utility method.

The economy is inhabited by a risk neutral principal and a risk averse agent. Time is discrete, and the state of the world follows an observable Markov process $\{s_t\}_{t=0}^{\infty}$, where $s_t \in S$, and $\text{card}(S) = I$. The realizations of the process are public information. Denote the single realizations with subscripts, and the histories with superscripts:

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

In 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 $a_t(s^t) \in A \subseteq \mathbb{R}_+$, and A is bounded. I will refer to $a_t(s^t)$ as action or effort.

The costly action affects the future probability distribution of the state of the world. For simplicity, let $\widehat{s}_i, i = 1, 2, \dots, I$ be the possible realizations of $\{s_t\}$ and let them be ordered such that $y(s_t = \widehat{s}_1) < y(s_t = \widehat{s}_2) < \dots < y(s_t = \widehat{s}_I)$. Let $\pi(s_{t+1} = \widehat{s}_i | s_t, a_t(s^t))$ be the probability that state tomorrow is $\widehat{s}_i \in S$ conditional on past state and effort exerted by the agent at the beginning of the period⁹, with $\pi(s_0 = \widehat{s}_I) = 1$. Assume $\pi(\cdot)$ is twice continuously differentiable in $a_t(s^t)$ with $\frac{\pi_a(\cdot)}{\pi(\cdot)}$ bounded, and has *full support*: $\pi(s_{t+1} = \widehat{s}_i | s_t, a) > 0 \forall i, \forall a, \forall s_t$. 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 probability of history s^{t+1} induced by the history of unobserved actions $a^t(s^t) \equiv (a_0(s^0), a_1(s^1), \dots, a_t(s^t))$.

The instantaneous utility of the agents is

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

conditions can be an issue. The researcher can control for this problem by starting from different initial conditions and checking if the algorithm always converges to the same solution. All examples presented in my paper are robust to this check.

⁹Notice that shocks can be persistent. In the numerical examples, the focus is on i.i.d. shocks, but it should be clear that persistence neither creates particular theoretical nor numerical problems.

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. The instantaneous utility is uniformly bounded. The agent does not accumulate assets autonomously: the only source of insurance is the principal. The budget constraint of the agent will be simply

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

Both principal and agent are fully committed once they sign the contract at time zero.

A feasible *contract* (or *allocation*) \mathcal{W} in this framework is a plan $(a^\infty, c^\infty, \tau^\infty) \equiv \{a_t(s^t), c_t(s^t), \tau_t(s^t) \mid \forall s^t \in S^{t+1}\}_{t=0}^\infty$ that belongs to the following set:

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

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

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

therefore the principal can characterize efficient contracts by 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:

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

$$\sum_{t=0}^{\infty} \sum_{s^t} \beta^t [u(c_t(s^t)) - v(a_t(s^t))] \Pi(s^t \mid s_0, a^{t-1}(s^{t-1})) \geq V^{out}. \quad (3)$$

Call this *the original problem*. Notice that the sequence of effort choices in (2) is the optimal solution of the agent's maximization problem, given the contract offered by the principal. If the agent's optimization problem is well-behaved, this sequence can be characterized by the first-order conditions of the agent's optimization problem. In that case, it is possible to use the agent's first-order conditions as constraints in the principal's dynamic problem. This solution strategy is commonly known in the

literature as the *first-order approach*. For this simple setup, there are well known conditions in the literature that guarantee the validity of the first-order approach, i.e. that guarantee that the problem with first-order conditions is equivalent to the original problem and therefore delivers the same solution. In the rest of this section assume that Rogerson (1985b) conditions of monotone likelihood ratio (MLRC) and convexity of the distribution (CDFC) are satisfied. These conditions are sufficient to guarantee the validity of the first-order approach in this simple setup¹⁰.

If the first-order approach is justified, the agent's first order conditions with respect to effort can be substituted into the principal's problem. The agent, given the principal's strategy profile $\tau^\infty \equiv \{\tau_t(s^t)\}_{t=0}^\infty$, solves

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

The first order condition for effort is

$$\begin{aligned} v'(a_t(s^t)) &= \sum_{j=1}^\infty \beta^j \sum_{s^{t+j}|s^t} \pi_a(s_{t+1} | s_t, a_t(s^t)) \times \\ &\times [u(c_{t+j}(s^{t+j})) - v(a_{t+j}(s^{t+j}))] \Pi(s^{t+j} | s^{t+1}, a^{t+j}(s^{t+j} | s^{t+1})). \end{aligned} \quad (4)$$

Intuitively, the marginal cost of effort today (LHS) has to be equal to future expected benefits (RHS) in terms of expected future utility. The use of (4) is crucial, since it allows to write the Lagrangean of the principal's problem. In the following, for simplicity I refer to (4) as the *incentive-compatibility constraint* (ICC).

Rewrite the Pareto problem of the principal as

$$\begin{aligned} W(s_0) &= \max_{\{a_t(s^t), c_t(s^t)\}_{t=0}^\infty \in \Gamma^{MH}} \sum_{t=0}^\infty \sum_{s^t} \beta^t [y(s_t) - c_t(s^t)] \Pi(s^t | s_0, a^{t-1}(s^{t-1})) \\ s.t. \quad v'(a_t(s^t)) &= \sum_{j=1}^\infty \beta^j \sum_{s^{t+j}|s^t} \frac{\pi_a(s_{t+1} | s_t, a_t(s^t))}{\pi(s_{t+1} | s_t, a_t(s^t))} \times \\ &\times [u(c_{t+j}(s^{t+j})) - v(a_{t+j}(s^{t+j}))] \Pi(s^{t+j} | s^t, a^{t+j-1}(s^{t+j-1} | s^t)) \\ &\quad \forall s^t, t \geq 0 \\ &\sum_{t=0}^\infty \sum_{s^t} \beta^t [u(c_t(s^t)) - v(a_t(s^t))] \Pi(s^t | s_0, a^{t-1}(s^{t-1})) \geq V^{out}. \end{aligned} \quad (5)$$

¹⁰For static problems, Jewitt (1988) provides another set of sufficient conditions, which can be used in alternative to Rogerson's to guarantee the feasibility of a first-order approach. Notice that both Rogerson's and Jewitt's conditions are sufficient for dynamic agency setups with observable endogenous states. He (2010) suggests a fixed-point condition that justifies the first-order approach in static environments, which can potentially also be used in dynamic settings.

2.1 The Lagrangean approach

It is trivial to show that (3) must be binding in the optimum. Given this consideration, Problem (5) can be seen as the constrained maximization of a social welfare function, where the Pareto weight for the principal and the agent are, respectively, 1 and γ :

$$\begin{aligned}
W^{SWF}(s_0) = & \max_{\{a_t(s^t), c_t(s^t)\}_{t=0}^{\infty} \in \Gamma^{MH}} \sum_{t=0}^{\infty} \sum_{s^t} \beta^t [y(s_t) - c_t(s^t)] \Pi(s^t | s_0, a^{t-1}(s^{t-1})) + \\
& + \gamma \sum_{t=0}^{\infty} \sum_{s^t} \beta^t [u(c_t(s^t)) - v(a_t(s^t))] \Pi(s^t | s_0, a^{t-1}(s^{t-1}))
\end{aligned} \tag{6}$$

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

where γ is a function of V^{out} in the original problem¹¹. Let $\beta^t \lambda_t(s^t) \Pi(s^t | s_0, a^{t-1}(s^{t-1}))$ be the Lagrange multiplier associated to each ICC. The Lagrangean is:

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

The Lagrangean can be manipulated with simple algebra to get the following expression:

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

¹¹To see how we can rewrite the original problem as a social welfare maximization, notice that equation (3) 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. Therefore, we can associate a strictly positive Lagrange multiplier (say, γ) to (3), which will be a function of V^{out} . This Lagrange multiplier can be seen as a Pareto-Negishi weight on the agent's utility. I can fully characterize the Pareto frontier of this economy by solving the problem for different values of γ between zero and infinity. 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 and can be dropped.

where

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

The intuition is simple. For any s^t , $\lambda_t(s^t)$ is the shadow cost of implementing an incentive compatible allocation, i.e. the amount of resources that the principal must spend to implement an incentive compatible contract. The expression $\frac{\pi_a(s_{t+1}|s_t, a_t(s^t))}{\pi(s_{t+1}|s_t, a_t(s^t))}$ is a measure of the informativeness of output as a signal for effort, and therefore an indirect measure of the effect of effort on the observed result. Rewrite the definition of $\phi(s^t)$ as:

$$\begin{aligned} \phi_{t+1}(s^t, \hat{s}) &= \phi_t(s^t) + \lambda_t(s^t) \frac{\pi_a(s_{t+1} = \hat{s} | s_t, a_t(s^t))}{\pi(s_{t+1} = \hat{s} | s_t, a_t(s^t))} \quad \forall \hat{s} \in S \\ \phi_0(s^0) &= \gamma \end{aligned} \quad (7)$$

Therefore, from (7) we can see $\phi_t(s^t)$ as the *Pareto-Negishi weight* of the agent's lifetime utility, that evolves *endogenously* in order to track the agent's effort. The optimal contract promises that the weight in $t+1$ will differ from the weight in t by an amount equal to the shadow cost $\lambda_t(s^t)$ multiplied by a measure of the effect of effort on the output distribution.

2.2 Recursive formulation

Marcet and Marimon (2009) show that, for full information problems with forward-looking constraints, the Lagrangean has a recursive structure and can be used to find a solution of the original problem. The question is therefore whether the same arguments can also be used in the principal-agent framework. By the duality theory (see for example Luenberger (1969)), a solution of the original problem corresponds to a saddle point of the Lagrangean¹², i.e. the contract

$$(c^{\infty*}, a^{\infty*}, \tau^{\infty*}) = \{c_t^*(s^t), a_t^*(s^t), y(s_t) - c_t^*(s^t) \quad \forall s^t \in S^{t+1}\}_{t=0}^{\infty}$$

is a solution for the original problem if there exist a sequence $\{\lambda_t^*(s^t) \quad \forall s^t \in S^{t+1}\}_{t=0}^{\infty}$ of Lagrange multipliers such that $(c^{\infty*}, a^{\infty*}, \lambda^{\infty*}) = \{c_t^*(s^t), a_t^*(s^t), \lambda_t^*(s^t) \quad \forall s^t \in S^{t+1}\}_{t=0}^{\infty}$ satisfy:

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

Finding these sequences can be complicated. However, had this Lagrangean problem a recursive representation, it would be possible to characterize the solutions with standard numerical methods that exploit dynamic programming arguments. This is the focus

¹²Notice that, in my setup, the conditions stated by Marcet and Marimon (2009) for equivalence between the saddle-point solution of the Lagrangean and the solution of the original problem are satisfied.

of this section. In particular, value and policy functions (or correspondences, more generally) are shown to depend on the state of the world s_t and the Pareto-Negishi weight $\phi_t(s^t)$.

I follow the strategy of MM by showing that a generalized version of Problem (6) is recursive in an enlarged state space. The generalized version of (6) is:

$$\begin{aligned}
W_\theta^{SWF}(s_0) &= \max_{\{a_t(s^t), c_t(s^t)\}_{t=0}^\infty \in \Gamma^{MH}} \bar{\phi}^0 \sum_{t=0}^{\infty} \sum_{s^t} \beta^t [y(s_t) - c_t(s^t)] \Pi(s^t | s_0, a^{t-1}(s^{t-1})) + \\
&\quad + \gamma \sum_{t=0}^{\infty} \sum_{s^t} \beta^t (u(c_t(s^t)) - v(a_t(s^t))) \Pi(s^t | s_0, a^{t-1}(s^{t-1})) \\
s.t. \quad v'(a_t(s^t)) &= \sum_{j=1}^{\infty} \beta^j \sum_{s^{t+j}|s^t} \frac{\pi_a(s_{t+1} | s_t, a_t(s^t))}{\pi(s_{t+1} | s_t, a_t(s^t))} \times \\
&\quad \times [u(c_{t+j}(s^{t+j})) - v(a_{t+j}(s^{t+j}))] \Pi(s^{t+j} | s^t, a^{t+j-1}(s^{t+j-1} | s^t)) \\
&\quad \quad \quad \forall s^t, t \geq 0
\end{aligned}$$

Notice that if $\bar{\phi}^0 = 1$, then we are back to (6). Write down the Lagrangean of this problem by assigning a Lagrange multiplier $\beta^t \lambda_t(s^t) \Pi(s^t | s_0, a^{t-1}(s^{t-1}))$ to each ICC constraint:

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

Notice that $r(a, c, s) \equiv y(s) - c$ is uniformly bounded by natural debt limits, so there exists a lower bound $\underline{\kappa}$ such that $r(a, c, s) \geq \underline{\kappa}$. We can therefore define $\kappa < \frac{\underline{\kappa}}{1-\beta}$. Define $\varphi(\phi, \lambda, a, s') \equiv \phi + \lambda \frac{\pi_a(s' | s, a)}{\pi(s' | s, a)}$, $h_0^P(a, c, s) \equiv r(a, c, s)$, $h_1^P(a, c, s) \equiv r(a, c, s) - \kappa$, $h_0^{ICC}(a, c, s) \equiv u(c) - v(a)$, $h_1^{ICC}(a, c, s) \equiv -v'(a)$, $\theta \equiv [\phi^0 \quad \phi] \in \mathbb{R}^2$, $\chi \equiv [\lambda^0 \quad \lambda]$ and

$$\begin{aligned}
h(a, c, \theta, \chi, s) &\equiv \theta h_0(a, c, s) + \chi h_1(a, c, s) \\
&\equiv [\phi^0 \quad \phi] \begin{bmatrix} h_0^P(a, c, s) \\ h_0^{ICC}(a, c, s) \end{bmatrix} + [\lambda^0 \quad \lambda] \begin{bmatrix} h_1^P(a, c, s) \\ h_1^{ICC}(a, c, s) \end{bmatrix}
\end{aligned}$$

which is homogenous of degree 1 in (θ, χ) . The Lagrangean can be written as:

$$L_\theta(s_0, \gamma, c^\infty, a^\infty, \chi^\infty) = \sum_{t=0}^{\infty} \sum_{s^t} \beta^t h(a_t(s^t), c_t(s^t), \theta_t(s^t), \chi_t(s^t), s_t) \Pi(s^t | s_0, a^{t-1}(s^{t-1}))$$

where

$$\theta_{t+1}(s^t, \widehat{s}) = \varphi(\theta_t(s^t), \chi_t(s^t), a_t(s^t), \widehat{s}) \quad \forall \widehat{s} \in S$$

$$\theta_0(s^0) = \begin{bmatrix} \overline{\phi}^0 & \gamma \end{bmatrix}$$

The constraint defined by $h_1^P(a, c, s)$ is never binding by definition, therefore $\lambda_t^0(s^t) = 0$ and $\phi_t^0(s^t) = \overline{\phi}^0 \quad \forall s^t, t \geq 0$, which implies that the only relevant state variable is $\phi_t(s^t)$.

The next step is to show that all solutions of the Lagrangean have a recursive structure. This is done in two steps. Firstly, Proposition 1 proves that a particular functional equation (the *saddle point functional equation*) associated with the Lagrangean satisfies the assumptions of the Contraction Mapping Theorem. This functional equation is the equivalent of a Bellman equation for saddle point problems. Secondly, it must hold that solutions of the functional equation are solutions of the Lagrangean and viceversa. This is a trivial application of MM (Theorems 3 and 4) and therefore it is omitted¹³.

Associate the following saddle point functional equation to the Lagrangean

$$J(s, \theta) = \min_{\chi} \max_{a, c} \left\{ h(a, c, \theta, \chi, s) + \beta \sum_{s'} \pi(s' | s, a) J(s', \theta'(s')) \right\} \quad (8)$$

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

In order to show that there is a unique value function $J(s, \theta)$ that solves Problem (8), it is sufficient to prove that the operator on the right hand side of the functional equation is a contraction¹⁴.

There are two technical differences with the original framework in MM. Firstly, the law of motion for Pareto-Negishi weights depends (non-linearly) on the current allocation, while in MM it only depends (linearly) on the Lagrange multipliers. Secondly, the probability distribution of the future states is endogenous and depends on the optimal effort $a_t(s^t)$. Therefore, on a first inspection, the problem looks much more complicated than the standard MM setup. However, Proposition 1 shows that MM's arguments also work here.

¹³Messner and Pavoni (2004) show an example with full information in which the policy function that solves the saddle point functional equation can be suboptimal or even unfeasible. To avoid these issues, though, it is sufficient to impose that the policy function satisfies all the constraints of the original problem. Since I solve for the Lagrangean first-order conditions, I always impose all the constraints.

¹⁴In general, this problem will yield a unique value function and a policy correspondence. In the rest of the paper, assume the policy correspondence is single-valued, i.e. it is a policy function. Marimon, Messner and Pavoni (2009) generalize the arguments of MM for policy correspondence, and similar ideas can be used in my setup.

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

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

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

is a contraction.

Proof. Appendix A. ■

Proposition 1 shows that the saddle point problem is recursive in the state space $(s, \theta) \in S \times \mathbb{R}^2$. Notice that the result of Proposition 1 is valid for any $K > 0$. Moreover, whenever the Lagrangean has a solution, the Lagrange multipliers are bounded (see MM for further discussion of this issue). Hence, a recursive solution of Problem (8) is a solution of the Lagrangean, and more importantly it is a solution of the original problem. As a consequence, it is enough to restrict the search for optimal contracts to the set of policy functions that are Markovian in the space $(s, \theta) \in S \times \mathbb{R}^2$. But remember that the first element of θ is constant for any t and hence the only relevant endogenous state is $\phi_t(s^t)$. Therefore, from this point of view, finding the optimal contract has the same numerical complexity as finding the optimal allocations in a standard RBC model¹⁵.

2.3 The meaning of Pareto-Negishi weights

To better understand the role of $\phi_t(s^t)$, assume there are only two possible realizations of the state of nature: $s_t \in \{s_L, s_H\}$. At time t , the weight is equal to ϕ_t . In period $t+1$, given our assumption on the likelihood ratio, the Pareto-Negishi weight is higher than ϕ_t if the principal observes s_H , while it is lower than ϕ_t if she observes s_L (a formal proof of this fact is obtained in Lemma 1 in Appendix A). Therefore the principal promises that the agent will be rewarded with a higher weight in the social welfare function (i.e., the principal will care more about him) if a good state of nature is observed, while it will be punished with a lower weight (i.e., the principal will care less about him) if a bad state of nature happens.

Appendix A contains some standard results of dynamic agency theory obtained by using Pareto-Negishi weights. The famous *immiseration result*¹⁶ of Thomas and Worrall

¹⁵Notice that, since in the Lagrangean formulation the constant γV^{out} was eliminated, the value of the original problem is:

$$W(s_0) = W^{SWF}(s_0) - \gamma V^{out} = J(s_0, [1 \quad \gamma]) - \gamma V^{out}$$

where $V^{out} = V(s_0; \tau^{\infty*})$ is the agent's lifetime utility implied by the optimal contract.

¹⁶The immiseration result states that agent's consumption goes almost surely to its lower bound in an optimal contract.

(1990) is implied by Proposition 3, where I show that the Pareto-Negishi weight is a non-negative martingale which almost surely converges to zero.

2.4 A comparison with APS

It is interesting to compare this result with the promised utility approach. The APS recursive formulation in (9) is based on the continuation values of the agent:

$$P(U, s) = \max_{\{c, \{U(s')\}_{s' \in S}, a^*\}} \left[y(s) - c + \beta \sum_{s'} \pi(s' | s, a^*) P(U(s'), s') \right] \quad (9)$$

$$s.t. \quad u(c) - v(a^*) + \beta \sum_{s'} \pi(s' | s, a^*) U(s') = U \quad (10)$$

$$a^* = \arg \max_{a \in A} \left\{ u(c) - v(a) + \beta \sum_{s'} \pi(s' | s, a^*) U(s') \right\} \quad (11)$$

$$U(s') \in \mathcal{U} \quad (12)$$

where (10) is the promise-keeping constraint, (11) is the incentive compatibility constraint, and (12) constrains the continuation values to belong to the feasible set \mathcal{U} , which is the fixed point of the operator:

$$B(W) = \left\{ \begin{array}{l} U \in W : u(c) - v(a^*) + \beta \sum_{s'} \pi(s' | s, a^*) U(s') = U \\ a^* = \arg \max_{a \in A} \{ u(c) - v(a) + \beta \sum_{s'} \pi(s' | s, a^*) U(s') \} \end{array} \right\}$$

The principal enforces incentive compatible contracts by promising the agent a higher continuation value if a good state of nature is observed in the future, and a lower continuation value if a bad state is observed. The two methodologies, therefore, differ in the way they make and enforce promises, but they both have the same number of state variables. Notice however that the APS technique needs to characterize the feasible set for continuation values by solving the fixed point problem $B(\mathcal{U}) = \mathcal{U}$, while in the recursive Lagrangean approach the problem is well defined for any Pareto-Negishi weight in the real line. Therefore, because of this additional step in the promised utilities method, the Lagrangean approach is simpler than the APS one. However, in this framework, the characterization of \mathcal{U} is easy: the feasible set is an interval, and the extrema of this interval can be found with standard algorithms. Things become more complicated when there are several agents and endogenous state variables. This is the subject of the next section.

3 A more general theorem

In this section, I derive a generalization of Proposition 1 for the case in which there are observable endogenous state variables and several agents. Suppose that all the

assumptions in MM are satisfied. In the following, when needed, other assumptions on the primitives of the model will be specified.

Assume there are N agents indexed by $i = 1, \dots, N$. Each agent is subject to an observable Markov state process $\{s_{it}\}_{t=0}^{\infty}$, where $s_{it} \in S_i$, s_{i0} is known, and the process is common knowledge. The process is independent across agents. Let $S \equiv \prod_{i=1}^N S_i$ and $s_t \equiv \{s_{1t}, \dots, s_{Nt}\} \in S$ be the state of nature in the economy, let $s^t \equiv \{s_0, \dots, s_t\}$ be the history of these realizations. Let $w_t(s^t) \equiv (w_{1t}(s^t), \dots, w_{Nt}(s^t))$ for any generic variable w , and let $W = \prod_{i=1}^N W_i$ for any generic set W .

Each agent exerts a costly action $a_{it}(s^t) \in A_i$, where A_i is a convex subset of \mathbb{R} . This action is unobservable to other players, and it affects the next period distribution of states of nature. Let $\pi^i(s_{i,t+1} | s_{it}, a_{it}(s^t))$ be the probability that state is $s_{i,t+1}$ conditional on both the past state and the effort exerted by the agent i in period t . Therefore, since the processes are independent across agents, define $\Pi(s^{t+1} | s_0, a^t(s^t)) = \prod_{i=1}^N \prod_{j=0}^t \pi^i(s_{i,j+1} | s_{ij}, a_{ij}(s^j))$ to be the cumulated probability of history s^{t+1} given the whole history of unobserved actions $a^t(s^t) \equiv (a_0(s^0), a_1(s^1), \dots, a_t(s^t))$. Probabilities $\pi^i(s_{i,t+1} | s_{it}, a_{it}(s^t))$ are differentiable in $a_{it}(s^t)$ as many times as necessary. Denote the derivative with respect to $a_{it}(s^t)$ as $\pi_a^i(s_{i,t+1} | s_{it}, a_{it}(s^t))$, and assume the likelihood ratio is bounded. Allocations are indicated by the vector $\varsigma_{it}(s^t) \in \Upsilon_i$. Each agent is endowed with a vector of endogenous state variables $x_{it}(s^t) \in X_i$, $X_i \subseteq R^m$ convex, that evolve according to the following laws of motion:

$$x_{i,t+1}(s^t, s_{t+1}) = \ell^i(x_{it}(s^t), \varsigma_{it}(s^t), a_{it}(s^t), s_{i,t+1})$$

The (uniformly bounded) per-period payoff function of each agent is given by

$$r^i(\varsigma_i, a_i, x_i, s)$$

and $r^i : \Upsilon_i \times A_i \times X_i \times S \rightarrow R$ is non-decreasing in ς_i , decreasing in a_i , concave in x_i and strictly concave in (ς_i, a_i) , (at least) once continuously differentiable in (ς_i, x_i) and twice continuously differentiable in a_i . The resource constraint is¹⁷:

$$p(x_t(s^t), \varsigma_t(s^t), a_t(s^t), s_t) \geq 0$$

Notice that the standard principal-agent setup belongs to this class of models, if we set $N = 2$, $X_i = \emptyset$, $r^P(\varsigma_i, a_i, x_i, s) \equiv y(s) - c_A$, $r^A(\varsigma_i, a_i, x_i, s) \equiv u(c_A) - v(a_A)$, and we assume that the principal does not exert effort or her effort has no impact on the distribution of the state of nature. More generally, the result in this section can be extended to the case in which only a subset of agents has a moral hazard problem. However, the notation becomes burdensome, hence for expositional purposes it is better

¹⁷Constraints that involve future endogenous variables, like participation constraints or Euler equations, can be incorporated by following the standard MM approach. Since they only complicate the notation, they are not included in the analysis.

to stick with the case where all agents involved in the contract have a moral hazard problem.

A feasible contract \mathcal{W} is a triplet of sequences $(\varsigma^\infty, a^\infty, x^\infty) \equiv \{\varsigma_t(s^t), a_t(s^t), x_t(s^t)\}_{t=0}^\infty$ $\forall s^t \in S^{t+1}$ that belongs to the set:

$$\begin{aligned} \Gamma^{GT} \equiv & \{(\varsigma^\infty, a^\infty, x^\infty) : a_t(s^t) \in A, \varsigma_t(s^t) \in \Upsilon, x_t(s^t) \in X, \\ & x_{i,t+1}(s^t, s_{t+1}) = \ell^i(x_{it}(s^t), \varsigma_{it}(s^t), a_{it}(s^t), s_{i,t+1}) \forall i, \\ & p(x_t(s^t), \varsigma_t(s^t), a_t(s^t), s_t) \geq 0 \quad \forall s^t \in S^{t+1}, t \geq 0\} \end{aligned}$$

Let $\omega \equiv \{\omega_i\}_{i=1}^N \in R^N$ be a vector of initial Pareto-Negishi weights, and assume the use of the first-order approach (FOA) is justified¹⁸. To avoid burdensome notation, in the following I do not explicitly indicate the measurability of each allocation with respect to history s^t . Since FOA is valid, we can use the first-order conditions of the agents' problems with respect to hidden actions as incentive compatibility constraints:

$$\begin{aligned} r_a^i(\varsigma_{it}, a_{it}, x_{it}, s_t) + \sum_{j=1}^{\infty} \sum_{s^{t+j}} \beta^j \frac{\pi_a^i(s_{i,t+1} | s_{it}, a_{it})}{\pi^i(s_{i,t+1} | s_{it}, a_{it})} \times \\ \times r^i(\varsigma_{i,t+j}, a_{i,t+j}, x_{i,t+j}, s_{t+j}) \Pi(s^{t+j} | s^{t+j-1}, a_{t+j-1}) = 0 \quad \forall i = 1, \dots, N \end{aligned} \quad (13)$$

The constrained efficient allocation is the solution of the following maximization problem:

$$\begin{aligned} P(s_0) = \max_{\mathcal{W} \in \Gamma^{GT}} \left\{ \sum_{i=1}^N \sum_{t=0}^{\infty} \beta^t \sum_{s^t} \omega_i r^i(\varsigma_{it}, a_{it}, x_{it}, s_t) \Pi(s^t | s^{t-1}, a_{t-1}) \right\} \\ s.t. \quad (13) \end{aligned}$$

Let $\beta^t \lambda_{it}(s^t) \Pi(s^t | s^{t-1}, a_{t-1})$ be the Lagrange multiplier for the incentive-compatibility constraint (13) of agent i . Substitute for the resource constraint and write the Lagrangean as:

$$\begin{aligned} L(s_0, \omega, \mathcal{W}, \lambda^\infty) = \\ = \sum_{i=1}^N \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \{ \phi_{it} r^i(\varsigma_{it}, a_{it}, x_{it}, s_t) + \\ + \lambda_{it} r_a^i(\varsigma_{it}, a_{it}, x_{it}, s_t) \} \Pi(s^t | s^{t-1}, a_{t-1}) \end{aligned}$$

where, for any i ,

$$\begin{aligned} x_{i,t+1}(s^t, s_{t+1}) &= \ell^i(x_{it}(s^t), \varsigma_{it}(s^t), a_{it}(s^t), s_{i,t+1}) \\ \phi_{i,t+1}(s^t, s_{t+1}) &= \phi_{it}(s^t) + \lambda_{it}(s^t) \frac{\pi_a^i(s_{i,t+1} | s_{it}, a_{it}(s^t))}{\pi^i(s_{i,t+1} | s_{it}, a_{it}(s^t))} \\ \phi_{i0}(s_0) &= \omega_i, \quad x_{i0} \text{ given} \end{aligned}$$

¹⁸It is easy to see that, in this setup as well, standard sufficient conditions for the static principal-agent problem will justify the validity of the first-order approach.

The newly defined variables $\phi_{it}(s^t)$, $i = 1, \dots, N$, are endogenously evolving Pareto-Negishi weights which have the same interpretation as in the previous section: they are optimally chosen by the planner to implement an incentive compatible allocation and they summarize the contract's (history-dependent) promises for each agent.

3.1 Recursivity

Notice that this problem is already in the form of a social welfare function maximization.

Let $\varphi^i(\phi_i, \lambda_i, a_i, s^i) \equiv \phi_i + \lambda_i \frac{\pi_a^i(s_i^i | s_i, a_i)}{\pi^i(s_i^i | s_i, a_i)}$, $h_0^i(\varsigma, a, x, s) \equiv r^i(\varsigma_i, a_i, x_i, s)$, $h_1^i(\varsigma, a, x, s) \equiv r_a^i(\varsigma_i, a_i, x_i, s)$, and

$$h(\varsigma, a, x, \phi, \lambda, s) \equiv \phi h_0(\varsigma, a, x, s) + \lambda h_1(\varsigma, a, x, s)$$

which is homogenous of degree 1 in (ϕ, λ) . The Lagrangean can be written as:

$$L(s_0, \omega, \varsigma^\infty, a^\infty, x^\infty, \lambda^\infty) = \sum_{t=0}^{\infty} \sum_{s^t} \beta^t h(\varsigma_t, a_t, x_t, \phi_t, \lambda_t, s_t) \Pi(s^t | s_0, a^{t-1}(s^{t-1}))$$

where

$$\begin{aligned} x_{t+1}(s^t, \widehat{s}) &= \ell(x_t(s^t), \varsigma_t(s^t), a_t(s^t), \widehat{s}) \\ \phi_{t+1}(s^t, \widehat{s}) &= \varphi(\phi_t(s^t), \lambda_t(s^t), a_t(s^t), \widehat{s}) \quad \forall \widehat{s} \in S \end{aligned}$$

$$\phi_0(s^0) = \omega, \quad x_{i0} \text{ given}$$

The corresponding saddle point functional equation is

$$\begin{aligned} J(s, \phi, x) &= \min_{\lambda} \max_{\varsigma, a} \left\{ h(\varsigma, a, x, \phi, \lambda, s) + \beta \sum_{s'} \pi(s' | s, a) J(s', \phi'(s'), x'(s')) \right\} \quad (14) \\ \text{s.t.} \quad x'(s') &= \ell(x, \varsigma, a, s') \\ \phi'(s') &= \varphi(\phi, \lambda, a, s') \quad \forall s' \end{aligned}$$

Proposition 2 shows that the operator on the RHS of (14) is a contraction. The proof is a simple repetition of the steps followed to prove Proposition 1, in a different functional space.

Proposition 2 *Fix an arbitrary constant $K > 0$ and let $K_\theta = \max\{K, K \|\phi\|\}$. The operator*

$$\begin{aligned} (T_K f)(s, \phi, x) &\equiv \min_{\{\lambda > 0: \|\lambda\| \leq K_\theta\}} \max_{\varsigma, a} \left\{ h(\varsigma, a, x, \phi, \lambda, s) + \beta \sum_{s'} \pi(s' | s, a) f(s', \phi'(s'), x'(s')) \right\} \\ \text{s.t.} \quad x'(s') &= \ell(x, \varsigma, a, s') \\ \phi'(s') &= \varphi(\phi, \lambda, a, s') \quad \forall s' \end{aligned}$$

is a contraction.

Proof. Straightforward by repeating the steps to prove Proposition 1 in the following space of functions:

$$M = \{f : S \times \mathbb{R}^N \times X \longrightarrow \mathbb{R} \text{ s.t.} \\ \begin{array}{l} \text{a) } \forall \alpha > 0 \quad f(\cdot, \alpha\phi, \cdot) = \alpha f(\cdot, \phi, \cdot) \\ \text{b) } f(s, \cdot, \cdot) \text{ is continuous and bounded} \end{array} \}$$

with norm

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

■

Using the same arguments as in section 2, a recursive solution of the original problem can be found by solving the functional equation (14), provided that the optimal policy correspondence is single-valued.

Notice that this problem has $N(m+1)$ state variables. However, the value function of the problem is homogenous of degree one in the vector of endogenous weights ϕ . This fact implies:

$$\frac{1}{\phi_1} J(s, \phi_1, \dots, \phi_N, x) = J\left(s, 1, \frac{\phi_2}{\phi_1}, \dots, \frac{\phi_N}{\phi_1}, x\right) \equiv \tilde{J}\left(s, \frac{\phi_2}{\phi_1}, \dots, \frac{\phi_N}{\phi_1}, x\right)$$

therefore the dimension of the state space is reduced to $N(m+1) - 1$. Moreover, the individual continuation values for each agent i are homogeneous of degree zero with respect to the vector of endogenous weights ϕ ¹⁹. These two facts are helpful in computational applications.

3.2 A comparison with APS

In this more general framework, the promised utility approach gives a recursive formulation which uses a new state space including continuation values U_t^i and the natural states variables x_t of the problem:

¹⁹This is a consequence of the homogeneity of degree one of the planner's value function. MM show that individual continuation values must satisfy an individual saddle-point functional equation, and they must be homogeneous of degree zero in order to satisfy the functional equation (14). The same argument holds in the current setup.

$$\begin{aligned}
P\left(\{U_i, x_i\}_{i=1, \dots, N}, s\right) &= \max_{\{c_i, a_i^*, \{U^i(s'), x'_i(s')\}_{s' \in S}\}_{i=1, \dots, N}} \\
&\left\{ \left[\sum_i \omega_i r^i(\varsigma_i, a_i^*, x_i, s) \right] + \beta \sum_{s'} \pi(s' | s, a^*) P\left(\{U^i(s'), x'_i(s')\}_{i=1, \dots, N}, s'\right) \right\} \\
s.t. \quad r^i(\varsigma_i, a_i^*, x_i, s) + \beta \sum_{s'} \pi(s' | s, a^*) U^i(s') &= U_i \quad i = 1, \dots, N \tag{15}
\end{aligned}$$

$$a_i^* = \arg \max_{a_i \in A_i} \left\{ r^i(\varsigma_i, a_i, x_i, s) + \beta \sum_{s'} \pi(s' | s, (a_i, a_{-i}^*)) U^i(s') \right\} \quad i = 1, \dots, N \tag{16}$$

$$x'_i(s') = \ell^i(x_i, \varsigma_i, a_i^*, s') \quad i = 1, \dots, N, p(x, \varsigma, a^*, s) \geq 0 \quad \forall s \in S, U^i(s') \in \mathcal{U}(x) \tag{17}$$

where (15) is the promise-keeping constraint, (16) is the incentive compatibility constraint, and the *value correspondence* $\mathcal{U}(x) \subset R^N$, $\mathcal{U} : X \rightrightarrows R^N$ in (17) is the fixed point of the operator:

$$B(W(x)) = \left\{ \begin{array}{l} \{U_i\}_{i=1, \dots, N} \in W(x) : r^i(\varsigma_i, a_i^*, x_i, s) + \beta \sum_{s'} \pi(s' | s, a^*) U^i(s') = U_i \\ a_i^* = \arg \max_{a_i \in A_i} \{ r^i(\varsigma_i, a_i, x_i, s) + \beta \sum_{s'} \pi(s' | s, (a_i, a_{-i}^*)) U^i(s') \} \\ x'_i(s') = \ell^i(x_i, \varsigma_i, a_i^*, s') \quad i = 1, \dots, N, p(x, \varsigma, a^*, s) \geq 0 \quad \forall s \in S \end{array} \right\}$$

Notice that this formulation has $N(m+1)$ state variables like the recursive Lagrangean problem. However, the correspondence $\mathcal{U}(x)$ is very complicated to characterize even for small values of N and m . While in the case of section 2 the correspondence \mathcal{U} was actually one interval, here there is a different interval for any point of the state space X . Computing this family of intervals is already a formidable task for the case $N+m=3$. There are algorithms that allow an efficient computation of the approximated correspondence (see e.g. Sleet and Yeltekin (2003)), but the complexity of the task increases exponentially with the number of agents and the number of endogenous state variables. This does not happen with the Lagrangean approach, where the characterization of the feasible set is absent.

4 Hidden endogenous states

Proposition 2 refers to cases in which all the endogenous state variables are observable. What happens in the case of unobservable endogenous states? In principle, it is possible to follow the same general idea of combining the first-order approach and the recursive Lagrangean: solve the agent's maximization problem with respect to unobservable variables (in this case, effort and the endogenous unobservable state variables) by taking first-order conditions, and use the latter as constraints in the planner's problem. In

general, first-order conditions for unobservable state variables will be forward-looking, and hence they will fit in the standard MM framework.

However, there is a caveat: the use of the first-order approach in these models is very restrictive (see Kocherlakota (2004) for an example). Moreover, to the best of my knowledge, there are no sufficient conditions that make sure the first-order approach is justified in dynamic models with unobservable endogenous states²⁰. One possibility is to verify numerically if the first-order approach is valid, along the lines of the verification algorithm suggested by Abraham and Pavoni (2009). In the rest of this section, as an example, the model in Abraham and Pavoni (2009) (AP from here on) with hidden effort and hidden assets is studied with recursive Lagrangeans (details on recursivity and the verification procedure are relegated in Appendix B). In the numerical section, I will provide a computed example in which the verification procedure guarantees we can use the first-order approach.

4.1 Repeated moral hazard with hidden saving and borrowing

Let $\{b_t(s^t)\}_{t=-1}^\infty$, b_{-1} given, be a sequence of one-period bond holdings, each of which costs the agent 1 today and returns R tomorrow. Assume that the principal cannot monitor the bond market, so that the asset accumulation is unobservable to her²¹. Then agent's budget constraint becomes:

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

while the instantaneous utility function for the agent is the same as in section 2. The agent's problem is:

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

²⁰Abraham and Pavoni (2008) and Koehne (2009) provide sufficient conditions for the two-period case, but their conditions cannot be easily extended to a multiperiod setting.

²¹Werning (2001, 2002) and AP analyze a model with hidden effort and hidden assets. This problem generates a continuum of incentive constraints (for each possible income realization, there is a continuum of possible asset positions for which we have to specify an incentive compatibility constraint). Hence the feasible set of continuation values has infinite dimension and APS techniques cannot be used. In order to overcome this issue, 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 solution of their auxiliary problem is characterized by three state variables (income, promised utility and consumption marginal utility), and can be solved recursively by value function iteration. AP also provide a numerical *ex-post* procedure to verify if the first-order approach delivers the true incentive compatible allocation. Even if their work is big step ahead in the analysis of this class of models, the use of APS arguments makes their numerical algorithm too slow for calibration purposes, and any extension of the model is computationally unmanageable.

where

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

The first-order approach, in this framework, amounts to taking first order conditions with respect to all unobservable variables, i.e. effort and bond holdings. The resulting constraints are equation (4) as in section 2, and the following Euler equation:

$$u'(c_t(s^t)) = \beta R \sum_{s_{t+1}} u'(c_{t+1}(s^t, s_{t+1})) \pi(s_{t+1} | s_t, a_t(s^t)) \quad (18)$$

The presence of hidden assets requires both (4) and (18) to be included in the set of constraints for the principal's problem.

4.2 The recursive Lagrangean

Let $\beta^t \eta_t(s^t) \Pi(s^t | s_0, a^{t-1}(s^{t-1}))$ be the Lagrange multiplier for equation (18), and $\beta^t \lambda_t(s^t) \Pi(s^t | s_0, a^{t-1}(s^{t-1}))$ the Lagrange multiplier for (4). The Lagrangean can be manipulated to get:

$$L(s_0, \gamma, c^\infty, a^\infty, \lambda^\infty, \eta^\infty) = \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \{y(s_t) - c_t(s^t) + \phi_t(s^t) [u(c_t(s^t)) - v(a_t(s^t))] + \\ -\lambda_t(s^t) v'(a_t(s^t)) + [\eta_t(s^t) - R\zeta_t(s^t)] u_c(c_t(s^t))\} \Pi(s^t | s_0, a^{t-1}(s^{t-1})) \quad (19)$$

where

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

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

This problem is characterized by two costate variables: the Pareto weight $\phi_t(s^t)$ and the new costate $\zeta_t(s^t)$, which keeps track of the Euler equation. Using the same arguments of Proposition 1, it is possible to show that the problem is recursive in the state space that includes (s, ϕ, ζ) as states variables (see Proposition 4 in Appendix B for details). As mentioned above, since it is not sure that the first-order approach is justified, it is necessary to verify numerically that the agent actually likes the optimal contract, i.e. that there are no profitable deviations for the agent. Appendix B suggests a numerical algorithm based on AP's verification procedure that checks the validity of the first-order approach.

5 Numerical examples

In this section, I describe the algorithm and I provide four computed examples.

5.1 The algorithm

For simplicity, the Markov process has only two possible realizations ($S_i \equiv \{s^L, s^H\}$ for any i , $s^L < s^H$). Assume the state is i.i.d., and use the simpler notation $\pi_j(a_{it}) = \pi^i(s_{i,t+1} = s^j | a_{it})$, $j = L, H$. Define a generic state of the economy as $\hat{s} \in S$ where $S \equiv \prod_{i=1}^N S_i$, and let $\pi(\hat{s} | a_t) \equiv \pi(s_{t+1} = \hat{s} | a_t) = \prod_{i=1}^N \pi^i(s_{i,t+1} = \hat{s}_i | a_{it})$. The numerical procedure is a collocation algorithm (see Judd (1998)) over the first-order conditions of the Lagrangean. From the recursive formulation we know that policy functions depend on the natural states of the problem and on the costates (i.e., Pareto weights) that come out from the Lagrangean approach. Let ς be the vector of allocations (including hidden actions), χ be the vector of Lagrange multipliers, $x \in X$ be the vector of natural states, and $\theta \in \Theta$ be the vector of costates, and define $R(s, \varsigma, \chi, x, \theta)$ as the objective function in the saddle point functional equation, and $r^i(s, \varsigma, \chi, x, \theta)$ as the instantaneous utility function for the agent i . The algorithm therefore is the following:

1. Fix ω_i , $i = 1, \dots, N$ and define a discrete grid $G \subset S \times X \times \Theta$ for natural states and costates.
2. Approximate policy functions for allocations ς and Lagrange multipliers χ , the value function of the principal J and agents' continuation value U^i using cubic splines or Chebychev polynomials, and set initial conditions for the approximation coefficients.
3. For any $(s, x, \theta) \in G$, use a nonlinear solver ²² to solve for the Lagrangean first order conditions and the following equations for the continuation value U^i and the value function J :

$$U^i(s, x, \theta) = r^i(s, \varsigma, \chi, x, \theta) + \beta \sum_{\hat{s}} \pi(\hat{s} | a_t) U^i(\hat{s}, x', \theta'(\hat{s})) \quad (20)$$

$$J(s, x, \theta) = R(s, \varsigma, \chi, x, \theta) + \beta \sum_{\hat{s}} \pi(\hat{s} | a_t) J(\hat{s}, x', \theta'(\hat{s})) \quad (21)$$

I use the Miranda-Fackler Compecon toolbox for function approximation. In all applications, steps 1-3 are applied first to a grid with very few gridpoints, and then the accuracy of the approximation is increased by applying steps 1-3 to a finer grid. Typically, a good approximation is obtained with few grid points. Due to the use of a non-linear equation solver, it is crucial to find good initial conditions for the parameters

²² In all applications presented in this paper, I use a version of the Broyden algorithm coded by Michael Reiter.

of the interpolants. In general, it is a good idea to start from the solution of a simpler model (e.g., for the hidden effort and hidden assets problem, start from the solution of the basic repeated moral hazard model). Homotopy methods help if the latter is not enough. The algorithm is coded in Matlab²³.

5.2 Examples

5.2.1 Repeated moral hazard

In order to make the algorithm clear, the first example of a standard repeated moral hazard setup is explained with all the details. Let simplify the notation by writing a generic variable as x_t instead of $x_t(s^t)$. Assume that the income process has two possible realizations ($y^L = y(s^L)$ and $y^H = y(s^H)$).

The Lagrangean first-order conditions are

$$c_t : \quad u'(c_t) = \frac{1}{\phi_t} \quad (22)$$

$$a_t : \quad 0 = -\lambda_t v''(a_t) - \phi_t v'(a_t) + \beta \pi_a(a_t) [J(y^H, \phi_{t+1}) - J(y^L, \phi_{t+1})] + \quad (23)$$

$$+ \beta \lambda_t \left\{ \pi(a_t) \frac{\partial \left(\frac{\pi_a(a_t)}{\pi(a_t)} \right)}{\partial a_t} [u(c_{t+1}) - v(a_{t+1}) \mid y_{t+1} = y^H] + \right.$$

$$\left. + (1 - \pi(a_t)) \frac{\partial \left(\frac{-\pi_a(a_t)}{1 - \pi(a_t)} \right)}{\partial a_t} [u(c_{t+1}) - v(a_{t+1}) \mid y_{t+1} = y^L] \right\}$$

$$\lambda_t : \quad 0 = -v'(a_t) + \beta \pi_a(a_t) [U(y^H, \phi_{t+1}^H) - U(y^L, \phi_{t+1}^L)] \quad (24)$$

Fix γ and choose a discrete grid for ϕ_t that contains γ . Approximate with cubic splines a , λ , U and J on each grid node. Consumption is obtained directly from ϕ by using (22): $c = u'^{-1}(\phi^{-1})$. There are four non-linear equations left: (23), (24), (20) and (21).

I choose the following functional forms:

$$u(c) = \frac{c^{1-\sigma}}{1-\sigma}$$

$$v(a) = \alpha a^\varepsilon$$

$$\pi(a) = a^\nu, \quad a \in (0, 1)$$

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

²³The basic code can be downloaded from my website or I can send it by email.

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 . 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 in Appendix A. 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 plots the parameterized policy functions for transfers, the continuation value of the agent and the value function of the principal. Transfers are increasing in ϕ , as is agent's lifetime utility. On the contrary, the planner's value is monotone decreasing and convex in the Pareto weight.

Figure 3 and 4 show the average allocations across 50 thousands independent simulations for 200 periods, starting with $y_0 = y^H$. In general, these simulations are in line with previous studies: average consumption decreases while average effort increases. 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. Interestingly, ϕ does not show a monotone pattern. To understand the last plot of Figure 4, notice that it is possible to derive the asset holdings implied by optimal allocations (Appendix C shows the details). According to the simulations, average assets must decrease across time²⁴.

Finally, Figure 5 shows the Pareto frontier: it is decreasing and strictly concave.

5.2.2 Hidden assets

This is a computed example for the model presented in section 4. Functional forms and parameters are the same as in the previous example. Policy functions for consumption, agent lifetime utility and λ are depicted in Figure 6 and 7 are strictly increasing and concave in both costates, while effort is strictly decreasing and convex.

The simulated series in Figure 8 and 9 confirm the results in Abraham and Pavoni: on average, consumption and lifetime utility increase across time, while effort decreases. Asset holdings (see Appendix C to see how they are calculated) also increase on average.

Finally, Figure 10 shows the Pareto frontier for different ζ_0 (the natural one is zero): it is decreasing and strictly concave. An application of the verification procedure described in the Appendix B shows that the first-order approach is justified.

5.2.3 Risk sharing

Two identical agents that must share their income in an endowment economy (hence there are no endogenous state variables). There is two-sided moral hazard: they can exert unobservable effort that affect the future distribution of income realization. In

²⁴The asset holdings in the simulation can be interpreted as the saving pattern of an agent in a decentralization of the optimal contract.

terms of the Proposition 2, let $N = 2$, $\varsigma_i \equiv c_i$, $r^i(\varsigma_i, a_i, s) \equiv u(c_i) - v(a_i)$. Theoretical and numerical results for this model are analyzed in detail in Mele (2009), therefore I report a synthesis of them.

I solve the model for the case where agents have the same initial weight in the social welfare function, with the same functional forms and parameters of the previous examples, except for income realizations:

α_i	ε_i	ν_i	σ_i	y_i^L	y_i^H	β	ω_i
0.5	2	0.5	2	.4	.6	0.95	0.5

It is possible to show that, due to the homogeneity properties of value and policy functions, the relevant state variable in this economy is the ratio of endogenous Pareto weights for agent 1 and 2: $\theta \equiv \frac{\phi_2}{\phi_1}$. From the Lagrangean's first-order conditions I obtain $\theta = \frac{u'(c_1)}{u'(c_2)}$ and it can be shown that θ is a submartingale²⁵. The variable θ can be interpreted as a measure of consumption inequality, and given the submartingale characterization, it should be very persistent. These results are in line with theoretical and numerical findings in Zhao (2007) and Friedman (1998). Figures 11 and 12 show that agent 1's consumption and lifetime utility are decreasing in θ for any possible state of the world while effort is increasing in θ . Obviously, the contrary is true for agent 2. Figure 13 and 14 show a sample path of 200 periods. Notice that θ is very persistent as expected. Finally, Figure 15 shows a decreasing, strictly concave Pareto frontier.

5.2.4 Risk sharing in a production economy

This example extends the risk sharing model to a production economy. As for the endowment economy, I present a summary of the results contained in Mele (2009) for more detailed analysis. Each agent can now produce income by using capital. The production function is subject to idiosyncratic productivity shocks, and their distribution is affected by unobservable effort. The law of motion for capital is standard. I keep the same functional forms of the risk sharing example, and I choose the following production function for both agents:

$$f(k_{it}) = A_{it}k_{it}^\rho$$

where A_t is the productivity shock which is affected by the unobservable effort. The baseline parameters are summarized in the following table:

α_i	ε_i	ν_i	σ_i	A_i^L	A_i^H	β	ω_i	δ_i	ρ_i	k_0^i
0.05	2	0.1	2	0.45	0.55	0.95	0.5	0.06	0.3	3.1

Also in this case, we can use the homogeneity properties of value and policy functions to reduce the state space: the relevant state variables are the ratio of Pareto weights

²⁵Work in progress is trying to characterize the long run properties of θ .

$\theta \equiv \frac{\phi_2}{\phi_1}$ and the capital holdings of each agent k_i , $i = 1, 2$. The main difference with respect to the endowment economy is that the persistence in consumption inequality has long-run consequences on the optimal path for capital, and therefore on the long-run path for production.

The following simulation results assume that agents are identical and equally weighted at time zero. Figures 16 and 17 show a simulated sample path for this setup. Both consumption and investment are very volatile. Notice also that consumption inequality is very persistent, and this is reflected in the path of expected discounted utilities of each agent.

The average allocations based on 50000 simulations with a horizon of 500 periods are presented in Figure 18 and 19. The main result is the divergence of capital in the long run. This is due to the history dependence of investment: in each period, it is better to invest a little more in the production technology that has a better history of shocks, i.e. in the richest country. The ratio of Pareto weights θ keeps track of the history, it is increasing on average. It is therefore different from 1 in the long run. The agent that accumulates the most capital is also the agent that exerts the most effort.

5.3 Computational speed and accuracy

The following tables present results for several performance tests. In order to test the computational speed of the algorithm and the accuracy of the approximated solution, the codes solve the examples for different number of grid points. Let M be the number of grid points in each dimension of the state space, e.g. with three endogenous state variables the grid has a total of M^3 grid points. The general message of this exercise is that it is possible to get an accurate solution in few seconds even with relatively few grid points. The hardware is a HP Pavilion dv6700 Notebook PC, with a processor Intel Core2 Duo T5450 at 1.66 GHz and 3 GB RAM.

The accuracy of the approximated solution can be tested by defining a large grid (with roughly 100000 linearly spaced grid points) and calculating the error of the Lagrangean first-order conditions for each grid point under the approximated solution. In the following tables, there are two statistics that measure accuracy: the maximum error and the norm of the error vector.

Table 1: Speed and Accuracy. Repeated Moral Hazard

Grid points	Time (sec)	Max Error	Norm(Error)
10	4.54	5.468001e-005	1.102151e-002
15	6.23	7.766462e-006	1.830439e-003
20	6.93	2.689196e-006	4.700367e-004
30	8.56	3.956188e-007	8.931410e-005
50	12.38	3.828380e-008	6.437146e-006
100	25.20	3.382069e-009	5.187055e-007

Table 1 reports results for the simplest repeated moral hazard model. The computational time is in the order of few seconds, and a fairly good accuracy (i.e., the maximum error is of the order of less than 10^{-5}) is obtained with few grid points.

Table 2: Speed and Accuracy. Hidden Assets

Grid points	Time (sec)	Max Error	Norm(Error)
4	3.61	8.185706e-004	1.366256e-001
6	5.70	6.107623e-004	6.481781e-002
8	9.10	1.347988e-004	1.511452e-002
10	13.55	5.534577e-005	5.425800e-003
12	24.80	2.373655e-005	2.409307e-003
15	84.05	7.876450e-006	8.739442e-004
20	132.72	5.343009e-006	3.026376e-004

Table 2 refers to the case with hidden assets. Also in this case, the computational time is in the order of few seconds. As before, a high accuracy does not require a very fine grid. It is worth mentioning again that the Fortran code of Abraham and Pavoni (2009) runs for around 15 hours before finding a solution. Therefore, the gain in terms of computational intensity is huge (remember that the code for the Lagrangean approach is written in Matlab, which is a much slower programming language than Fortran).

Table 3: Speed and Accuracy. Risk Sharing, Endowment Economy

Grid points	Time (sec)	Max Error	Norm(Error)
10	5.29	5.181706e-006	8.094645e-004
15	6.92	1.228476e-006	1.589214e-004
20	7.85	4.318931e-007	5.363575e-005
30	9.77	8.595712e-008	1.136224e-005
50	13.92	1.175558e-008	1.166124e-006
100	27.06	5.406727e-008	1.177096e-006

The two-agents risk sharing model in an endowment economy has the same level of difficulty than the standard repeated moral hazard model, as table 3 shows. With 10 grid points, the maximum error is less than 10^{-5} . Again, the computational time is in the order of few seconds.

Table 4: Speed and Accuracy. Risk Sharing, Production Economy

Grid points	Time (sec)	Max Error	Norm(Error)
2	2.03	8.194660e-002	1.313794e+001
4	11.34	4.366386e-003	6.193256e-001
6	209.68	2.613091e-004	3.645618e-002
7	773.75	6.527705e-005	9.096389e-003
8	2541.47	1.638221e-005	2.294711e-003

Finally, table 4 presents the statistics for the last example of risk sharing in a production economy. This model has three endogenous state variables, therefore it is more complicated to solve. However, also in this case we do not need a very fine grid to get decent levels of accuracy. Computational time increases, but it is still at tolerable levels (42 minutes with 8 grid points for each dimension). I conjecture that the performance of the algorithm can be improved by combining collocation with the Smolyak algorithm (see for example Malin et al. (2010)). In particular, Smolyak can be useful for more complicated models, since it is well known that the collocation method does not perform well for state spaces with more than 3 endogenous states variables.

6 Conclusions

The use of recursive Lagrangeans as a solution strategy is common for dynamic environment with full information, but not for private information setups. Sleet and Yeltekin (2008a) open the way for applications with privately observed shocks. This paper does the same for models with privately observed actions, and in particular proposes an algorithm which is much faster than the traditional APS technique. This methodology allows the researcher to deal with models with many states, and to calibrate simulated series to real data in a reasonable amount of time. A large class of models which are practically intractable under standard techniques can be easily addressed with the techniques discussed here.

This method has many possible applications. Given the speed, the algorithm can also be useful (as a time-saving technique) for solving those models that are tractable with traditional techniques, but computationally burdensome. Dynamic agency problems with hidden effort and hidden assets are a good example: while there is a good qualitative idea of the main predictions of this model, to the best of my knowledge a quantitative assessment in a calibrated economy is still missing. Mainly this is due to numerical difficulties. The Lagrangean approach offers a chance to overcome these limits: it is easy to calibrate models and match data even in the hidden effort, hidden assets economy. These techniques can be potentially helpful in the analysis of several issues such as e.g. consumption-saving anomalies, optimal unemployment insurance with assets accumulation or DSGE models with financial frictions.

However, the main gain of the Lagrangean method can be seen in more compli-

cated setups, which are practically intractable with current state-of-the-art algorithms. Models of repeated moral hazard with heterogeneous agents and endogenous states are a good example: they require us to solve the problem of each agent and aggregate the resulting individual optimal choices, before iterating until a general equilibrium is found. In these cases, APS techniques are unmanageable even with just two endogenous states, while with my approach it is a simple computational task. Other problems for which the Lagrangean approach has a potential advantage are optimal taxation theory in economies with hidden states and hidden actions, models of CEO compensation, and models of banking and credit markets.

References

- [1] Abraham, Á. and N. Pavoni (2008), "Principal-Agent Relationships with Hidden Borrowing and Lending: The First-Order Approach in Two Periods", mimeo, UCL
- [2] Abraham, Á. and N. Pavoni (2009), "Efficient allocations with moral hazard and hidden borrowing and lending: A recursive formulation, *Review of Economic Dynamics*, Volume 11, Issue 4, October 2008, Pages 781-803
- [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] Atkeson, A. and H. Cole (2008), "A Dynamic Theory of Optimal Capital Structure and Executive Compensation," mimeo, UCLA
- [5] Chien, Y. and H. Lustig (Forthcoming) "The Market Price of Aggregate Risk and the Wealth Distribution" *Review of Financial Studies*
- [6] Clementi, G. L., Cooley, T. and C. Wang (2006), "Stock Grants as a Commitment Device", *Journal of Economic Dynamics and Control* 30(11): 2191-2216
- [7] Clementi, G. L., Cooley, T. and S. Di Giannatale (2008a), "Total executive compensation", mimeo
- [8] Clementi, G. L., Cooley, T. and S. Di Giannatale (2008b), "A theory of firm decline", mimeo
- [9] Fernandes, A. and C. Phelan (2000), "A Recursive Formulation for Repeated Agency with History Dependence", *Journal of Economic Theory* 91(2): 223-247
- [10] Friedman, E. (1998), "Risk sharing and the dynamics of inequality", mimeo, Northwestern University
- [11] Hopenhayn, H. A. and Nicolini, J. P. (1997), "Optimal Unemployment Insurance", *Journal of Political Economy* 105(2): 412-438

- [12] Jewitt, I. (1988) "Justifying the First-Order Approach to Principal-Agent Problems," *Econometrica*, Econometric Society, vol. 56(5), pages 1177-90, September.
- [13] Judd K. (1998), "Numerical Methods in Economics", MIT Press, Cambridge (MA)
- [14] Judd K., J. Conklin and Sevin Yeltekin (2003) "Computing Supergame Equilibria", *Econometrica*, 2003 71(4): 1239- 1255.
- [15] Ke, R. (2010), "A Fixed-Point Method for Validating the First-Order Approach: Necessary and Sufficient Condition and its Implications", mimeo
- [16] Koehne, S. (2009), "The First-Order Approach to Moral Hazard Problems with Hidden Saving", mimeo, University of Mannheim
- [17] Lehnert A., Ligon E. and R. M. Townsend (1999), "Liquidity Constraints and Incentive Contracts", *Macroeconomic Dynamics* 3: 1-47.
- [18] Luenberger D. G. (1969), "Optimization by vector space methods", Wiley and Sons, New York
- [19] Marcet, A. and R. Marimon (2009), "Recursive contracts", mimeo, IAE
- [20] Mele, A. (2009), "Dynamic risk sharing and moral hazard", work in progress
- [21] Messner, M. and N. Pavoni (2004). "On the Recursive Saddle Point Method: A Note", IGIER Working Paper n. 255
- [22] Mirrlees, J. A. (1975), "The Theory of Moral Hazard and Unobservable Behaviour: Part I", published in: *The Review of Economic Studies*, Vol. 66, No. 1, Special Issue: Contracts (Jan., 1999), pp. 3-21
- [23] Paulson, A. L., Karaivanov, A. and Townsend R. M. (2006), "Distinguishing Limited Liability from Moral Hazard in a Model of Entrepreneurship", *Journal of Political Economy* 144(1): 100-144
- [24] Pavoni, N. (2007), "On optimal unemployment compensation", *Journal of Monetary Economics* 54(6): 1612-1630
- [25] Pavoni, N. (forthcoming), "Optimal Unemployment Insurance with Human Capital Depreciation and Duration Dependence", *International Economic Review*
- [26] Phelan, C. and R. M. Townsend (1991), "Computing Multi-Period, Information Constrained Equilibria", *Review of Economic Studies* 58(5): 853-881
- [27] Quadrini, V. (2004), "Investment and liquidation in renegotiation-proof contracts with moral hazard", *Journal of Monetary Economics*, 51(4): 713-751
- [28] Rogerson, W. (1985a), "Repeated Moral Hazard", *Econometrica*, 53: 69-76

- [29] Rogerson, W. (1985b), "The First-Order Approach to Principal-Agent Problems", *Econometrica*, 53 (6): 1357-1368
- [30] Sleet C. and S. Yeltekin (2003), "On the Approximation of Value Correspondences", mimeo, Carnegie Mellon University
- [31] Sleet C. and S. Yeltekin (2006), "Credibility and Endogenous Societal Discounting", *Review of Economic Dynamics* 9, 2006; 410-437.
- [32] Sleet C. and S. Yeltekin (2008a), "Solving private information models", mimeo, Carnegie Mellon University
- [33] Sleet C. and S. Yeltekin (2008b), "Politically Credible Social Insurance", *Journal of Monetary Economics* 55, 2008; 129-151
- [34] Shimer, R. and I. Werning (forthcoming), "Liquidity and insurance for the unemployed", *American Economic Review*
- [35] Spear, S. and S. Srivastava (1987), "On Repeated Moral Hazard with Discounting", *Review of Economic Studies* 54(4): 599-617
- [36] 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
- [37] Werning, I. (2001), "Repeated Moral-Hazard with Unmonitored Wealth: A Recursive First-Order Approach", mimeo, MIT
- [38] Werning, I. (2002), "Optimal Unemployment Insurance with Unobservable Savings", mimeo, MIT
- [39] Zhao, R. (2007), "Dynamic risk-sharing with two-sided moral hazard", *Journal of Economic Theory* 136: 601-640.

A Proofs

In this Appendix A, I collect the proof of Proposition 1 and the characterization of the optimal contract for the simple principal-agent model in section 2.

A.1 Proof of Proposition 1

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

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

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

is a contraction.

Proof. The space

$$M = \left\{ f : S \times \mathbb{R}^2 \longrightarrow \mathbb{R} \quad s.t. \right.$$

$$\left. \begin{array}{l} a) \quad \forall \alpha > 0 \quad f(\cdot, \alpha\theta) = \alpha f(\cdot, \theta) \\ b) \quad f(s, \cdot) \text{ is continuous and bounded} \end{array} \right\}$$

will be our candidate, with norm

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

Marcet and Marimon (2009) show that M is a nonempty complete metric space. I have to show that $T_K : M \longrightarrow M$. Notice that

$$(T_K f)(s, \theta) = \theta h_0(a^*, c^*, s) + \chi^* h_1(a^*, c^*, s) + \beta \sum_{s'} \pi(s' | s, a^*) f(s', \theta^{*'}(s'))$$

hence by Schwartz's inequality

$$\begin{aligned} \|(T_K f)(s, \theta)\| &\leq \|\theta\| \|h_0(a^*, c^*, s)\| + \max \{K, K \|\theta\|\} \|h_1(a^*, c^*, s)\| \\ &\quad + \beta \left(\max \{K, K \|\theta\|\} \left\| \frac{\pi_a(s' | s, a^*)}{\pi(s' | s, a^*)} \right\| + \|\theta\| \right) \left\| f \left(s', \frac{\theta^{*'}(s')}{\|\theta^{*'}(s')\|} \right) \right\| \end{aligned}$$

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 (a^*, c^*, χ^*) be such that

$$(T_K f)(s, \theta) = h(a^*, c^*, \theta, \chi^*, s) + \beta \sum_{s'} \pi(s' | s, a^*) f(s', \theta^{*'}(s'))$$

Then for any $\alpha > 0$ we get

$$\alpha (T_K f)(s, \theta) = \alpha \left[h(a^*, c^*, \theta, \chi^*, s) + \beta \sum_{s'} \pi(s' | s, a^*) f(s', \theta^{*'}(s')) \right]$$

Therefore

$$\begin{aligned} & h(a^*, c^*, \alpha\theta, \alpha\chi^*, s) + \beta \sum_{s'} \pi(s' | s, a^*) f(s', \alpha\theta^{*'}(s')) = \\ & = \alpha \left[h(a^*, c^*, \theta, \chi^*, s) + \beta \sum_{s'} \pi(s' | s, a^*) f(s', \theta^{*'}(s')) \right] \end{aligned}$$

Now take a generic χ , then define $\theta'_{a^*}(s') = \varphi(\alpha\theta, \chi, a^*, s')$ and $\theta'_a(s') = \varphi(\theta, \chi^*, a, s')$ for a feasible a . We can write:

$$\begin{aligned} & h(a^*, c^*, \alpha\theta, \chi, s) + \beta \sum_{s'} \pi(s' | s, a^*) f(s', \theta'_{a^*}(s')) \\ & = \alpha \left[h\left(a^*, c^*, \theta, \frac{\chi}{\alpha}, s\right) + \beta \sum_{s'} \pi(s' | s, a^*) f\left(s', \frac{\theta'_{a^*}(s')}{\alpha}\right) \right] \\ & \geq \alpha \left[h(a^*, c^*, \theta, \chi^*, s) + \beta \sum_{s'} \pi(s' | s, a^*) f(s', \theta^{*'}(s')) \right] \\ & \geq \alpha \left[h(a, c, \theta, \chi^*, s) + \beta \sum_{s'} \pi(s' | s, a) f(s', \theta'_a(s')) \right] \end{aligned}$$

and therefore

$$\begin{aligned} (T_K f)(s, \alpha\theta) &= h(a^*, c^*, \alpha\theta, \alpha\chi^*, s) + \beta \sum_{s'} \pi(s' | s, a^*) f(s', \alpha\theta^{*'}(s')) \\ &= \alpha (T_K f)(s, \theta) \end{aligned}$$

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

$$\begin{aligned} & \max_{a,c} \left\{ h(a, c, \theta, \chi, s) + \beta \sum_{s'} \pi(s' | s, a) g(s', \theta'(s')) \right\} \\ & \leq \max_{a,c} \left\{ h(a, c, \theta, \chi, s) + \beta \sum_{s'} \pi(s' | s, a) u(s', \theta'(s')) \right\} \end{aligned}$$

and then

$$\begin{aligned} & \min_{\{\chi \geq 0: \|\chi\| \leq K_\theta\}} \max_{a,c} \left\{ h(a, c, \theta, \chi, s) + \beta \sum_{s'} \pi(s' | s, a) g(s', \theta'(s')) \right\} \\ & \leq \min_{\{\chi \geq 0: \|\chi\| \leq K_\theta\}} \max_{a,c} \left\{ h(a, c, \theta, \chi, s) + \beta \sum_{s'} \pi(s' | s, a) u(s', \theta'(s')) \right\} \end{aligned}$$

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

$$\begin{aligned} & \max_{a,c} \left\{ h(a, c, \theta, \chi, s) + \beta \sum_{s'} \pi(s' | s, a) (g + k)(s', \theta'(s')) \right\} \\ &= \max_{a,c} \left\{ h(a, c, \theta, \chi, s) + \beta \sum_{s'} \pi(s' | s, a) g(s', \theta'(s')) + \beta k \right\} \\ &= \max_{a,c} \left\{ h(a, c, \theta, \chi, s) + \beta \sum_{s'} \pi(s' | s, a) g(s', \theta'(s')) \right\} + \beta k \end{aligned}$$

Hence we get

$$\begin{aligned} T_K(f + k)(s, \theta) &= \\ &= \min_{\{\chi \geq 0: \|\chi\| \leq K_\theta\}} \max_{a,c} \left\{ h(a, c, \theta, \chi, s) + \beta \sum_{s'} \pi(s' | s, a) (f + k)(s', \theta'(s')) \right\} \\ &= \min_{\{\chi \geq 0: \|\chi\| \leq K_\theta\}} \max_{a,c} \left\{ h(a, c, \theta, \chi, s) + \beta \sum_{s'} \pi(s' | s, a) f(s', \theta'(s')) \right\} + \beta k \\ &= (T_K f)(s, \theta) + \beta k \end{aligned}$$

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

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

and then

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

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

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

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. ■

A.2 Characterization of the optimal contract

In this section I show some properties of the optimal contract. These properties are the analogue, under the Lagrangean approach, of well known results in the literature. Let us go back to the problem with $\bar{\phi}^0 = 1$. We can take the first order conditions of the Lagrangean:

$$c_t(s^t) : \quad 0 = -1 + \phi_t(s^t) u_c(c_t(s^t)) \quad (25)$$

$$\begin{aligned} a_t(s^t) : \quad 0 = & -\lambda_t(s^t) v''(a_t(s^t)) - \phi_t(s^t) v'(a_t(s^t)) + \quad (26) \\ & + \sum_{j=1}^{\infty} \beta^j \sum_{s^{t+j}|s^t} \frac{\pi_a(s_{t+1} | s_t, a_t(s^t))}{\pi(s_{t+1} | s_t, a_t(s^t))} \{y(s_t) - c_t(s^t) - \lambda_{t+j}(s_{t+j}) v'(a_{t+j}(s^{t+j})) - \\ & + \phi_{t+j}(s^{t+j}) [u(c_{t+j}(s^{t+j})) - v(a_{t+j}(s^{t+j}))]\} \Pi(s^{t+j} | s^t, a^{t+j-1}(s^{t+j-1})) + \\ & + \beta \lambda_t(s^t) \sum_{s^{t+1}|s^t} \frac{\partial \left(\frac{\pi_a(\cdot)}{\pi(\cdot)} \right)}{\partial a} [u(c_{t+1}(s^{t+1})) - v(a_{t+1}(s^{t+1}))] \pi(s_{t+1} | s_t, a_t(s^t)) \end{aligned}$$

and

$$\begin{aligned} \lambda_t(s^t) : \quad 0 = & -v'(a_t(s^t)) + \sum_{j=1}^{\infty} \sum_{s^{t+j}|s^t} \frac{\pi_a(s_{t+1} | s_t, a_t(s^t))}{\pi(s_{t+1} | s_t, a_t(s^t))} \times \quad (27) \\ & \times [\beta^j [u(c_{t+j}(s^{t+j})) - v(a_{t+j}(s^{t+j}))] \Pi(s^{t+j} | s^t, a^{t+j-1}(s^{t+j-1} | s^t))] \end{aligned}$$

Lemma 1 makes clear how $\phi_t(s^t)$ incorporates the promises of the principal. From (25) 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 today²⁶.

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 (27), $0 = v'(a_t(s^t))$ 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. ■

The following Proposition characterizes the long run properties of the Pareto Negishi weight.

²⁶Thomas and Worrall (1990) prove the same property with APS techniques.

Proposition 3 $\phi_t(s^t)$ is a martingale that converges to zero.

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

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

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

$$E_t^a [\phi_{t+1} | s^t] = \phi_t(s^t) \quad (28)$$

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 (28) by using (25):

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

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, $f(\phi_{t(k)}(s^{t(k)}), \hat{s}_I) = \phi_{t(k)+1}(s^{t(k)}, \hat{s}_I)$, 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, we have 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. ■

Proposition 3 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 almost surely as $t \rightarrow \infty$ for any possible sequence of realizations of the income shock.

The result in Proposition 3 is obtained by using the law of motion of $\phi_t(s^t)$ and (25), 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.

B Hidden assets

B.1 Recursivity

Define the following generalized version of the problem:

$$\begin{aligned} W_\theta^{SWF}(s_0) = & \max_{\{a_t(s^t), c_t(s^t)\}_{t=0}^\infty \in \Gamma^{HA}} \bar{\phi}^0 \sum_{t=0}^{\infty} \sum_{s^t} \beta^t [y(s_t) - c_t(s^t)] \Pi(s^t | s_0, a^{t-1}(s^{t-1})) + \\ & + \gamma \sum_{t=0}^{\infty} \sum_{s^t} \beta^t (u(c_t(s^t)) - v(a_t(s^t))) \Pi(s^t | s_0, a^{t-1}(s^{t-1})) \\ \text{s.t. } v'(a_t(s^t)) = & \sum_{j=1}^{\infty} \beta^j \sum_{s^{t+j}|s^t} \frac{\pi_a(s_{t+1} | s_t, a_t(s^t))}{\pi(s_{t+1} | s_t, a_t(s^t))} \times \\ & \times [u(c_{t+j}(s^{t+j})) - v(a_{t+j}(s^{t+j}))] \Pi(s^{t+j} | s^t, a^{t+j-1}(s^{t+j-1} | s^t)) \\ & \quad \quad \quad \forall s^t, t \geq 0 \\ u'(c_t(s^t)) = & \beta R \sum_{s_{t+1}} u'(c_{t+1}(s^t, s_{t+1})) \pi(s_{t+1} | s_t, a_t(s^t)) \end{aligned}$$

The Lagrangean is:

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

Notice that $r(a, c, s) \equiv y(s) - c$ is uniformly bounded by debt limits, therefore there exists a lower bound $\underline{\kappa}$ such that $r(a, c, s) \geq \underline{\kappa}$. As before, we can define $\kappa < \frac{\underline{\kappa}}{1-\beta}$, $\varphi^1(\phi, \lambda, s') \equiv \phi + \lambda \frac{\pi_a(s'|s, a)}{\pi(s'|s, a)}$, $\varphi^2(\zeta, \eta, s') \equiv \eta$, $\Psi(\phi, \zeta, \lambda, \eta, s') \equiv \begin{bmatrix} \varphi^1(\phi, \lambda, s') \\ \varphi^2(\zeta, \eta, s') \end{bmatrix}$, $h_0^P(a, c, s) \equiv r(a, c, s)$, $h_1^P(a, c, s) \equiv r(a, c, s) - \kappa$, $h_0^{ICC}(a, c, s) \equiv u(c) - v(a)$, $h_1^{ICC}(a, c, s) \equiv -v'(a)$, $h_0^{EE}(a, c, s) \equiv -Ru'_c(c)$, $h_1^{EE}(a, c, s) \equiv u'_c(c)$, $\theta \equiv [\phi^0 \ \phi \ \zeta] \in \mathbb{R}^3$, $\chi \equiv [\lambda^0 \ \lambda \ \eta]$ and

$$\begin{aligned}
h(a, c, \theta, \chi, s) &\equiv \theta h_0(a, c, s) + \chi h_1(a, c, s) \\
&\equiv [\phi^0 \ \phi \ \zeta] \begin{bmatrix} h_0^P(a, c, s) \\ h_0^{ICC}(a, c, s) \\ h_0^{EE}(a, c, s) \end{bmatrix} + [\lambda^0 \ \lambda \ \eta] \begin{bmatrix} h_1^P(a, c, s) \\ h_1^{ICC}(a, c, s) \\ h_1^{EE}(a, c, s) \end{bmatrix}
\end{aligned}$$

which is homogenous of degree 1 in (θ, χ) . The Lagrangean can be written as:

$$\begin{aligned}
L_\theta(s_0, \gamma, c^\infty, a^\infty, \chi^\infty) &= \\
&= \sum_{t=0}^{\infty} \sum_{s^t} \beta^t h(a_t(s^t), c_t(s^t), \theta_t(s^t), \chi_t(s^t), s_t) \Pi(s^t | s_0, a^{t-1}(s^{t-1}))
\end{aligned}$$

where

$$\theta_{t+1}(s^t, \hat{s}) = \Psi(\theta_t(s^t), \chi_t(s^t), \hat{s}) \quad \forall \hat{s} \in S$$

$$\theta_0(s^0) = \begin{bmatrix} \bar{\phi}^0 \\ \gamma \\ 0 \end{bmatrix}$$

We can associate a saddle point functional equation to this Lagrangean

$$J(s, \theta) = \min_{\chi} \max_{a,c} \left\{ h(a, c, \theta, \chi, s) + \beta \sum_{s'} \pi(s' | s, a) J(s', \theta'(s')) \right\} \quad (29)$$

$$s.t. \quad \theta'(s') = \Psi(\theta, \chi, s') \quad \forall s'$$

The following Proposition shows that the RHS operator is a contraction mapping.

Proposition 4 *Fix an arbitrary constant $K > 0$ and let $K_{\theta} = \max\{K, K \|\theta\|\}$. The operator*

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

$$s.t. \quad \theta'(s') = \Psi(\theta, \chi, s') \quad \forall s'$$

is a contraction.

Proof. *Straightforward by repeating the steps to prove Proposition 1 in the following space of functions:*

$$M = \left\{ f : S \times \mathbb{R}^3 \longrightarrow \mathbb{R} \quad s.t. \right.$$

$$\left. \begin{array}{l} a) \quad \forall \alpha > 0 \quad f(\cdot, \alpha\theta) = \alpha f(\cdot, \theta) \\ b) \quad f(s, \cdot) \text{ is continuous and bounded} \end{array} \right\}$$

with norm

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

■

B.2 The verification procedure

No conditions are known under which the first-order approach is guaranteed to be valid in the framework with hidden effort and hidden assets. Therefore, we cannot be sure that the first-order approach delivers the correct optimal allocation: it is possible that the solution obtained does not satisfies the true incentive compatibility constraint of the original problem. However we can verify it by a simple numerical procedure similar to the one proposed by Abraham and Pavoni (2009): we remaximize the lifetime utility of the agent, by taking as given the optimal transfer scheme implied by the solution of the Pareto problem; if remaximization delivers a welfare gain to the agent, the solution obtained with the first-order approach does not satisfy incentive compatibility. Instead, if no gain is possible, then the first-order approach is valid.

We solve the following problem:

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

$$\text{s.t. } c_t^V(s^t) + b_t^V(s^t) = y(s_t) + T(s_t, \phi_t(s^t), \zeta_t(s^t)) + Rb_{t-1}^V(s^{t-1})$$

$$b_{-1} \text{ given}$$

$$\phi_{t+1}(s^t, \hat{s}) = \varphi^1(\hat{s}, \phi_t(s^t), \zeta_t(s^t)) \quad \forall \hat{s} \in S \quad \text{and} \quad \phi_0(s^0) = \gamma$$

$$\zeta_{t+1}(s^t, \hat{s}) = \varphi^2(\hat{s}, \phi_t(s^t), \zeta_t(s^t)) \quad \forall \hat{s} \in S \quad \text{and} \quad \zeta_0(s^0) = 0$$

where $T(\cdot)$, $\varphi^1(\cdot)$ and $\varphi^2(\cdot)$ are the policy functions derived from Lagrangean (19), and are exogenous from the point of view of the agent (they define the transfer policy of the principal). It is obvious that this problem is recursive in the state space (s, ϕ, ζ, b) , but notice that ϕ and ζ are exogenous states. As in Abraham and Pavoni (2009), I solve this dynamic optimization problem by value function iteration on collocation nodes with linear interpolation, to be sure I do not force the code to yield a smooth value function (this is important if the problem is not concave). Once we get the value function of the agent's problem, we can calculate the welfare gain from reoptimization with respect to the optimal allocation obtained with the first-order approach. In particular, we compare the value obtained with the verification procedure and the value implied by the Lagrangean approach: if their difference is zero (in numerical terms), then the Lagrangean first-order method delivers the solution of the original problem. As Abraham and Pavoni (2009) suggest, there can be approximation issues when comparing the two value functions²⁷, therefore a non-zero cut-off value must be carefully chosen to take into account this problem.

C Bond holdings

I show how to recover bond holdings from the solution of the Lagrangean problem, for the simplest case of a dynamic principal-agent model and for the model with hidden

²⁷Notice that we can end up with a very different accuracy in the two procedures due to hardware limitations. In general, the Lagrangean approach (in which we solve nonlinear equations) has a high degree of accuracy even with few grid points (around ten for each state variable in a rectangular grid), while the value function iteration used in the verification procedure needs many grid points to get a decent degree of approximation (say around 1000 for each state variable to get a level of accuracy of the same magnitude of the Lagrangean approach). See for example Judd (1998) for a discussion of this issue.

assets.

C.1 Repeated moral hazard

We can define bond holdings recursively as:

$$\begin{aligned}
b_t(s^t) &= -E_t^a \sum_{j=1}^{\infty} \beta^j (y_{t+j} - c_{t+j}) = \\
&= -E_t^a \sum_{j=1}^{\infty} \beta^j \{ (y_{t+j} - c_{t+j}) + \phi_{t+j} [u(c_{t+j}) - v(a_{t+j})] - \lambda_{t+j} v'(a_{t+j}) \} + \\
&+ E_t^a \sum_{j=1}^{\infty} \beta^j \{ \phi_{t+j} [u(c_{t+j}) - v(a_{t+j})] - \lambda_{t+j} v'(a_{t+j}) \} \\
&= -\beta E_t^a J(y_{t+1}, \phi_{t+1}) + E_t^a \sum_{j=1}^{\infty} \beta^j \{ \phi_{t+j} [u(c_{t+j}) - v(a_{t+j})] - \lambda_{t+j} v'(a_{t+j}) \} \\
&= -\beta E_t^a J(y_{t+1}, \phi_{t+1}) + E_t^a \sum_{j=1}^{\infty} \beta^j \{ \phi_{t+j} [u(c_{t+j}) - v(a_{t+j})] \} \\
&- E_t^a \sum_{j=1}^{\infty} \beta^j \lambda_{t+j} E_{t+1}^a \sum_{k=0}^{\infty} \beta^k \frac{\pi_a(a_{t+j+1})}{\pi(a_{t+j+1})} [u(c_{t+j+k+1}) - v(a_{t+j+k+1})] \\
&= -\beta E_t^a J(y_{t+1}, \phi_{t+1}) + E_t^a \sum_{j=1}^{\infty} \beta^j \{ \phi_{t+j} [u(c_{t+j}) - v(a_{t+j})] \} \\
&- E_t^a \sum_{j=1}^{\infty} \beta^j \lambda_{t+j} \frac{\pi_a(a_{t+j})}{\pi(a_{t+j})} U(s_{t+j+1}, \phi_{t+j+1})
\end{aligned}$$

and notice that

$$\begin{aligned}
\phi_t^* U(s_t, \phi_t^*) &= \phi_t^* [u(c_t^*) - v(a_t^*)] + \phi_t^* \beta E_t^a U(s_{t+1}, \phi_{t+1}^*) \\
&= \phi_t^* [u(c_t^*) - v(a_t^*)] + \beta E_t^a \phi_t^* \frac{\phi_{t+1}^*}{\phi_{t+1}^*} U(s_{t+1}, \phi_{t+1}^*) \\
&= \phi_t^* [u(c_t^*) - v(a_t^*)] + \beta E_t^a \frac{\phi_t^*}{\phi_{t+1}^*} \phi_{t+1}^* U(s_{t+1}, \phi_{t+1}^*) \\
&= \phi_t^* [u(c_t^*) - v(a_t^*)] + \beta E_t^a \phi_{t+1}^* U(s_{t+1}, \phi_{t+1}^*) = \\
&= \phi_t^* [u(c_t^*) - v(a_t^*)] + E_t^a \sum_{j=1}^{\infty} \beta^j \phi_{t+j}^* [u(c_{t+j}^*) - v(a_{t+j}^*)]
\end{aligned}$$

due to homogeneity of degree zero of the policy functions and of $U(s, \cdot)$. Therefore

$$\begin{aligned} b_t(s^t) &= -\beta E_t^a J(y_{t+1}, \phi_{t+1}) + \beta E_t^a \phi_{t+1} U(y_{t+1}, \phi_{t+1}) \\ &\quad - E_t^a \sum_{j=1}^{\infty} \beta^j \lambda_{t+j} \frac{\pi_a(a_{t+j})}{\pi(a_{t+j})} U(s_{t+j+1}, \phi_{t+j+1}) \end{aligned}$$

by Abel's formula $= -\beta E_t^a J(y_{t+1}, \phi_{t+1}) + \beta E_t^a \phi_{t+1} U(y_{t+1}, \phi_{t+1})$

$$\begin{aligned} &- E_t^a \sum_{j=1}^{\infty} \beta^j (\phi_{t+j+1} - \phi_{t+1}) [u(c_{t+j+1}^*) - v(a_{t+j+1}^*)] \\ &= -\beta E_t^a J(y_{t+1}, \phi_{t+1}) + \beta E_t^a \phi_{t+1} U(y_{t+1}, \phi_{t+1}) \\ &\quad - E_t^a \sum_{j=1}^{\infty} \beta^j (\phi_{t+j+1}) [u(c_{t+j+1}^*) - v(a_{t+j+1}^*)] \\ &\quad + E_t^a \sum_{j=1}^{\infty} \beta^j \phi_{t+1} [u(c_{t+j+1}^*) - v(a_{t+j+1}^*)] \\ &= -\beta E_t^a J(y_{t+1}, \phi_{t+1}) + \beta E_t^a \phi_{t+1} [U(y_{t+1}, \phi_{t+1}) + E_{t+1}^a U(y_{t+2}, \phi_{t+2})] \\ &\quad - \beta E_t^a \phi_{t+2} U(y_{t+2}, \phi_{t+2}) \\ &= -\beta E_t^a J(y_{t+1}, \phi_{t+1}) + \beta E_t^a \phi_{t+1} [U(y_{t+1}, \phi_{t+1})] \\ &\quad - E_t^a \lambda_{t+1} \frac{\pi_a(a_{t+1})}{\pi(a_{t+1})} \beta E_{t+1}^a U(y_{t+2}, \phi_{t+2}) \end{aligned}$$

which can be rewritten as

$$\begin{aligned} b_t(s^t) &= -\beta E_t^a J(y_{t+1}, \phi_{t+1}) + \beta E_t^a \phi_{t+1} U(y_{t+1}, \phi_{t+1}) \\ &\quad - E_t^a \lambda_{t+1} \frac{\pi_a(a_{t+1})}{\pi(a_{t+1})} [U(y_{t+1}, \phi_{t+1}) - u(c_{t+1}^*) + v(a_{t+1}^*)] \end{aligned}$$

where the second line is due to the optimality of the contract.

C.2 Hidden assets

Starting from the previous result, in this case we can write

$$\begin{aligned}
b_t(s^t) &= -E_t^a \sum_{j=1}^{\infty} \beta^j (y_{t+j} - c_{t+j}) = \\
&\quad -\beta E_t^a J(y_{t+1}, \phi_{t+1}) + \beta E_t^a \phi_{t+1} U(y_{t+1}, \phi_{t+1}) - \\
&\quad - E_t^a \lambda_{t+1} \frac{\pi_a(a_{t+1})}{\pi(a_{t+1})} [U(y_{t+1}, \phi_{t+1}) - u(c_{t+1}^*) + v(a_{t+1}^*)] - \\
&\quad - E_t^a \sum_{j=1}^{\infty} \beta^j [\eta_{t+j} - \beta^{-1} \zeta_{t+j}] u'(c_{t+j}) \\
&= -\beta E_t^a J(y_{t+1}, \phi_{t+1}) + \beta E_t^a \phi_{t+1} U(y_{t+1}, \phi_{t+1}) - \\
&\quad - E_t^a \lambda_{t+1} \frac{\pi_a(a_{t+1})}{\pi(a_{t+1})} [U(y_{t+1}, \phi_{t+1}) - u(c_{t+1}^*) + v(a_{t+1}^*)] \\
&\quad - \underbrace{E_t^a \sum_{j=1}^{\infty} \beta^j [\eta_{t+j} - \beta^{-1} \zeta_{t+j}] u'(c_{t+j}) - E_t^a \zeta_{t+1} u'(c_{t+1}) + E_t^a \zeta_{t+1} u'(c_{t+1})}_{=0 \text{ by definition}} \\
&= -\beta E_t^a J(y_{t+1}, \phi_{t+1}) + \beta E_t^a \phi_{t+1} U(y_{t+1}, \phi_{t+1}) - \\
&\quad - E_t^a \lambda_{t+1} \frac{\pi_a(a_{t+1})}{\pi(a_{t+1})} [U(y_{t+1}, \phi_{t+1}) - u(c_{t+1}^*) + v(a_{t+1}^*)] + \\
&\quad + E_t^a \zeta_{t+1} u'(c_{t+1})
\end{aligned}$$

D Figures

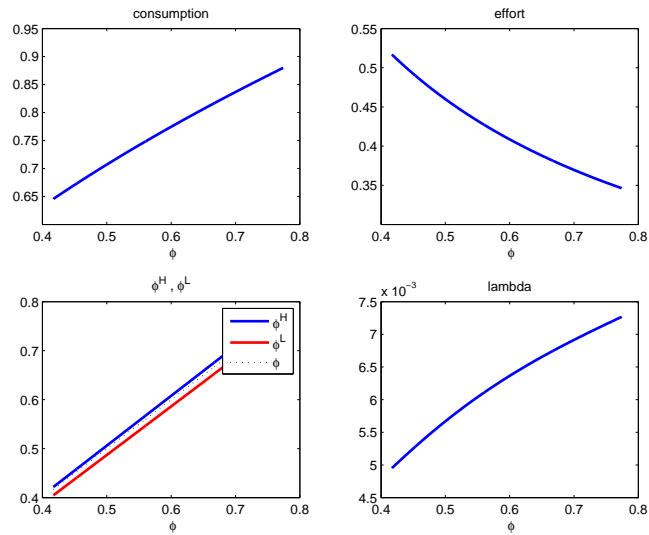


Figure 1: Pure moral hazard: policy functions

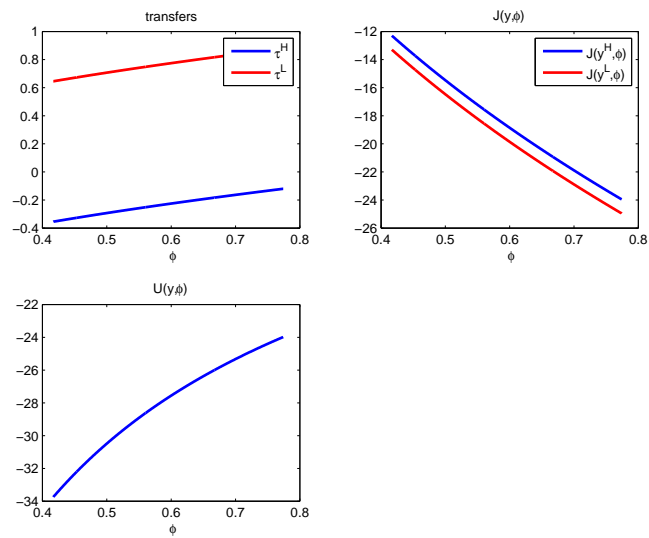


Figure 2: Pure moral hazard: policy functions (cont.)

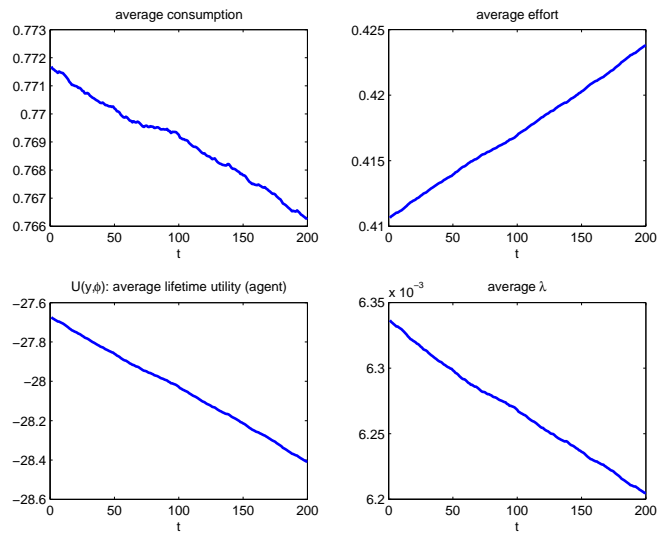


Figure 3: Pure moral hazard, average over 50000 independent simulations

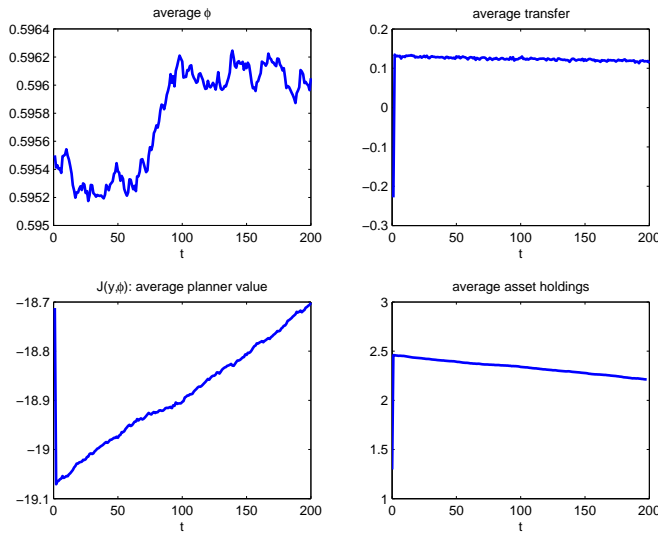


Figure 4: Pure moral hazard, average over 50000 independent simulations (cont.)




Figure 5: Pure moral hazard: Pareto frontier

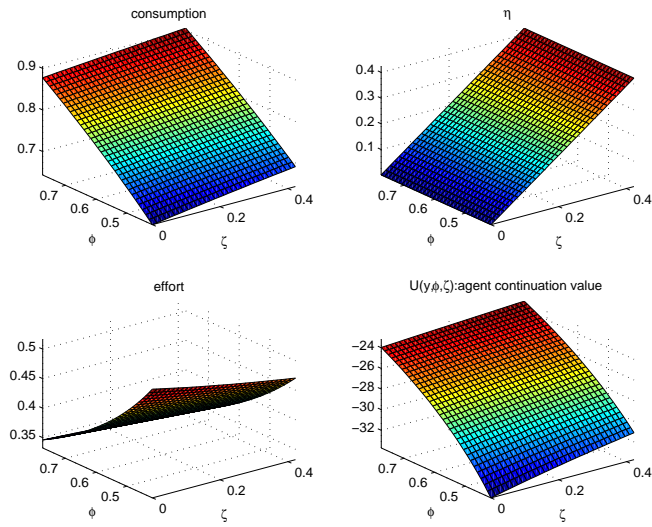


Figure 6: Moral hazard with hidden assets, policy functions




Figure 7: Moral hazard with hidden assets, policy functions (cont.)




Figure 8: Moral hazard with hidden assets, average over 50000 independent simulations




Figure 9: Moral hazard with hidden assets, average over 50000 independent simulations (cont.)




Figure 10: Pure moral hazard: Pareto frontier

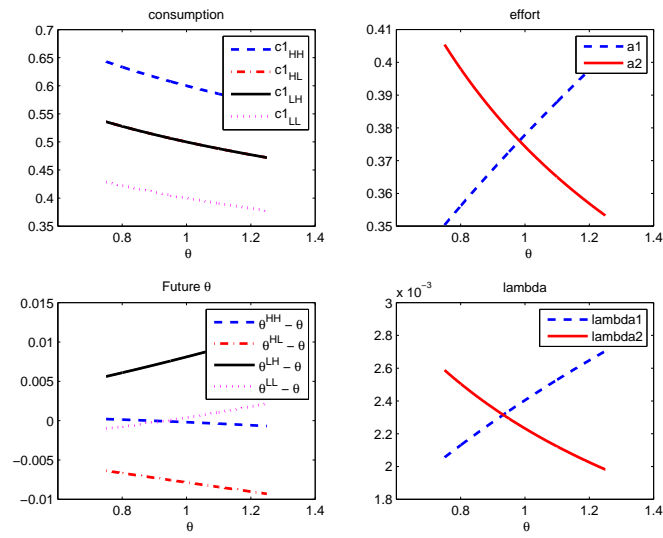


Figure 11: Risk sharing with moral hazard, policy functions (2 agents)




Figure 12: Risk sharing with moral hazard, policy functions (2 agents) (cont.)

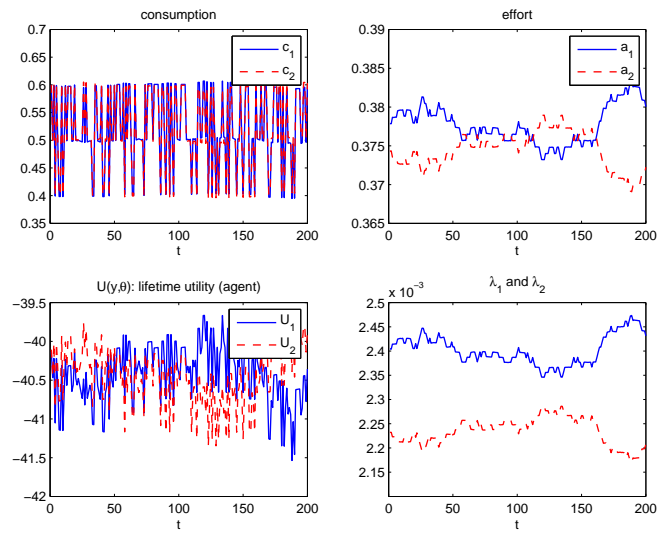


Figure 13: Risk sharing with moral hazard, sample path (2 agents)

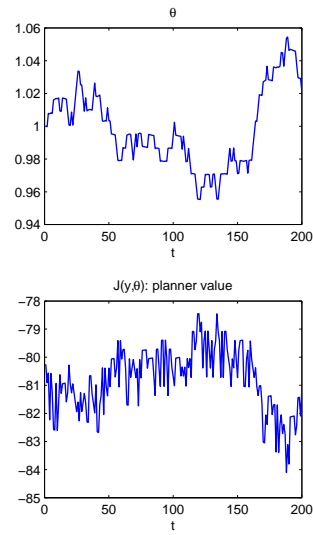


Figure 14: Risk sharing with moral hazard, sample path (2 agents) (cont.)

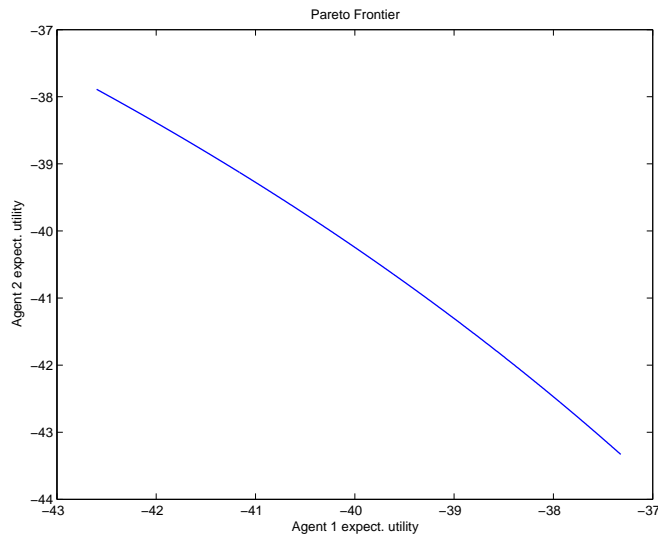


Figure 15: Risk sharing with moral hazard, Pareto frontier (2 agents)

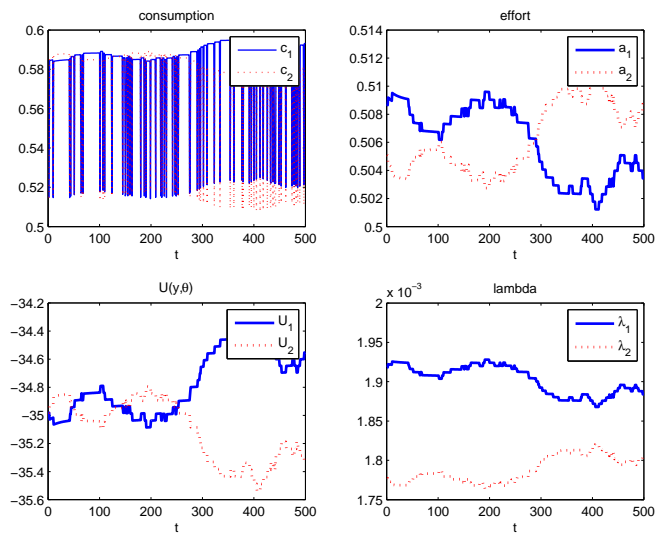


Figure 16: Production economy: sample path

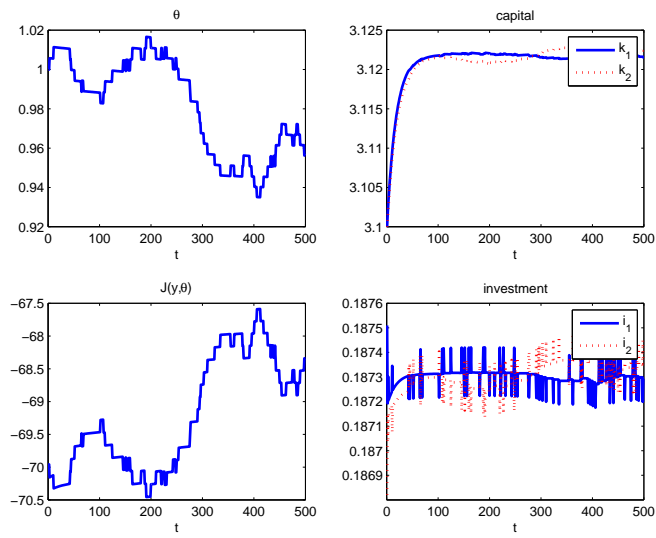


Figure 17: Production economy: sample path (cont.)




Figure 18: Production economy: average over 50000 simulations




Figure 19: Production economy: average over 50000 simulations (cont.)