

Repeated Moral Hazard and Recursive Lagrangeans

Antonio Mele*,[†]

School of Economics, University of Surrey

March 17, 2014

Abstract

This paper shows how to solve dynamic agency models by extending recursive Lagrangean techniques *à la* Marcet and Marimon (2011) 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.

Keywords: Repeated moral hazard; recursive Lagrangean; computational methods; collocation

JEL Classification Numbers: C61, C63, D82.

***Acknowledgements:** I am grateful to Albert Marcet for his suggestions and long fruitful discussions on the topic, and the editor Paul Klein and two anonymous referees for their comments that contributed to improve the paper. 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, to Chris Sleet for pointing out a mistake in a previous version of the work, and to an anonymous referee for highlighting an aspect of the methodology that was not mentioned in the paper. This paper has benefited from comments by Klaus Adam, Sofia Bauducco, Toni Braun, Filippo Brutti, Andrea Caggese, Francesco Caprioli, Martina Cecioni, Federica Dispenza, José 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.

[†]Corresponding author. Address: School of Economics, University of Surrey, GU2 7XH Guildford (Surrey), United Kingdom, **Phone:** +44 (0)1483 682 774 **email:** meleantonio@gmail.com

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 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. These models do not usually have closed form solutions, and they must be solved numerically. The most popular method is the *promised utilities approach* due to Abreu, Pearce, and Stacchetti (1990) (APS henceforth). However, this approach is very challenging, if not infeasible, for models with either several agents or many state variables. Therefore, if we use this approach, a large class of models cannot be analyzed even with numerical methods.

This paper provides a way to overcome the limits of the APS technique: under assumptions that justify the use of the first-order approach, it extends the recursive Lagrangean techniques developed in Marcat and Marimon (2011) (MM henceforth) to dynamic agency models. 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 (2008) 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 approach.

In order to illustrate the method, a dynamic principal-agent model such as the one in Spear and Srivastava (1987), where no endogenous state variables are present, is presented in section 2. As Rogerson (1985b) highlighted, these models exhibit history dependence, and therefore a standard Bellman equation representation is not possible, with the consequence of not being able to use dynamic programming techniques. The recursive Lagrangean approach allows a “recursification” of the problem, and 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. This Pareto-Negishi weight² becomes the new state variable

¹Many contributions have focused on the case in which agent’s consumption is observable (see for example Rogerson (1985b), 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 (Abraham and Pavoni (2008)); 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, Cooley, and Giannatale (2010)). Among applications, a non-exhaustive list includes unemployment insurance (Hopenhayn and Nicolini (1997), Shimer and Werning (2008), Pavoni (2007), Pavoni (2009)), executive compensation (Clementi, Cooley, and Giannatale (2010), Clementi, Cooley, and Wang (2006), Atkeson and Cole (2005)), entrepreneurship Quadrini (2004), Paulson, Townsend, and Karaivanov (2006)), credit markets (Lehnert, Ligon, and Townsend (1999)), and many more.

²Chien and Lustig (2010) 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

that "recursifies" 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 the state of nature, either "good" or "bad". The contract promises that, if tomorrow "good" 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. Hence, the contract rewards the agent for "good" realizations, and punishes him for "bad" ones. An optimal contract chooses the sequence of Pareto-Negishi weights (i.e., rewards and punishments) that implements an incentive compatible allocation.

The APS approach, instead, would "recursify" the problem by using the agent's continuation value as an auxiliary state variable. At each point in time, the agent's continuation value incorporates the past promises made by the contract to the agent, and by keeping track of those promises the principal has all the information he needs to give the right incentives to the agents in the current period. Analogously to the Lagrangean formulation with two possible contingencies, future continuation values increase if "good" is observed, and decrease otherwise. However, in order to make sure the continuation values are consistent with the original problem, a feasible set for them must be characterized before solving the dynamic program. This additional step is the main difference between the two approaches, and it is the reason why APS becomes intractable for models with many state variables. While in Spear and Srivastava (1987)'s model this set is easy to characterize, the task becomes exponentially more complicated if we increase the dimensionality of the state space.

The basic logic of the Lagrangean approach is applied to a more general class of dynamic agency problems with several agents and several observable endogenous state variables in section 3. Under the Lagrangean approach, the dynamic optimization problem has a recursive formulation based on the Pareto-Negishi weights of the agents, and the endogenous state variables from the original problem. These weights are updated in each period to enforce an incentive compatible allocation, while the endogenous states follow their own law of motion. In this more general class of models, the additional APS step makes a big difference for computational purposes. In the recursive Lagrangean approach, the social welfare function maximization problem is well defined for any non-negative real-valued weight³. Therefore, for any original dynamic agency model it is possible to characterize the optimal allocation with the recursive Lagrangean, provided that we make sure the Pareto-Negishi weights are non-negative⁴. However, with APS it might be impossible even to characterize the feasible set for the dynamic program. In fact, given an endogenous state space of dimension D for each agent and N agents, it is easy to show that the feasible set for continuation values

each agent evolves stochastically in order to keep track of occasionally binding enforcement constraints. Sleet and Yeltekin use the same terminology in their papers.

³This is also valid for the recursive Lagrangeans approach in dynamic optimization problems with full information and limited enforcement. For a discussion of this issue, see Marcet and Marimon (2011).

⁴Usually some form of the Inada condition guarantees that the weights are always positive.

has dimension $N(D + 1)$ ⁵, and its characterization might be extremely time consuming or infeasible even for small D and/or N . While in principle the Lagrangean approach can deal with models with tens of state variables and agents⁶, APS would require an unreasonable amount of CPU time in such cases, even with the efficient techniques developed by Sleet and Yeltekin (2007). Moreover, for models that have been solved with APS, the Lagrangean approach would save computational and coding time and therefore it might be a more suitable choice for calibration purposes.

There is a second and subtler reason for which the Lagrangean approach is more efficient⁷, and it has to do with the number of choice variables and the number of constraints involved in the dynamic program. Imagine there are I possible exogenous states of nature, and N agents. For each point in the state space, both approaches require computing allocations and endogenous states. However, both methods also need to compute the auxiliary state variables that recursify the problem. In APS one has to calculate $N \times I$ continuation values (one for each agent and for each future state of nature). In the Lagrangean approach instead, calculation of the N values of the current Lagrange multipliers deliver the full set of future Pareto-Negishi weights. In addition to that, any APS problem has N constraints (the promise keeping constraint for each agent, which are needed to keep track of past and future promises) which are absent in the Lagrangean approach. Summarizing, at each iteration of the dynamic programming algorithm, the APS maximization has to characterize $N(I - 1)$ additional variables under N additional constraints. As the number of agents or the number of states of nature grows, the APS program becomes much more difficult to solve than the Lagrangean one. Therefore, even if we can characterize the APS feasible set for continuation values, the Lagrangean approach is still more efficient⁸.

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 an example, a risk sharing problem with moral hazard with two agents

⁵In other words, the feasible set for continuation values is a correspondence that maps sets of the natural state space (i.e. the state variables of the original problem) into feasible continuation values' sets. For Spear and Srivastava (1987)'s model, this feasible set is an interval, and therefore its characterization is a simple, if not trivial, numerical task. However, in more complicated setups, we will need to characterize an interval for each agent and each value of the natural state space, therefore exponentially increasing the complexity of the numerical problem.

⁶The numerical algorithm proposed in this paper is based on projection techniques and therefore the ideas in Judd, Maliar, and Maliar (2010) can be applied to solve models with many agents and states. In their paper, Judd, Maliar, and Maliar (2010) solve models with hundreds of agents.

⁷I would like to thank an anonymous referee for bringing this point to my attention.

⁸As the anonymous referee pointed out, this is a feature widely exploited in the literature with full information and limited commitment.

in a production economy is solved. This model has 5 state variables, and therefore it is at the limit of reach of APS techniques. A Matlab code finds the solution in short time (a matter of minutes) with good accuracy, even with a relatively coarse grid. Therefore, models that are extremely challenging with APS become easy and fast to solve with the Lagrangean approach. This opens the doors for exploration of a large class of dynamic agency setups that were reputed too complicated even for numerical analysis.

Recent work on the Lagrangean approach includes Messner, Pavoni, and Sleet (2013) and Cole and Kubler (2012). The first proposes a more general Lagrangean approach to dynamic problems with incentive constraints that can be applied to a very large class of models, including the ones presented in this paper. The crucial differences lie in the state space used to derive recursivity, and in the techniques to characterize contractivity of the recursive dual problem. For the first, they include laws of motion of natural state variables in the Lagrangean by assigning a Lagrange multiplier to each of them, and then they only rely on Lagrange multipliers (or functions of Lagrange multipliers) for the recursive characterization. This is different from the strategy followed here (and in Marcet and Marimon (2011)) where "natural" state variables of the original problem are included together with Pareto weights derived from Lagrange multipliers, and recursivity is recovered in this enlarged state space⁹. For the second, they show contractivity by bounding the value function of the dual problem between two "sandwich" functions, and using the fact that the value function is convex. Messner, Pavoni, and Sleet (2013) make use of recent results in convex analysis with weaker assumptions on the primitives of the model, and therefore their approach is applicable to a larger class of setups, and in particular it is able to solve models that cannot be dealt with MM-like techniques. While this is an important theoretical result per se, their approach does not imply any specific computational advantage with respect to the one presented in this paper: their recursive problem has the same state space dimensionality as mine.

Cole and Kubler (2012) instead focus on solving models with flat spots on the Pareto frontier. In these cases, the MM approach fails to determine a unique policy function, even if the value of the problem is unique. Cole and Kubler (2012) propose to select the optimal policy by using an additional state variable, which is the realization of a end-of-period lottery over the extreme values on the flat part of the Pareto frontier. This approach thus combines aspects of MM and APS together. An alternative to it is Marimon, Messner, and Pavoni (2011), which extends MM method to problems lacking concavity (and therefore suffering from non-uniqueness of the policy function) by using the sign of the last period Lagrange multipliers as additional costates. In this paper I assume uniqueness of the policy function for simplicity of exposition, however it is straightforward to apply the suggestions of Marimon, Messner, and Pavoni (2011) for cases in which uniqueness fails.

⁹As an example, in a model with capital accumulation my approach would use capital as a state, while Messner, Pavoni, and Sleet (2013) would treat the law of motion for capital as a constraint and use the Lagrange multiplier associated with it as a state variable

2 An illustration with a simple dynamic agency model

Sleet and Yeltekin (2008) has shown how to extend MM's ideas to dynamic economies with hidden shocks¹⁰. However, many economic situations can be better represented as dynamic principal-agent relationship: unemployment insurance, worker's compensation, bank lending, etc. It turns out that the MM arguments apply to these economies too.

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). Theirs is an infinitely repeated version of a standard moral hazard problem, in which a risk neutral principal hires a risk averse agent to perform a task subject to uncertainty for infinitely many periods. The output of the task is observable, but the effort that the agent puts in the process is not. Therefore, the principal must structure the sequence of payments to the agent in order to preserve incentives to put high effort. To do that, the agent bears part of the risk involved in the uncertain output.

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 produces 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 period¹¹, with $\pi(s_0 = \widehat{s}_I) = 1$. Assume $\pi(\cdot)$ is twice continuously differentiable in $a_t(s^t)$ with $\frac{\pi_a(\cdot)}{\pi(\cdot)}$ uniformly 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

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

¹⁰This paper is different from Sleet and Yeltekin (2008) 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 (2008) 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 mostly 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.

¹¹Notice that shocks can be persistent.

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

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

and $c_t(s^t) \in C \subseteq \mathbb{R}_+$, with C bounded. 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:

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

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:

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

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

Constraint (3) assumes there is an exogenous value V^{out} that the agent would get if she decides not to get into a contractual relationship with the principal. Hence V^{out} can be interpreted as a reservation value for the agent in order to sign the contract.

Assume that the agent's maximization problem is a concave program. Hence it has an interior solution for effort for which first-order conditions are necessary and sufficient, and they hold with equality. We can therefore substitute the constraint (2) with the first order

conditions with respect to effort. 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, for every $s^t \in S^{t+1}$

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

Intuitively, the marginal cost of effort today (LHS) has to be equal to future expected benefits (RHS) in terms of expected future utility. Those expected benefits are defined over the entire future histories: today's effort affects directly the distribution of tomorrow's states, and indirectly (given independence across time) the distribution of future states of nature from tomorrow onwards. 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)*.

It is trivial to show that (3) must be binding in the optimum. Given this consideration, the planner's problem 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 :

$$\begin{aligned} W^{SWF}(s_0) &= \max_{\{a_t(s^t), c_t(s^t)\}_{t=0}^\infty \in \Gamma^{MH}} \sum_{t=0}^\infty \sum_{s^t} \beta^t [y(s_t) - c_t(s^t)] \Pi(s^t | s_0, a^{t-1}(s^{t-1})) + \\ &+ \gamma \sum_{t=0}^\infty \sum_{s^t} \beta^t [u(c_t(s^t)) - v(a_t(s^t))] \Pi(s^t | s_0, a^{t-1}(s^{t-1})) \end{aligned} \quad (5)$$

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

where γ is a function of V^{out} in the original problem¹², and equation (4) has been rearranged by using the definition of the cumulative probability of a history. Let $\beta^t \lambda_t(s^t) \Pi(s^t | s_0, a^{t-1}(s^{t-1}))$

¹² To 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.

be the Lagrange multiplier associated to each ICC. The Lagrangean, with some simple algebra and the use of Abel's formula¹³, yields the following expression:

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

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 the likelihood ratio and is a measure of the effect of effort on the observed output. In order to better grasp the role of the newly defined variable ϕ_t , rewrite the previous expression as:

$$\begin{aligned} \phi_{t+1}(s^t, \hat{s}) &= \phi_t(s^t) + \lambda_t(s^t) \frac{\pi_a(s_{t+1} = \hat{s} | s_t, a_t(s^t))}{\pi(s_{t+1} = \hat{s} | s_t, a_t(s^t))} \quad \forall s^t \in S^{t+1}, \forall \hat{s} \in S \quad (6) \\ \phi_0(s^0) &= \gamma \end{aligned}$$

Therefore, from (6) 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 of the ICC $\lambda_t(s^t)$ multiplied by a measure of the effect of effort on the output distribution. In order to fix ideas, imagine there are only two possible realizations of the shock, "good" and "bad", where the output is higher in the "good" state than in the "bad" state. In that case, it is easy to prove that the likelihood ratio $\frac{\pi_a(s_{t+1}|s_t, a_t(s^t))}{\pi(s_{t+1}|s_t, a_t(s^t))}$ is positive for the "good" state, and negative for the "bad" state. In other words, given that $\lambda_t(s^t)$ is always positive¹⁴, after each history the planner observes the current realization of the shock: if it is "good" (resp. "bad"), then it rewards the agent by increasing (resp. decreasing) the Pareto-Negishi weight by the marginal benefit (resp. cost) that the principal bears because of the change in the distribution of future histories that his hidden effort has caused.

Equation (6) also suggests the possibility of having a recursive formulation of the problem in an enlarged state space that includes $\phi_t(s^t)$. If this intuition is right, then we can apply standard dynamic programming methods for this optimization problem. In the next section, this is formalized in a more general setup with several agents and endogenous observable states. Recursivity of the Lagrangean is proven for this wider class of models that nest the

¹³Abel's formula states that $\sum_{t=0}^{\infty} a_t (\sum_{s=0}^{\infty} b_{t+s}) = \sum_{t=0}^{\infty} (\sum_{s=0}^t a_s) b_t$ and notice that, by defining $\mu_t \equiv (\sum_{s=0}^t a_s)$ we can write $\mu_{t+1} = \mu_t + a_t$.

¹⁴This fact is a well known result for an interior solution due to Rogerson (1985a).

setup presented in this section as a special case. Given recursivity, the Lagrangean problem is easy to solve: if well behaved, we can use Lagrangean first order conditions to characterize the optimal contract. In other words, under assumptions that justify the substitution of the original incentive compatibility constraint (2) with the agent's first order conditions with respect to effort, this dynamic agency problem can be easily solved with standard computational techniques. In fact, this is a problem with just one endogenous state variable (ϕ_t) and one shock, and therefore at the same level of difficulty of the simplest stochastic growth model.

3 Dynamic economies with moral hazard

The same line of reasoning can be applied to a more general class of dynamic agency models with several agents and several endogenous state variables. This section describes the class of models of interest, and derives a recursive formulation of the contractual problem. Notably, this class of models is suitable for heterogeneous agents and for general equilibrium analysis. For this purpose, we are going to assume that all the assumptions in MM are satisfied. In the following, when needed, other assumptions on the primitives of the model will be specified.

There are N agents¹⁵ 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 closed 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 $a^t(s^t) \equiv (a_0(s^0), a_1(s^1), \dots, a_t(s^t))$. Notice that a change in the action profile $a_t(s^t)$ affects all the cumulated probabilities of histories that follow the change. In other words, if agent i changes his action $a_{it}(s^t)$, then the distribution of states of nature for any subsequent history $\{s^{t+j}\}_{j=1}^{\infty}$ is shifted. In this assumption rests the intertemporal nature of the moral hazard problem.

Probabilities $\pi^i(s_{i,t+1} | s_{it}, a_{it}(s^t))$ are continuously 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 (i.e., $\frac{\pi_a^i(s_{i,t+1} | s_{it}, a_{it}(s^t))}{\pi^i(s_{i,t+1} | s_{it}, a_{it}(s^t))}$) is uniformly bounded with $\left\| \frac{\pi_a^i(s_{i,t+1} | s_{it}, a_{it}(s^t))}{\pi^i(s_{i,t+1} | s_{it}, a_{it}(s^t))} \right\| \leq$

¹⁵Extension of the arguments for the case with a continuum of agents is straightforward and available upon request. For expositional purposes, it is easier to present the case with finite number of agents, since all measure theoretic details are avoided.

D for some $D \in R_+$. 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:

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

This is a general and standard formulation for the law of motion of endogenous states in many dynamic macroeconomic models with several agents and idiosyncratic shocks. In particular, the assumption is that only variables under the control of agent i , or its own idiosyncratic shock, can affect the sequence x_{it} .

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 is ¹⁶:

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

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:

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

Let $\omega \equiv \{\omega_i\}_{i=1}^N \in R^N$ be a vector of initial Pareto-Negishi weights. The contractual problem is therefore

$$P(x_0, s_0) = \max_{\mathcal{W} \in \Gamma^{GT}} \left\{ \sum_{i=1}^N \sum_{t=0}^{\infty} \beta^t \sum_{s^t} \omega_i r^i(\varsigma_{it}, a_{it}, x_{it}, s_t) \Pi(s^t | s_0, a^{t-1}) \right\} \quad (7)$$

$$s.t. \quad \{a_{it}\}_{t=0}^\infty \in \arg \max_{\{\tilde{a}_{it}\}_{t=0}^\infty} \sum_{t=0}^{\infty} \beta^t \sum_{s^t} r^i(\varsigma_{it}, \tilde{a}_{it}, x_{it}, s_t) \Pi(s^t | s_0, a^{t-1}) \quad \forall i \quad (8)$$

The optimal contract maximizes the weighted sum of individual lifetime payoffs. The constraints (8) make sure that the contract keeps into account that, given a contract, agents choose hidden actions' profiles that maximize their payoffs, i.e. they optimally respond to the incentives that the contract gives them.

¹⁶Constraints 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.

Call the program (7)-(8) *the original problem*. If each agent's optimization problem is well-behaved, the sequence of constraints (8) can be characterized by the corresponding first-order conditions. In that case, it is possible to use the agent's first-order conditions as constraints in the dynamic contract problem. This solution strategy is commonly known in the literature as the *first-order approach* (FOA). For the simple principal-agent model in Section 2 there are well known sufficient conditions that guarantee that the problem with first-order conditions as constraints is equivalent to the original problem, and therefore delivers the same solution. It is easy to see that in the more general framework presented here, those same conditions are also sufficient, and therefore for the rest of this section assume that Rogerson (1985a) conditions of monotone likelihood ratio (MLRC) and convexity of the distribution (CDFC) are satisfied for each agent¹⁷.

Since FOA is valid, we can use the first-order conditions of the agents' problems with respect to hidden actions as incentive compatibility constraints:

$$r_a^i(\varsigma_{it}, a_{it}, x_{it}, s_t) + \sum_{j=1}^{\infty} \sum_{s^{t+j}} \beta^j \frac{\pi_a^i(s_{i,t+1} | s_{it}, a_{it})}{\pi^i(s_{i,t+1} | s_{it}, a_{it})} \times \\ \times r^i(\varsigma_{i,t+j}, a_{i,t+j}, x_{i,t+j}, s_{t+j}) \Pi(s^{t+j} | s^{t+j-1}, a^{t+j-1}) = 0 \quad \forall i = 1, \dots, N, \forall s^t \in S^{t+1} \quad (9)$$

where $r_a^i(\varsigma_{it}, a_{it}, x_{it}, s_t)$ is the derivative of the per-period payoff function with respect to the hidden action. The intuition behind equation (9) is straightforward: the current cost (in terms of lower instantaneous payoff) that a marginal increase in the hidden action implies must be compensated by a marginal increase in the benefits from it. Those benefits are obtained by the higher likelihood of future "good" realizations of the state of nature (since higher a shifts the distribution of the states of nature towards better outcomes for any future history), which increases the expected discounted lifetime utility from $t+1$ onwards. Hence, the constrained efficient allocation is the solution of the following maximization problem:

$$P(s_0) = \max_{\mathcal{W} \in \Gamma^{GT}} \left\{ \sum_{i=1}^N \sum_{t=0}^{\infty} \beta^t \sum_{s^t} \omega_i r^i(\varsigma_{it}, a_{it}, x_{it}, s_t) \Pi(s^t | s_0, a^{t-1}) \right\} \\ s.t. \quad (9)$$

Let $\beta^t \lambda_{it}(s^t) \Pi(s^t | s_0, a^{t-1})$ be the Lagrange multiplier for the incentive-compatibility constraint (9) of agent i . Substitute for the resource constraint and use Abel's formula to write the Lagrangean as:

$$L(s_0, \omega, \mathcal{W}, \lambda^\infty) = \sum_{i=1}^N \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \left\{ \phi_{it} r^i(\varsigma_{it}, a_{it}, x_{it}, s_t) + \right. \\ \left. + \lambda_{it} r_a^i(\varsigma_{it}, a_{it}, x_{it}, s_t) \right\} \Pi(s^t | s_0, a^{t-1}) \quad (10)$$

¹⁷For static principal-agent models, 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. Both Rogerson's and Jewitt's conditions are sufficient for dynamic agency setups with observable endogenous states. Ke (2013) suggests a fixed-point condition that justifies the first-order approach in static environments, which can potentially also be used in dynamic settings.

where, for any i and any $s^t \in S^{t+1}$,

$$\begin{aligned} x_{i,t+1}(s^t, s_{t+1}) &= \ell^i(x_{it}(s^t), \varsigma_{it}(s^t), s_{i,t+1}) \\ \phi_{i,t+1}(s^t, s_{t+1}) &= \phi_{it}(s^t) + \lambda_{it}(s^t) \frac{\pi_a^i(s_{i,t+1} | s_{it}, a_{it}(s^t))}{\pi^i(s_{i,t+1} | s_{it}, a_{it}(s^t))} \end{aligned} \quad (11)$$

$$\phi_{i0}(s_0) = \omega_i, \quad x_{i0} \text{ given} \quad (12)$$

The newly defined variables $\phi_{it}(s^t)$, $i = 1, \dots, N$, are endogenously evolving Pareto-Negishi weights which are optimally chosen by the planner to implement an incentive compatible allocation and they summarize the contract's (history-dependent) promises for each agent. The intuition is the same as in Section 2. Given the weight inherited from previous period, the optimal contract rewards or punishes the agent by increasing or decreasing its weight in the next period by an amount that internalizes the externality produced by moral hazard. Thus, the weight summarizes the history of past promises (either rewards or punishments), and it is then natural to consider $\phi_t(s^t)$ as a probable sufficient statistic that can make the problem recursive: given the current value for the Pareto weight, the planner can choose current allocations and future Pareto weights that solve the maximization. A formal proof of this intuition is obtained in the next subsection.

3.1 Recursivity

Notice that the Lagrangean problem (10) looks like a social welfare function maximization. The only difference with a standard one is the fact that Pareto weights change through time. This feature can be exploited to show that the original problem has a recursive formulation.

By duality theory (see for example Luenberger (1969)), a solution of the original problem corresponds to a saddle point of the Lagrangean, i.e. the contract \mathcal{W}^* is a solution for the original problem if there exist a sequence $\lambda_\infty^* \equiv \{\lambda_t^*(s^t) \quad \forall s^t \in S^{t+1}\}_{t=0}^\infty$ of Lagrange multipliers such that $(\mathcal{W}^*, \lambda_\infty^*)$ satisfy:

$$L(s_0, \omega, \mathcal{W}', \lambda_\infty^*) \leq L(s_0, \omega, \mathcal{W}^*, \lambda_\infty^*) \leq L(s_0, \omega, \mathcal{W}^*, \lambda_\infty') \quad \forall \mathcal{W}', \forall \lambda_\infty'$$

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 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 , the vector of endogenous state variables $x_t(s^t)$ and the Pareto-Negishi weights $\phi_t(s^t)$.

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 homogeneous 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

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

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

We need to show that all solutions of the Lagrangean have a recursive structure, and that can be solved with standard dynamic programming techniques. This is done in two steps. Firstly, Proposition 1 proves that, under the assumption that there is a unique allocation that solves the original problem, a particular functional equation (the *saddle point functional equation*) associated with the Lagrangean satisfies the assumptions of the Contraction Mapping Theorem¹⁸. 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. The latter is a trivial application of MM (Theorems 3 and 4) and therefore the proof is omitted¹⁹.

The saddle point functional equation associated with the Lagrangean is

$$\begin{aligned} J(s, \phi, x) &= \min_{\lambda} \max_{\varsigma, a} \left\{ h(\varsigma, a, x, \phi, \lambda, s) + \beta \sum_{s'} \pi(s' | s, a) J(s', \phi'(s'), x'(s')) \right\} \quad (13) \\ \text{s.t.} \quad x'(s') &= \ell(x, \varsigma, s') \\ \phi'(s') &= \varphi(\phi, \lambda, a, s') \quad \forall s' \end{aligned}$$

¹⁸In 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. Messner 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 infeasible. This happens more generally when the Pareto frontier of the problem is not strictly concave, i.e. it presents flat spots, and therefore even if the planner value is uniquely determined, agents' individual continuation values are in fact indeterminate. Cole and Kubler (2012) get around this problem by using an additional state variable, which is the realization of an end-of-period lottery over the extreme values on the flat part of the Pareto frontier, hence combining together aspects of MM and APS methodologies. Marimon, Messner, and Pavoni (2011) generalize the arguments of MM for policy correspondence by using the sign of the last positive Lagrange multiplier as an additional co-state variable, and similar ideas can be used in the repeated moral hazard setup.

¹⁹When the solution for allocations is unique, then Marcat and Marimon (2011) shows that the ‘‘promise-keeping’’ equations of the form

$$U^i(s, x, \phi) = r^i(\varsigma_i, a_i, x_i, s) + \beta \sum_{s'} \pi(s' | s, a) U^i(s', x', \phi'(s')) \quad \forall i, (s, x, \phi)$$

hold for optimality, where $U^i(s, x, \phi)$ is the continuation value of agent i at state (s, x, ϕ) . If allocations are not uniquely determined, those equations must be imposed for optimality. Although for the proof of Proposition 1 we assume that allocations are uniquely determined and therefore we do not need to impose these equations, we include them in the numerical algorithm.

Candidate value functions will be taken from the following space of functions:

$$M = \{f : S \times \mathbb{R}^N \times X \longrightarrow \mathbb{R} \quad s.t. \begin{array}{l} a) \quad f(s, \cdot, \cdot) \text{ is continuous, and } f(s, \cdot, \phi) \text{ is bounded when } \|\phi\| \leq 1 \\ b) \quad f(s, x, \cdot) \text{ is convex and homogeneous of degree 1} \end{array} \}$$

with norm

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

We first need to establish a few preliminary results, which are simple applications of Lemmata in Marcet and Marimon (2011) and therefore the proof is omitted:

1. M is a complete metric space;
2. Functions from the set M are differentiable in ϕ (thanks to Euler's formula), and can be represented as $f(s, \phi, x) = \phi U(s, x, \phi)$, where $U(s, \phi, x) \equiv \{U^i(s, \phi, x)\}_{i=1}^N$ are the agents' continuation values. Therefore, agents' continuation values are homogeneous of degree zero in ϕ ;
3. Optimal Lagrange multipliers are uniformly bounded, i.e., there exists a K such that, for any (s, ϕ, x) , we have $\|\lambda^*(s, \phi, x)\| \leq K \|\phi\|$. Together with the assumption of bounded likelihood ratios (i.e., there exists a $D \in R_+$ for which $\left\| \frac{\pi_a^i(s'_i | s'_i, a_i)}{\pi^i(s'_i | s'_i, a_i)} \right\| \leq D$), this implies that, for any s' , we have $\|\phi'(s')\| \leq (1 + KD)\|\phi\|$, i.e. the feasibility correspondence generated by the law of motion for Pareto weights is compact-valued.

Proposition 1 shows that the operator on the RHS of (13) is a contraction. 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: Blackwell's conditions are satisfied, and therefore the operator is a contraction.

Proposition 1. *The operator*

$$(Tf)(s, \phi, x) \equiv \min_{\lambda} \max_{\varsigma, a} \left\{ h(\varsigma, a, x, \phi, \lambda, s) + \beta \sum_{s'} \pi(s' | s, a) f(s', \phi'(s'), x'(s')) \right\}$$

$$s.t. \quad x'(s') = \ell(x, \varsigma, s')$$

$$\phi'(s') = \varphi(\phi, \lambda, a, s') \quad \forall s'$$

is a contraction.

Proof. See Appendix A □

This result therefore establishes the existence of a unique value function that solves the saddle point functional equation. Theorem 3 and 4 in MM guarantee that, provided that the optimal policy correspondence for allocations is single-valued, a recursive solution of Problem (13) 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 states $(s, \phi, x) \in S \times \mathbb{R}^N \times X$.

Notice that this problem has $N(m + 1)$ endogenous state variables. However, the value function of the problem is homogeneous of degree one in the vector of endogenous weights ϕ and individual continuation values for each agent i are homogeneous of degree zero with respect to the vector of endogenous weights ϕ . This fact implies:

$$\frac{1}{\phi_1} J(s, \phi_1, \dots, \phi_N, x) = J\left(s, 1, \frac{\phi_2}{\phi_1}, \dots, \frac{\phi_N}{\phi_1}, x\right) \equiv \tilde{J}\left(s, \frac{\phi_2}{\phi_1}, \dots, \frac{\phi_N}{\phi_1}, x\right) \quad (14)$$

$$U^i(s, \phi_1, \dots, \phi_N, x) = U^i\left(s, 1, \frac{\phi_2}{\phi_1}, \dots, \frac{\phi_N}{\phi_1}, x\right) \equiv \tilde{U}^i\left(s, \frac{\phi_2}{\phi_1}, \dots, \frac{\phi_N}{\phi_1}, x\right) \quad \forall i \quad (15)$$

therefore the dimension of the state space can be reduced to $N(m + 1) - 1$. This result is sometimes helpful in computational applications, and it is exploited in the example in the next section.

4 A numerical example

In this section, I describe the algorithm and I provide one computed example of a production economy with two agents. This example features two agents and two endogenous state variables (capital for each agent). It is therefore a challenging problem with the APS approach, and I am not aware of any paper that solves a similarly complicated model. However, it turns out that it is a relatively easy model for the Lagrangean approach. The proposed algorithm can find a quite accurate solution in less than 20 minutes on a state-of-the-art laptop.

4.1 The algorithm

Here I describe the algorithm for the general setup in section 3. Define a generic state of the economy as $\hat{s} \in S$ where $S \equiv \times_{i=1}^N S_i$. For simplicity, assume the shocks are not persistent, 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. The algorithm therefore is the following:

1. Fix $\omega_i, i = 1, \dots, N$ and define a discrete grid $G \subset S \times X \times \mathbf{R}_+^N$ for natural states and costates.

2. Approximate policy functions for allocations ς , hidden actions a , Lagrange multipliers λ , 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 $(s, x, \phi) \in G$, use a non-linear solver²⁰ to solve for the Lagrangean first order conditions and the following equations for the continuation value U^i and the value function J :

$$U^i(s, x, \phi) = r^i(\varsigma_i, a_i, x_i, s) + \beta \sum_{\hat{s}} \pi(\hat{s} | a) U^i(\hat{s}, x', \phi'(\hat{s})) \quad (16)$$

$$J(s, x, \phi) = h(\varsigma, a, x, \phi, \lambda, s) + \beta \sum_{\hat{s}} \pi(\hat{s} | a) J(\hat{s}, x', \phi'(\hat{s})) \quad (17)$$

I use the Miranda-Fackler Compecon toolbox for function approximation. 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 of the interpolants. In general, it is a good idea to start from the solution of a simpler model. Homotopy methods help if the latter is not enough. The algorithm is coded in Matlab²¹.

4.2 Risk sharing in a production economy

The proposed example is a production economy in which two agents (or countries, depending on the interpretation) must share their income. For simplicity of the notation, I omit the dependence on the history, hence a generic variable w_{it} is to be read as $w_{it}(s^t)$. In each period, they produce output y_{it} with their own capital k_{it} . However, the production function is subject to idiosyncratic productivity shocks, the distribution of which is affected by unobservable effort a_{it} . The production function is

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

where A_{it} is the i.i.d. productivity shock (hence, in terms of the notation of section 3, we have $s_t = (A_{1t} A_{2t})$). I assume that each A_{it} can assume only two possible values $\{A_i^L, A_i^H\}$ and the probability of the high realization is given by $\pi(a_{it}) = a_{it}^{\nu_i}$, $a_{it} \in (0, 1)$, where ν_i is a parameter that can possibly differ between agents. Agents accumulate capital according to the following standard law of motion:

$$k_{i,t+1} = A_{it} k_{it}^{\rho_i} - c_{it} + (1 - \delta_i) k_{it}$$

²⁰I use a version of the Broyden algorithm coded by Michael Reiter.

²¹The basic code can be downloaded from my website (where the interested reader can find other examples included in a previous version of the paper) or I can send it by email.

The utility function is given by

$$u(c_{it}) - v(a_{it}) = \frac{c_{it}^{1-\sigma_i}}{1-\sigma_i} - \alpha_i a_{it}^{\varepsilon_i}$$

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

We can now define the Pareto problem for the planner. The feasibility constraint becomes

$$\sum_{i=1}^2 (c_{it} + k_{i,t+1} - (1 - \delta_i) k_{it}) \leq \sum_{i=1}^2 A_{it} f(k_{it})$$

The Pareto-constrained allocation can be found by solving:

$$P(s_0, k_0) = \max_{\{c_{it}, a_{it}, k_{it}\}_{i=1,2}} \left\{ E_t^a \sum_{i=1}^2 \sum_{t=0}^{\infty} \beta^t \omega_i [u(c_{it}) - v(a_{it})] \right\}$$

$$s.t. \quad v'(a_{it}(s^t)) = E_{t+1}^a \sum_{j=1}^{\infty} \beta^j \frac{\pi_{a_i}(s_{i,t+1} | s_{it}, a_{it})}{\pi(s_{i,t+1} | s_{it}, a_{it})} [u(c_{i,t+j}) - v(a_{i,t+j})] \quad (18)$$

$$\sum_{i=1}^2 (c_{it} + k_{i,t+1} - (1 - \delta_i) k_{it}) \leq \sum_{i=1}^2 A_{it} f(k_{it})$$

where the expectational operator E_t^a uses the probability distribution defined by $\pi(a_{it})$. We can now write down the Lagrangean of this problem. We first substitute the feasibility constraint in the objective function. Let $\lambda_{it}(s^t)$ be the Lagrange multiplier of (18) for agent i , then the Lagrangean is:

$$L(\zeta^\infty, \lambda^\infty, k^\infty, \phi^\infty) =$$

$$= \sum_{i=1}^2 \sum_{t=0}^{\infty} \sum_{s^t} \beta^t \{ \phi_{it}(s^t) [u(c_{it}(s^t)) - v(a_{it}(s^t))] -$$

$$- \lambda_{it}(s^t) v'(a_{it}(s^t)) \} \Pi(s^t | s_0, a^{t-1}(s^{t-1}))$$

where

$$\phi_{i,t+1}(s^t, s_{t+1}) = \phi_{it}(s^t) + \lambda_{it}(s^t) \frac{\pi_{a_i}(s_{i,t+1} | s_{it}, a_{it}(s^t))}{\pi(s_{i,t+1} | s_{it}, a_{it}(s^t))}$$

$$\phi_{i0}(s_0) = \omega_i$$

and recursivity is therefore obtained by using ϕ_{it} and k_{it} as state variables. Therefore the state variables are 5 (the 4 endogenous states plus the state of nature). However, it is possible

to reduce the number of states by using the homogeneity properties of the value functions. Remember from equation (14) in section 3 that the planner's value function is homogeneous of degree one. Moreover, given that the continuation value U^i of each agent is the derivative of the planner's value function with respect to ϕ_i , then U^i is homogeneous of degree zero by Euler's theorem with respect to the Pareto-Negishi weights. Hence it is possible to derive the Lagrangean first-order conditions and rewrite them as a function of the Pareto weights' ratio $\Phi_t \equiv \frac{\phi_{2t}}{\phi_{1t}}$, thus reducing the number of states to 3 endogenous states and one state of nature (the steps to obtain this simplification are available upon request).

4.3 Computational speed and accuracy

The following table 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, using orthogonal collocation with Chebichev polynomials. Let M be the number of grid points in each dimension of the state space, e.g. with three endogenous states the grid has a total of M^3 grid points. The Miranda-Fackler Compecon toolbox automatically chooses the degree of the polynomials equal to the number of grid points, and the gridpoints to be the zeros of the Chebichev polynomials.

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 64-bit laptop with a processor Intel Core i7-2640 at 2.80 GHz and 6 GB RAM. The accuracy of the approximated solution can be tested by defining a large grid (with $46^3 = 97336$ 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 table, there are two statistics that measure accuracy: the maximum error and the norm of the error vector. These statistics are calculated at the absolute value, and at the free-unit value. The latter is calculated by using the Bellman equation of each agent, and defining the unit-free Bellman equation error BEE_i for agent i as:

$$BEE \equiv \frac{\left[(1 - \sigma) (v(a_i) - \beta \sum_{s'} U^i(\Phi', k', s'))^{\frac{1}{1-\sigma}} + U^i(\Phi, k, s) \right]}{c_i} - 1 \quad (19)$$

The interpretation in terms of consumption is then straightforward: if $BEE = .01$, then the solution is making a mistake of 0.01 for every unit of consumption. In general, free-unit and absolute errors are similar and very small.

Table 1: Speed and Accuracy

Grid	Time (sec)	Max Error	Norm(Error)	Max Error (FU)	Norm Error (FU)
2	2.56	2.718581e-02	4.103847e+00	7.179348e-03	9.106260e-01
4	9.92	1.317713e-04	2.097053e-02	3.351783e-05	4.317211e-03
6	39.20	8.078817e-07	1.250894e-04	2.380975e-07	2.259080e-05
7	87.66	4.455317e-07	2.805273e-05	1.421670e-07	4.145154e-06
8	209.29	4.563646e-07	2.558593e-05	1.206049e-07	3.468939e-06
9	486.82	8.332295e-07	4.427910e-05	1.328316e-07	4.171927e-06
10	1152.13	3.854335e-07	2.926761e-05	6.606006e-08	3.322667e-06

Table 1 shows that it is possible to obtain a good accuracy (of the order of less than 10^{-6}) even with few grid points. Computational time is very low (less than 20 minutes with 10 grid points per dimension). The performance of the algorithm could be plausibly improved by combining collocation with the Smolyak algorithm (see for example Malin, Krueger, and Kubler (2011)). 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 state variables. Another useful technique that can help in large dimensional problems is the endogenous grid method of Judd, Maliar, and Maliar (2010), which can be applied to problems with more than 2 agents. Judd, Maliar, and Maliar (2010) show how to choose a non-rectangular grid in the ergodic set of the model at hand, dramatically reducing the number of grid points needed for a good accuracy. They produce examples of risk sharing models with hundred of agents and state variables, solved with collocation techniques.

5 Discussion

In this section I provide a more detailed comparison of APS and the Lagrangean approach for dynamic agency problems. In particular, the two methodologies differ in the way in which they make the problem recursive and the computational complexity for finding a solution of the recursive problem. I then discuss the applicability of the Lagrangean approach.

5.1 Keeping promises in APS vs MM

In order to be able to compare the two methodologies, APS is shortly described in the following lines. 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:

$$\begin{aligned}
P\left(\{U_i, x_i\}_{i=1, \dots, N}, s\right) &= \max_{\left\{\{\varsigma_i, a_i^*, \{U^i(s'), x'_i(s')\}_{s' \in S}\}_{i=1, \dots, N}\right\}} \\
&\left\{ \left[\sum_i \omega_i r^i(\varsigma_i, a_i^*, x_i, s) \right] + \beta \sum_{s'} \pi(s' | s, a^*) P\left(\{U^i(s'), x'_i(s')\}_{i=1, \dots, N}, s'\right) \right\} \\
s.t. \quad r^i(\varsigma_i, a_i^*, x_i, s) + \beta \sum_{s'} \pi(s' | s, a^*) U^i(s') &= U_i \quad i = 1, \dots, N \tag{20}
\end{aligned}$$

$$a_i^* = \arg \max_{a_i \in A_i} \left\{ r^i(\varsigma_i, a_i, x_i, s) + \beta \sum_{s'} \pi(s' | s, (a_i, a_{-i}^*)) U^i(s') \right\} \quad i = 1, \dots, N \tag{21}$$

$$\begin{aligned}
x'_i(s') &= \ell^i(x_i, \varsigma_i, s') \quad i = 1, \dots, N, p(x, \varsigma, a^*, s) \geq 0 \quad \forall s \in S, \\
\left\{ \{U^i(s'), x'_i(s')\}_{i=1, \dots, N} \right\}_{s' \in S} &\in \mathcal{U} \tag{22}
\end{aligned}$$

where (20) is the promise-keeping constraint, (21) is the incentive compatibility constraint.

APS therefore “recursifies” the problem by using the continuation values of the agents. The planner promises a higher continuation values for higher realizations of the state of nature. In order to make sure that promises are kept, the promise keeping (which is absent in the original problem) constraint must be satisfied: this constraint plays the role of the law of motion for the continuation values. For simplicity, imagine again that we only have two possible states of nature for each agent: “good” and “bad”. Therefore, given the current continuation value, the optimal contract will promise a future continuation value that is higher than the current one if tomorrow is “good”, and lower if tomorrow is “bad”. This seems analogous to the recursive Lagrangean approach, where the promises were implemented by using higher or lower Pareto-Negishi weights. Intuitively, Pareto-Negishi weights and continuation values are two ways of fulfilling the same purpose: rewarding or punishing the agents for their effort. However, even if the two methodologies look similar in the way they implement incentive compatibility, there are remarkable differences for finding a solution.

5.2 Computational differences

The set \mathcal{U} is the feasible set for continuation values. APS showed that it can be obtained by iterating on a particular set-valued operator defined as

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

This operator transforms sets W in new sets W' . The contribution of APS is to show that the sequence of sets W_t generated by iterating on the operator B converges to the set of feasible

payoffs of the original problem. Hence, when applying APS technique to a dynamic agency problem, one 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 vector of Pareto-Negishi weights in R_+^N . Therefore, because of this additional step in the promised utilities method, the Lagrangean approach is generally simpler than the APS one.

Moreover, the dimensionality of the feasible set \mathcal{U} is given by $N(m + 1)$, where N is the number of agents, and m is the dimensionality of the state space for each agent. The computational task therefore is to pick every possible combination of continuation values and states that belong to an initial set W_0 , apply to them the operator $B(W_0)$ repeatedly until convergence. While there are various algorithms that perform this task in an efficient way (e.g. Sleet and Yeltekin (2007)), the problem is a nightmare even for low N and m . It is easy to see that for the simple model in section 2 the correspondence \mathcal{U} is actually one interval. However, in the more general framework presented in section 3 with N agents with endogenous states $x \in X$, there are N different intervals for each point of the natural state space X , i.e. this feasible set for continuation values is the multidimensional graph of a correspondence. Computing this correspondence is already a formidable task for the case $N + m = 3$. Hence, 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. While the Lagrangean approach can potentially be applied to frameworks with hundreds of agents and state variables, the APS technique is much more limited and I am not aware, to the best of my knowledge, of any paper that solves a case with $N + m = 4$.

The second difference concerns the number of choice variables and the constraints involved in the dynamic program²². Imagine there are I possible states of nature. Then, in both approaches, for any point in the state space one needs to compute allocations ς , hidden actions a and endogenous states x' (s'). However, in APS $N \times I$ continuation values (one for each agent and for each future state of nature) must be found, while with my approach it is sufficient to calculate the N values in λ in order to have the full set of future Pareto weights worked out. Moreover, in APS the maximization problem involves N more constraints (the promise keeping constraint (20) for each agent). Hence, in APS the maximization step of each iteration of the dynamic programming algorithm involves $N(I - 1)$ additional variables to be computed and N additional constraints with respect to the Lagrangean approach. It is then easy to see that this feature becomes relevant in terms of computational complexity and speed as N and I become large.

5.3 Applicability

There are two main caveats to the recursive Lagrangean method for dynamic agency problems: the applicability of the technique, and its local nature.

²²I would like to thank an anonymous referee that pointed out this crucial advantage to me.

5.3.1 Hidden endogenous states

Proposition 1 refers to cases in which all the endogenous state variables are observable. However, there are many situations that are better modelled with unobservable endogenous states. One important contribution in this respect is the work of Abraham and Pavoni (2008) about dynamic agency with hidden savings, where they use APS techniques and the first order approach. In principle, it is possible to follow the same general idea of combining the first-order approach and the recursive Lagrangean for these models: solve the agent's maximization problem with respect to all unobservable 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 big 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 states. 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 (2008)²³.

Recursive Lagrangean techniques must therefore be used with caution for these models. Even if an ex-post verification algorithm can tell if the solution satisfies incentive compatibility, it should be thought as a tool for validating the use of FOA when one has already a reasonable expectation that FOA would work. It is indeed a risky strategy to start using the Lagrangean approach just to discover that the FOA is not valid in that particular application or, even worse, with that particular calibration that matches data.

Nevertheless, the recent work of Abraham, Koehne, and Pavoni (2011) on two-period repeated moral hazard with hidden savings suggests a proof strategy for the validity of FOA in multiperiod models that could potentially be pursued for specific applications.

5.3.2 Local vs global

The second caveat is that the numerical algorithm is a local method, since it is based on Kuhn-Tucker necessary conditions. One can check if, starting from different initial conditions, the algorithm always delivers the same solution. This indeed is a standard check in dynamic optimal taxation literature. However, such a procedure does not guarantee that the solution is a global optimum.

This problem can be addressed if one is ready to compromise speed with global results. The suggested algorithm is not the only way to find a solution. The main benefit of the algorithm is its speed and the simple implementation, however the big advantage of the recursive Lagrangean approach (the absence of a characterization step for the feasible set of costate variables) does not depend on it. If one has a strong reason to believe that the Kuhn-Tucker conditions are not sufficient, the saddle point can be found by iterating over the value function

²³A previous version of the paper presented the model with hidden savings and solved it with Lagrangean techniques. The Matlab code is available upon request.

and using a global optimization procedure (e.g., direct search or genetic algorithms). While this computational strategy would lose the gains in terms of speed, it still retains the advantage of not needing a characterization of the costates' feasible set. With modern parallel computing tools, the cost to pay might not be excessively high.

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 (2008) 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. 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 complicated 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 effort and several assets, models of CEO compensation, and models of banking and credit markets.

References

- ABRAHAM, A., S. KOEHNE, AND N. PAVONI (2011): "On the first-order approach in principal-agent models with hidden borrowing and lending," *Journal of Economic Theory*, 146(4), 1331–1361.
- ABRAHAM, A., AND N. PAVONI (2008): "Efficient Allocations with Moral Hazard and Hidden Borrowing and Lending: A Recursive Formulation," *Review of Economic Dynamics*, 11(4), 781–803.
- ABREU, D., D. PEARCE, AND E. STACCHETTI (1990): "Toward a theory of discounted

- repeated games with imperfect monitoring,” *Econometrica: Journal of the Econometric Society*, pp. 1041–1063.
- ATKESON, A., AND H. COLE (2005): “A Dynamic Theory of Optimal Capital Structure and Executive Compensation,” NBER Working Papers 11083, National Bureau of Economic Research, Inc.
- CHIEN, Y., AND H. LUSTIG (2010): “The Market Price of Aggregate Risk and the Wealth Distribution,” *Review of Financial Studies*, 23(4), 1596–1650.
- CLEMENTI, G. L., T. COOLEY, AND S. D. GIANNATALE (2010): “A Theory of Firm Decline,” *Review of Economic Dynamics*, 13(4), 861–885.
- CLEMENTI, G. L., T. F. COOLEY, AND C. WANG (2006): “Stock grants as a commitment device,” *Journal of Economic Dynamics and Control*, 30(11), 2191–2216.
- COLE, H., AND F. KUBLER (2012): “Recursive Contracts, Lotteries and Weakly Concave Pareto Sets,” *Review of Economic Dynamics*, 15(4), 479–500.
- FERNANDES, A., AND C. PHELAN (2000): “A Recursive Formulation for Repeated Agency with History Dependence,” *Journal of Economic Theory*, 91(2), 223–247.
- FRIEDMAN, E. (1998): “Risk Sharing and the Dynamics of Inequality,” Discussion Papers 1235, Northwestern University, Center for Mathematical Studies in Economics and Management Science.
- HOPENHAYN, H. A., AND J. P. NICOLINI (1997): “Optimal Unemployment Insurance,” *Journal of Political Economy*, 105(2), 412–38.
- JEWITT, I. (1988): “Justifying the First-Order Approach to Principal-Agent Problems,” *Econometrica*, 56(5), 1177–90.
- JUDD, K. L. (1998): *Numerical Methods in Economics*, vol. 1 of *MIT Press Books*. The MIT Press.
- JUDD, K. L., L. MALIAR, AND S. MALIAR (2010): “A Cluster-Grid Projection Method: Solving Problems with High Dimensionality,” NBER Working Papers 15965, National Bureau of Economic Research, Inc.
- KE, R. (2013): “A Fixed-Point Method for Validating the First-Order Approach,” working paper.
- KOCHERLAKOTA, N. (2004): “Figuring out the Impact of Hidden Savings on Optimal Unemployment Insurance,” *Review of Economic Dynamics*, 7(3), 541–554.
- LEHNERT, A., E. LIGON, AND R. M. TOWNSEND (1999): “Liquidity Constraints And Incentive Contracts,” *Macroeconomic Dynamics*, 3(01), 1–47.

- LUENBERGER, D. (1969): *Optimization by Vector Space Methods*, Professional Series. Wiley.
- MALIN, B. A., D. KRUEGER, AND F. KUBLER (2011): “Solving the multi-country real business cycle model using a Smolyak-collocation method,” *Journal of Economic Dynamics and Control*, 35(2), 229–239.
- MARCET, A., AND R. MARIMON (2011): “Recursive Contracts,” Discussion paper.
- MARIMON, R., M. MESSNER, AND N. PAVONI (2011): “Solving recursive contracts with non-unique solutions,” working paper.
- MESSNER, M., AND N. PAVONI (2004): “On the Recursive Saddle Point Method,” Discussion paper.
- MESSNER, M., N. PAVONI, AND C. SLEET (2013): “The Dual Approach to Recursive Optimization: Theory and Examples,” Discussion paper.
- PAULSON, A. L., R. M. TOWNSEND, AND A. KARAIVANOV (2006): “Distinguishing Limited Liability from Moral Hazard in a Model of Entrepreneurship,” *Journal of Political Economy*, 114(1), 100–144.
- PAVONI, N. (2007): “On optimal unemployment compensation,” *Journal of Monetary Economics*, 54(6), 1612–1630.
- (2009): “Optimal Unemployment Insurance, With Human Capital Depreciation, And Duration Dependence,” *International Economic Review*, 50(2), 323–362.
- PHELAN, C., AND R. M. TOWNSEND (1991): “Computing Multi-period, Information-Constrained Optima,” *Review of Economic Studies*, 58(5), 853–81.
- QUADRINI, V. (2004): “Investment and liquidation in renegotiation-proof contracts with moral hazard,” *Journal of Monetary Economics*, 51(4), 713–751.
- ROGERSON, W. P. (1985a): “The First-Order Approach to Principal-Agent Problems,” *Econometrica*, 53(6), 1357–67.
- (1985b): “Repeated Moral Hazard,” *Econometrica*, 53(1), 69–76.
- SHIMER, R., AND I. WERNING (2008): “Liquidity and Insurance for the Unemployed,” *American Economic Review*, 98(5), 1922–42.
- SLEET, C., AND S. YELTEKIN (2007): “On the Approximation of Value Correspondences,” Discussion paper.
- (2008): “Solving private information models,” working paper, Carnegie Mellon University.

SPEAR, S. E., AND S. SRIVASTAVA (1987): “On Repeated Moral Hazard with Discounting,” *Review of Economic Studies*, 54(4), 599–617.

THOMAS, J., AND T. WORRALL (1990): “Income fluctuation and asymmetric information: An example of a repeated principal-agent problem,” *Journal of Economic Theory*, 51(2), 367–390.

ZHAO, R. R. (2007): “Dynamic risk-sharing with two-sided moral hazard,” *Journal of Economic Theory*, 136(1), 601–640.

A Proof of Proposition 1

The proof strategy is to show that the operator T maps M into itself, and then that Blackwell’s conditions of discounting and monotonicity are satisfied for T .

In order to show that $T : M \rightarrow M$, notice that

$$(Tf)(s, \phi, x) = \phi h_0(\zeta^*, a^*, x^*, s) + \lambda^* h_1(\zeta^*, a^*, x^*, s) + \beta \sum_{s'} \pi(s' | s, a^*) f\left(s', \phi^*(s'), x^*(s')\right)$$

hence by Schwartz’s inequality

$$\begin{aligned} \|(Tf)(s, \phi, x)\| &\leq \|\phi\| \|h_0(\zeta^*, a^*, x^*, s)\| + K \|\phi\| \|h_1(\zeta^*, a^*, x^*, s)\| \\ &\quad + \beta \left(K \|\phi\| \left\| \frac{\pi_a(s' | s, a^*)}{\pi(s' | s, a^*)} \right\| + \|\phi\| \right) \left\| f\left(s', \frac{\phi^*(s')}{\|\phi^*(s')\|}, x^*(s')\right) \right\| \end{aligned}$$

and therefore $(Tf)(s, \phi, x)$ is bounded. A generalized Maximum Principle argument gives continuity of $(Tf)(s, \phi, x)$. To check for homogeneity properties, let $(\zeta^*, a^*, x^*, \lambda^*)$ be such that

$$(Tf)(s, \phi, x) = h(\zeta^*, a^*, x^*, \phi, \lambda^*, s) + \beta \sum_{s'} \pi(s' | s, a^*) f\left(s', \phi^*(s'), x^*(s')\right)$$

Then for any $\alpha > 0$ we get

$$\alpha (Tf)(s, \phi, x) = \alpha \left[h(\zeta^*, a^*, x^*, \phi, \lambda^*, s) + \beta \sum_{s'} \pi(s' | s, a^*) f\left(s', \phi^*(s'), x^*(s')\right) \right]$$

Therefore

$$\begin{aligned} &h(\zeta^*, a^*, x^*, \alpha\phi, \alpha\lambda^*, s) + \beta \sum_{s'} \pi(s' | s, a^*) f\left(s', \alpha\phi^*(s'), x^*(s')\right) = \\ &= \alpha \left[h(\zeta^*, a^*, x^*, \phi, \lambda^*, s) + \beta \sum_{s'} \pi(s' | s, a^*) f\left(s', \phi^*(s'), x^*(s')\right) \right] \end{aligned} \quad (23)$$

Now take a generic λ and notice that we can write:

$$\begin{aligned}
& h(\zeta^*, a^*, x^*, \alpha\phi, \lambda, s) + \beta \sum_{s'} \pi(s' | s, a^*) f\left(s', \varphi(\alpha\phi, \lambda, a^*, s'), x^{*'}(s')\right) \\
&= \alpha \left[h\left(\zeta^*, a^*, x^*, \phi, \frac{\lambda}{\alpha}, s\right) + \beta \sum_{s'} \pi(s' | s, a^*) f\left(s', \frac{\varphi(\alpha\phi, \lambda, a^*, s')}{\alpha}, x^{*'}(s')\right) \right] \text{ (by homogeneity)} \\
&\geq \alpha \left[h(\zeta^*, a^*, x^*, \phi, \lambda^*, s) + \beta \sum_{s'} \pi(s' | s, a^*) f\left(s', \phi^{*'}(s'), x^{*'}(s')\right) \right] \text{ (by definition of saddle point)} \\
&\geq \alpha \left[h(\zeta, a, x, \phi, \lambda^*, s) + \beta \sum_{s'} \pi(s' | s, a) f\left(s', \varphi(\phi, \lambda^*, a, s'), x'(s')\right) \right]
\end{aligned}$$

and using (23)

$$\begin{aligned}
(Tf)(s, \alpha\phi, x) &= h(\zeta^*, a^*, x^*, \alpha\phi, \alpha\lambda^*, s) + \beta \sum_{s'} \pi(s' | s, a^*) f\left(s', \alpha\phi^{*'}(s'), x^{*'}(s')\right) \\
&= \alpha (Tf)(s, \phi, x)
\end{aligned}$$

and therefore the operator preserves the homogeneity properties. Convexity is then obvious.

To see monotonicity, let $g, u \in M$ such that $g \leq u$. Therefore

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

and then

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

which implies $(Tg)(s, \phi, x) \leq (Tu)(s, \phi, x)$.

To see discounting, let $k \in \mathbb{R}_+$. Therefore, for any $\lambda > 0$:

$$\begin{aligned}
& \max_{\zeta, a} \left\{ h(\zeta, a, x, \phi, \lambda, s) + \beta \sum_{s'} \pi(s' | s, a) (g + k)(s', \phi'(s'), x'(s')) \right\} \\
&= \max_{\zeta, a} \left\{ h(\zeta, a, x, \phi, \lambda, s) + \beta \sum_{s'} \pi(s' | s, a) g(s', \phi'(s'), x'(s')) + \beta k \sum_{s'} \pi(s' | s, a) \|\varphi(\phi, \lambda, a, s')\| \right\} \\
&\leq \max_{\zeta, a} \left\{ h(\zeta, a, x, \phi, \lambda, s) + \beta \sum_{s'} \pi(s' | s, a) g(s', \phi'(s'), x'(s')) \right\} + \beta k (1 + KD) \|\phi\|
\end{aligned}$$

For any feasible (ς, a) , let then $\widehat{\lambda}(\varsigma, a)$ be the solution of

$$\begin{aligned} & \min_{\{\lambda: \|\lambda\| \leq K\|\phi\|\}} \left\{ h(\varsigma, a, x, \phi, \lambda, s) + \beta \sum_{s'} \pi(s' | s, a) (g + k)(s', \phi'(s'), x'(s')) \right\} \\ &= h(\varsigma, a, x, \phi, \widehat{\lambda}(\varsigma, a), s) + \beta \sum_{s'} \pi(s' | s, a) g(s', \varphi(\phi, \widehat{\lambda}(\varsigma, a), a, s'), x'(s')) + \\ & \quad + \beta k \sum_{s'} \pi(s' | s, a) \|\varphi(\phi, \widehat{\lambda}(\varsigma, a), a, s')\| \\ &\leq h(\varsigma, a, x, \phi, \widehat{\lambda}(\varsigma, a), s) + \beta \sum_{s'} \pi(s' | s, a) g(s', \varphi(\phi, \widehat{\lambda}(\varsigma, a), a, s'), x'(s')) + \\ & \quad + \beta k (1 + KD) \|\phi\| \end{aligned}$$

For any feasible (ς, a) , let then $\lambda^*(\varsigma, a)$ be the solution of

$$\begin{aligned} & \min_{\{\lambda: \|\lambda\| \leq K\|\phi\|\}} \left\{ h(\varsigma, a, x, \phi, \lambda, s) + \beta \sum_{s'} \pi(s' | s, a) (g)(s', \phi'(s'), x'(s')) \right\} \\ &= h(\varsigma, a, x, \phi, \lambda^*(\varsigma, a), s) + \beta \sum_{s'} \pi(s' | s, a) g(s', \varphi(\phi, \lambda^*(\varsigma, a), a, s'), x'(s')) \end{aligned}$$

Therefore we can write

$$\begin{aligned} & h(\varsigma, a, x, \phi, \widehat{\lambda}(\varsigma, a), s) + \beta \sum_{s'} \pi(s' | s, a) g(s', \varphi(\phi, \widehat{\lambda}(\varsigma, a), a, s'), x'(s')) + \\ & \quad + \beta k (1 + KD) \|\phi\| \\ &\leq h(\varsigma, a, x, \phi, \lambda^*(\varsigma, a), s) + \beta \sum_{s'} \pi(s' | s, a) g(s', \varphi(\phi, \lambda^*(\varsigma, a), a, s'), x'(s')) + \\ & \quad + \beta k (1 + KD) \|\phi\| \end{aligned}$$

Hence we get

$$\begin{aligned} & T(f + k)(s, \phi, x) = \\ &= \min_{\lambda} \max_{\varsigma, a} \left\{ h(\varsigma, a, x, \phi, \lambda, s) + \beta \sum_{s'} \pi(s' | s, a) (f + k)(s', \phi'(s'), x'(s')) \right\} \\ &\leq h(\varsigma^*, a^*, x, \phi, \lambda^*(\varsigma^*, a^*), s) + \beta \sum_{s'} \pi(s' | s, a^*) f(s', \varphi(\phi, \lambda^*(\varsigma^*, a^*), a, s'), x'(s')) + \\ & \quad + \beta k (1 + KD) \|\phi\| \\ &= T(f)(s, \phi, x) + \beta k (1 + KD) \|\phi\| \end{aligned}$$

and by homogeneity we can always choose a ϕ such that $\|\phi\| \leq (1 + KD)^{-1}$, finally obtaining $T(f + k) \leq Tf + \beta k$. Now it is possible to use the previous results to show the contraction

property for the operator T . In order to see this, let $f, g \in M$. By homogeneity, we get

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

and then

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

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

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

which implies finally

$$\|Tf - Tg\| \leq \beta \|f - g\|$$

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