

Dynamic Principal–Agent Models*

Philipp Renner	Karl Schmedders
Hoover Institution	Universität Zürich and
Stanford University	Swiss Finance Institute
phrenner@gmail.com	karl.schmedders@business.uzh.ch

April 12, 2016

Abstract

This paper contributes to the theoretical and numerical analysis of discrete time dynamic principal–agent problems with continuous choice sets. We first provide a new and simplified proof for the recursive reformulation of the sequential dynamic principal–agent relationship. Next we prove the existence of a unique solution for the principal’s value function, which solves the dynamic programming problem in the recursive formulation, by showing that the Bellman operator is a contraction mapping. Therefore, the theorem also provides a convergence result for the value function iteration. To compute a solution for the problem we have to solve a collection of static principal–agent problems at each iteration. Under the assumption that the agent’s expected utility is a rational function of his action, we can transform the bi-level optimization problem into a standard nonlinear program (NLP). We can then solve these nonlinear problems with a standard NLP solver. The final results of our solution method are numerical approximations of the policy and value functions for the dynamic principal–agent model. We illustrate our solution method by solving variations of two prominent social planning models from the economics literature.

Keywords: Optimal unemployment tax, principal–agent model, repeated moral hazard.

JEL codes: C63, D80, D82.

*We thank Ken Judd, Tom Sargent, Chris Sleet, and Sevin Yeltekin for insightful discussions. We are also grateful to audiences at the University of Zurich, the Fall 2015 Midwest Economic Theory Meeting at Penn State University, and the University of Lancaster for helpful comments. We are indebted to Dave Brooks for a thorough editorial review of an earlier version. Philipp Renner and Karl Schmedders gratefully acknowledge financial support from the Swiss National Science Foundation (SNSF) under grant 148769 and Swiss Finance Institute, respectively.

1 Introduction

This paper contributes to the theoretical and numerical analysis of discrete time dynamic principal–agent problems with continuous choice sets. We first provide a new and simplified proof for the recursive reformulation of the sequential principal–agent relationship, which was first derived by Spear and Srivastava (1987). This recursive reformulation requires a one-dimensional state variable, the continuation utility for the agent. By providing a revised proof, we show that we can eliminate a redundant constraint in their description of the problem. Next we prove the existence of a unique solution for the principal’s value function, which solves the dynamic programming problem in the recursive formulation, by providing a contraction mapping theorem. We then show how to solve this dynamic programming problem via standard value function iteration. Within each iteration, this dynamic programming approach requires us to solve a collection of static principal–agent problems, each of which is a difficult bi-level optimization problem. Under the assumption that the agent’s expected utility is a rational function of his action, we can avoid the classical first-order approach and instead transform the bi-level optimization problem into a standard nonlinear program (NLP) using the polynomial approach of Renner and Schmedders (2015). We can solve these nonlinear problems with a standard NLP solver. The final results of our solution method are numerical approximations of the policy and value functions for the dynamic principal–agent model. We illustrate our solution method by solving two prominent models from the economics literature.

In a repeated principal–agent relationship, the history of the interaction between the principal and the agent generally affects the optimal contract. That is, the optimal contract in any period depends on the entire history of the relationship; see Lambert (1983) and Rogerson (1985). Clearly, for infinite-horizon principal–agent problems, this history dependence seemingly renders a sequential description of the optimal contract over time impossible. The fundamental insight of Spear and Srivastava (1987) is that the infinite-horizon problem is equivalent to a static variational problem. Specifically, they show, for a model without hidden state variables, that the sequential problem allows for a recursive reformulation with a one-dimensional state variable, the agent’s continuation utility. In their description they require three constraints, one of which involves the optimal policy for the continuation payoff. This makes a computational treatment using this formula difficult, if not impossible. We provide a revised formulation, which now requires solving a dynamic programming (DP) recursion with two constraints. The first constraint is the incentive-compatibility condition for the agent; the second constraint requires the agent to receive an expected discounted utility equal to the promised continuation utility.

The revised DP reformulation is still a rather daunting problem. It is a so-called bi-level optimization problem, since one of the constraints, the incentive-compatibility condition, is

itself an optimization problem. Spear and Srivastava (1987) assume (but do not prove) that under some strong simplifying assumptions the first-order approach (FOA) for static principal–agent problems also applies to their DP problem. Therefore, they can replace the agent’s utility maximization problem by a first-order condition. As a result, the DP problem is no longer a bi-level optimization problem and instead becomes a “normal” NLP. The first-order conditions for this NLP enable Spear and Srivastava (1987) to derive some characteristics of the optimal contract.

The FOA requires very restrictive assumptions on the probability distribution of outcomes, which fail to hold in many economic applications even for static problems. To the best of our knowledge, the conjecture of Spear and Srivastava (1987) that their assumptions are sufficient for an application of the FOA to dynamic models has not been proven. A more widely applicable solution approach for dynamic principal–agent problems is obviously desirable.

In this paper, we present a new solution method for the analysis of infinitely repeated moral hazard principal–agent problems. The new approach only requires the agent’s expected utility to be a rational function of the action. This assumption enables us to extend the polynomial approach of Renner and Schmedders (2015) for static principal–agent models to dynamic models. Simply put, we can numerically solve the DP problem by combining the standard dynamic programming technique of value function iteration, see Judd (1998, Chapter 12), with the polynomial solution method of Renner and Schmedders (2015) for static principal–agent models. The polynomial method allows us, similarly to the first-order approach, to transform the principal’s DP problem from a bi-level optimization problem into a standard constrained DP problem. And so a standard DP method, such as value function iteration, suffices to solve the infinite-horizon principal–agent problem. If the agent’s expected utility function is not a rational function, we can still apply this new approach after first approximating the expected utility by a polynomial. We demonstrate the applicability of our new method by solving two prominent dynamic principal–agent problems from the economics literature. First, we analyze a continuous action version of the social planning problem in Phelan and Townsend (1991). In this problem, the agent’s utility function is not rational in the effort; however, we can easily approximate it by a low-degree polynomial. The qualitative properties of the resulting value function are similar to those of the surplus function in Phelan and Townsend (1991) for the original model with discrete actions. Second, we study the unemployment insurance design problem of Hopenhayn and Nicolini (1997). We show that the optimal policy—that is, the level of unemployment benefits, is decreasing over time. We also consider a version of this model in which the principal has to provide a minimum payoff to the agent in every period. We find that in this situation the optimal policy may feature a region in which the policy is constant over time.

In this paper we restrict our attention to infinite-horizon principal–agent relationships in which the principal can completely control the agent’s consumption. The agent does not have (unobservable) access to a credit market for saving or borrowing and also does not possess a storage technology. We impose this greatly simplifying assumption for two reasons. First, there are very influential economic applications, such as those of Phelan and Townsend (1991) and Hopenhayn and Nicolini (1997), which satisfy this assumption. Second, repeated moral hazard problems with hidden states pose significant computational difficulties that extend far beyond the scope of the current paper. And while several computational methods for such general problems have been suggested in the economics literature, all of them exhibit serious limitations of one kind or another. Werning (2001) proposes a first-order approach for repeated moral hazard problems with an unmonitored wealth of the agent. However, Kocherlakota (2004) argues that there are no known sufficient conditions that ensure the validity of the FOA for problems with hidden savings. In fact, for an enhanced version of the Hopenhayn and Nicolini (1997) unemployment insurance model with hidden savings and linear disutility of search effort, he can show analytically that the FOA must fail. Doepke and Townsend (2006) present a recursive solution method for dynamic principal–agent environments with hidden states. However, they cannot allow for continuous choice sets and instead must restrict them to being finite. Abraham and Pavoni (2008) present a first-order approach for a model with continuous savings decisions. While they cannot guarantee that the approach finds a correct solution, they can verify *ex post* whether the obtained solution is in fact incentive compatible. In sum, to the best of our knowledge, no reliable solution method for infinitely repeated principal–agent relationships with continuous choice sets and hidden states exists. The development of such a method is a task for the future but beyond the scope of the present paper. Nevertheless we believe that the present paper is a significant step toward a future development of theoretical and numerical methods for such models with hidden states.

The remainder of this paper is organized as follows. In Section 2 we describe the dynamic principal–agent model, derive a recursive reformulation, and introduce the solution method for the recursive problem. Along the way we point out some incorrect statements in Spear and Srivastava (1987). In Section 3 we solve a continuous action version of the social planning problem in Phelan and Townsend (1991). Section 4 presents a solution to the design problem of an unemployment insurance system presented in Hopenhayn and Nicolini (1997). Section 5 concludes. The Appendix provides additional proofs and results.

2 The Model

In this section, we describe the dynamic principal–agent model in the spirit of Spear and Srivastava (1987).

2.1 The Dynamic principal–agent Problem

We consider an infinitely repeated principal–agent relationship in discrete time. Time is indexed by $t = 1, 2, \dots$. In each period t , the agent chooses an action (“effort level”) a_t from a set A . The principal cannot monitor the agent’s actions but only the resulting output. After the agent chooses his action the output y_t received by the principal at time t is an element of the sample space Y , a probability space. In many cases Y is chosen to be a discrete subset of the reals. Let $f(\cdot|a)$ be a parameterized time-invariant probability measure on the set of outputs Y . In each period t , the principal pays the agent a compensation (e.g., “wage”) $c_t \in C$ with C being the set of possible compensations. The timing within a period is as follows: first the agent chooses his action a , then the output is drawn from Y , and finally the agent receives compensation c from the principal. The principal has a Bernoulli utility function over her income, $v : \mathbb{R} \rightarrow \mathbb{R}$. For example, if the principal receives the output y_t and pays the compensation c_t , then she receives utility $v(y_t - c_t)$ in period t . The agent has a Bernoulli utility function over actions and compensations given by $u : A \times C \rightarrow \mathbb{R}$. Both the principal and the agent discount the future. For the analysis in this paper, we impose the following assumptions:

Assumption 1.

1. *The set A of actions is a non-empty, compact subset of a finite-dimensional Euclidean space, $A \subset \mathbb{R}^L$.*
2. *The set Y of outputs is a finite set of cardinality N .¹*
3. *The set C of compensations is non-empty and compact, $C = [0, \bar{c}]$.*
4. *For all $y \in Y$, $f(y|\cdot) : A \rightarrow [0, 1]$ is a continuous function.*
5. *For all $a \in A$, $\sum_{y \in Y} f(y|a) = 1$.*
6. *The Bernoulli utility functions $v : \mathbb{R} \rightarrow \mathbb{R}$ and $u : A \times C \rightarrow \mathbb{R}$ are continuous.*

¹We can extend our analysis to infinite output sets. However, doing so adds the additional complication of having to deal with integrability conditions for the involved functions. For simplicity we, therefore, assume the presence of only finitely many output values.

7. The principal and the agent have a common discount factor $\beta \in (0, 1)$.²

In repeated principal–agent relationships, the endogenous optimal choices of the principal and the agent at time t depend on the entire previous history of their relationship. This history dependence has been known (at least) since the work of Lambert (1983) and Rogerson (1985). Spear and Srivastava (1987) show that if we restrict ourselves to pure strategies for the two agents, then the history of output before time t is sufficient to describe the optimal choices at time t . For this reason, we define the output history at time t to be the $(t - 1)$ -tuple (an “ordered list”) $h^t = (y_1, y_2, \dots, y_{t-1})$ for $t \geq 2$ and $h^1 = ()$. We define the set of possible histories at time t by $H^t = Y^{t-1} = Y \times Y \times \dots \times Y$. In period t the agent chooses his action before the output in that period is realized. Therefore, his strategy is a function on the domain H^t , so $a_t : H^t \rightarrow A$. The principal’s compensation strategy is a function of both the history up to period t and the output in that period, so her strategy is a function $c_t : H^t \times Y \rightarrow C$. We denote the principal’s and the agent’s infinite horizon strategies by the sequences of strategy functions $\mathbf{c} = (c_1, c_2, \dots)$ and $\mathbf{a} = (a_1, a_2, \dots)$, respectively.

For each period t and history $h^t \in H^t$, the strategy function a_t of the agent leads to (marginal) probability distributions on the set H^t of histories,

$$\mu((h^t, y) | h^t, \mathbf{a}) = f(y | a_t(h^t)),$$

and by recursion for $\tau \geq 1$,

$$\mu((h^{t+\tau}, y) | h^t, \mathbf{a}) = f(y | a_{t+\tau}(h^{t+\tau})) \mu(h^{t+\tau-1} | h^t, \mathbf{a}).$$

Prior to period 1, the agent has the reservation utility \tilde{w} . We assume that both the agent and the principal have to commit to a strategy before period 1. This assumption implies that only the overall discounted expected utility for the agent has to be at least \tilde{w} . Otherwise we would have to add a constraint at each time t requiring that the agent has to receive utility of at least \tilde{w} . We describe and solve this extension of the standard problem in Section 2.4 below.

We are now in the position to define the *dynamic principal–agent problem*. The principal faces the following infinite horizon utility maximization problem:³

$$\max_{\mathbf{c}, \mathbf{a}} \mathbb{E} \left[\sum_{t=1}^{\infty} \beta^{t-1} v(Y_t - c_t(h^t, Y_t)) \mid h^1, \mathbf{a} \right] \quad (1)$$

²This last assumption, while standard, is not necessary for the theoretical framework.

³In line with the notation for the measure μ , we place the argument \mathbf{a} in the expectation next to the history of output h^t . A different way of expressing the dependence of the expectation on the strategy \mathbf{a} would be to write $\mathbb{E}^{\mathbf{a}}$. Here we follow the notation in Spear and Srivastava (1987).

subject to the constraints for any $\bar{t} \in \mathbb{N}$ and any history $h^{\bar{t}} \in H^{\bar{t}}$ that

$$(a_t)_{t \geq \bar{t}} \in \arg \max_{(\tilde{a}_t)_{t \geq \bar{t}}} \mathbb{E} \left[\sum_{t=\bar{t}}^{\infty} \beta^{t-1} u(\tilde{a}_t(h^t), c_t(h^t, Y_t)) \middle| h^{\bar{t}}, (\tilde{a}_t)_{t \geq \bar{t}} \right], \quad (2)$$

and

$$\tilde{w} \leq \mathbb{E} \left[\sum_{t=1}^{\infty} \beta^{t-1} u(a_t(h^t), c_t(h^t, Y_t)) \middle| h^1, \mathbf{a} \right], \quad (3)$$

where Y_t is a random variable with values in Y and $\mathbb{P}(Y_1 = y_1, \dots, Y_t = y_t) = \mu((y_1, \dots, y_t) | h^1, \mathbf{a})$.

An optimal compensation contract for the dynamic principal–agent problem is a sequence of compensation functions \mathbf{c}^* together with a sequence of action functions \mathbf{a}^* such that $(\mathbf{c}^*, \mathbf{a}^*)$ is an optimal solution to the optimization problem (1,2,3). At the optimal solution $(\mathbf{c}^*, \mathbf{a}^*)$, the principal maximizes her lifetime utility in period 1 under the condition that the agent’s incentive compatibility constraint holds for all possible histories $h^{\bar{t}} \in H^{\bar{t}}$ in each period $\bar{t} = 1, 2, \dots$ and the agent’s expected utility at the beginning of period 1 is at least \tilde{w} .

Observe that, in the language of game theory, solutions to the dynamic principal–agent problem are sequential equilibria for the supergame between the principal and the agent. This observation allows us to use the approach of Abreu, Pearce, and Stacchetti (1990) to derive a recursive formulation of the dynamic principal–agent problem.

2.2 A Recursive Formulation

The fundamental insight of Spear and Srivastava (1987, Section III) is that the infinite-horizon problem (1,2,3) is equivalent to a static variational problem. For completeness, we re-prove this result using the modern techniques of Abreu, Pearce, and Stacchetti (1990), which allow for the reduction of a repeated game to a static framework.

We can easily establish a lower and an upper bound on the expected lifetime utility of the agent. Observe that the principal can offer the agent a deterministic contract. This observation leads to the lower and upper bound

$$\underline{w} = \frac{1}{1 - \beta} \min_{c \in C} \max_{a \in A} u(a, c) \quad \text{and} \quad \bar{w} = \frac{1}{1 - \beta} \max_{a \in A, c \in C} u(a, c),$$

respectively, on the agent’s expected lifetime utility. Note that both extrema exist since, by Assumption 1, the action set A and the compensation set C are compact and the utility function u is continuous. Define the set of possible utility values for the agent by $W = [\underline{w}, \bar{w}]$.

For the presentation of the recursive formulation it is helpful to explicitly define the random variable of possible outcomes in any given period. We denote this random variable with values in Y and probability distribution $f(\cdot | a)$ by \bar{Y} .

Theorem 1. *Let a dynamic principal–agent problem with action set A , output set Y , compensation set C , and utility set W be given. Also let the reservation utility be $\tilde{w} \in W$. Suppose the functions*

$$V : W \rightarrow \mathbb{R}, \quad c : W \times Y \rightarrow C, \quad U : W \times Y \rightarrow W, \quad a : W \rightarrow A$$

solve the following dynamic programming problem:

$$\begin{aligned} V(w) &= \max_{\hat{c}, \hat{U}, \hat{a}} \mathbb{E} \left[v(\bar{Y} - \hat{c}(\bar{Y})) + \beta V(\hat{U}(\bar{Y})) \mid \hat{a} \right], \\ w &= \mathbb{E} \left[u(\hat{a}, \hat{c}(\bar{Y})) + \beta \hat{U}(\bar{Y}) \mid \hat{a} \right], \\ \hat{a} &\in \arg \max_{a \in A} \mathbb{E} \left[u(a, \hat{c}(\bar{Y})) + \beta \hat{U}(\bar{Y}) \mid a \right] \\ \hat{c}(\bar{Y}) &\in C, \quad \hat{U}(\bar{Y}) \in W. \end{aligned} \tag{4}$$

Also assume that the function V is upper semi-continuous. Then $\max_{w \in W, w \geq \tilde{w}} V(w)$ is the optimal value of problem (1,2,3).⁴

The statement of the theorem provides essentially the same solution as Spear and Srivastava (1987, Section III). However, we omit the redundant constraints that appear in their description, which requires us to know the policy function $U(w, y)$. We demonstrate below that this simplification allows us to reduce the problem to a DP that can be solved using standard value function iteration. The overall objective now is to prove that the recursive formulation (4) is equivalent to the sequential problem (1,2,3). In a first step, we show that any solution to the two constraints in the proposed recursive formulation describes a solution to the agent’s sequential problem (2) in the original problem. The proofs follow the same pattern as those in Abreu, Pearce, and Stacchetti (1990).

Consider any value $w \in W$. First, we want to prove the existence of well-defined functions $c : W \times Y \rightarrow C$ and $U : W \times Y \rightarrow W$ satisfying the following recursive relationship:

$$w = \max_{a_1(w) \in A} \mathbb{E} \left[u(a_1(w), c(w, \bar{Y})) + \beta U(w, \bar{Y}) \mid a_1(w) \right],$$

where the expectation is taken with respect to the probability distribution given by $f(\cdot \mid a_1(w))$. The role of the function U is to summarize all future (discounted) utilities to the agent.⁵ The function on the right-hand side is continuous in the agent’s decision variable since the probability $f(y \mid a)$ is continuous in a for all $y \in Y$. And so, the maximum exists and is attained.

⁴Since we chose Y to be discrete, we do not need any further restrictions on the functions c or U . For an uncountably infinite state space however we would need to assume at least integrability.

⁵If we assumed Y to be uncountably infinite, then we would have to show that the expectation exists, i.e., that U is indeed integrable over Y . However, for a finite sum such a proof is not necessary.

Lemma 1. Let $c : W \times Y \rightarrow C$, $U : W \times Y \rightarrow W$, and $a : W \rightarrow A$ be functions that solve the following problem, $\forall w \in W$:

$$\begin{aligned} w &= \mathbb{E} \left[u(a(w), c(w, \bar{Y})) + \beta U(w, \bar{Y}) \middle| a(w) \right], \\ a(w) &\in \arg \max_{\hat{a} \in A} \mathbb{E} \left[u(\hat{a}, c(w, \bar{Y})) + \beta U(w, \bar{Y}) \middle| \hat{a} \right]. \end{aligned}$$

Then for all $w \in W$, $a(w)$ and $U(w, y)$ determine an optimal solution for (2) with expected payoff w .

Proof. First we show that for any $w \in W$ the functions c , a , and U exist. Note that since Y is finite, U and c can be identified, for fixed w , as vectors in \mathbb{R}^N . Furthermore, the Bernoulli utility u is continuous in a and c , and the probability $f(y|a)$ is continuous in a ; also, the sets A , C^N , and W^N are compact. So by Berge's theorem (Berge 1963), and the mean value theorem, for any $w \in W$, there are $c(w, y) \in C^N$, $U(w, y) \in W^N$ such that

$$w = \max_{\hat{a} \in A} \mathbb{E} [u(\hat{a}, c(w, \bar{Y})) + \beta U(w, \bar{Y}) \mid \hat{a}].$$

Choose $a(w)$ as one of the solutions to this problem.

For a fixed w we can now define the optimal strategy for problem (2) recursively.

$$\begin{aligned} U_1(h^1, y) &= U(w, y), \\ U_t(h^t, y) &= U(U_{t-1}(h^{t-1}, y_{t-1}), y), \\ c_1(h^1, y) &= c(w, y), \\ c_t(h^t, y) &= c(U_{t-1}(h^{t-1}, y_{t-1}), y), \\ a_1(h^1) &= a(w), \\ a_t(h^t) &= a(U_{t-1}(h^{t-1}, y_{t-1})). \end{aligned}$$

By construction, for the chosen w , the expected value of the strategy $(a_t(h^t))_{t \geq 1}$ (at the beginning of period 1) is w . Also for $h^t = (h^{t-1}, y_{t-1})$ the following holds:

$$\begin{aligned} U_{t-1}(h^{t-1}, y_{t-1}) &= \max_{\hat{a} \in A} \mathbb{E} [u(\hat{a}, c_t(h^t, \bar{Y})) + \beta U_t(h^t, \bar{Y}) \mid h^t, \hat{a}] \\ &= \mathbb{E} [u(a_t(h^t), c_t(h^t, \bar{Y})) + \beta U_t(h^t, \bar{Y}) \mid h^t, a_t(h^t)]. \end{aligned}$$

Thus this strategy is unimprovable and therefore optimal. Iterating over t yields an optimal solution for (2) with expected payoff w . \square

Next we show that the agent's optimal sequential response to a sequential consumption function $(c_t(h^t, y_t))_{t \geq 1}$ with payoff w for problem (2) also solves the agent's recursive problem.

Lemma 2. Let $\mathbf{a} = (a_t)_{t \geq 1}$ be an optimal strategy for the agent's problem given function $\mathbf{c} = (c_t)_{t \geq 1}$ with payoff w . Then there is a function $U(w, y) : Y \rightarrow W$ and a value $a(w) \in A$ such that

$$\begin{aligned} w &= \mathbb{E} [u(a(w), c_1(h^1, \bar{Y})) + \beta U(w, \bar{Y}) \mid a(w)], \\ a(w) &\in \arg \max_{\hat{a} \in A} \mathbb{E} [u(\hat{a}, c_1(h^1, \bar{Y})) + \beta U(w, \bar{Y}) \mid \hat{a}]. \end{aligned}$$

Proof. Define the following functions:

$$\begin{aligned} a(w) &= a_1(h^1), \\ U(w, y) &= \mathbb{E} \left[\sum_{t=2}^{\infty} \beta^{t-2} u(a_t(h^t), c_t(h^t, Y_t)) \mid (h^1, y), \mathbf{a} \right]. \end{aligned}$$

Due to the definition of W , clearly $U(w, y) \in W$ for all y . By first using the tower property for conditional expectations and then the optimality of a we get the following:

$$\begin{aligned} w &= \mathbb{E} \left[\sum_{t=1}^{\infty} \beta^{t-1} u(a_t(h^t), c_t(h^t, Y_t)) \mid h^1, a \right] \\ &= \mathbb{E} [u(a(w), c_1(h^1, \bar{Y})) + \beta U(w, \bar{Y}) \mid h^1, a(w)] \\ &= \max_{\hat{a} \in A} \mathbb{E} [u(\hat{a}, c_1(h^1, \bar{Y})) + \beta U(w, \bar{Y}) \mid h^1, \hat{a}]. \end{aligned}$$

□

Now we can show how to find a solution for the principal's optimization problem.

Proof of Theorem 1. Lemmas 1 and 2 show that the agent's problem is equivalent, in the sense of payoff equivalence, to the recursive agent's problem. By the same argument, we can construct the principal's policy as in Lemma 1.⁶ In particular $\{V(w) \mid w \geq \tilde{w}, w \in W\}$ is the set of achievable payoffs to the principal. The principal now can pick how high the initial payoff w to the agent is as long as it is greater or equal to \tilde{w} and also implementable—that is, $w \in W$. The maximal possible payoff is then $\max_{w \in W, w \geq \tilde{w}} V(w)$. This maximum is well defined since V is assumed to be upper semi-continuous and the set W is compact. □

The following corollary now gives us a way of obtaining the value function for the original problem from the recursive one. The construction is an immediate consequence of the proof of Lemma 1.

⁶The converse is not necessarily true, i.e., an optimal strategy might not correspond to a recursive one. However there will always be one that gives the same payoff. This can happen if the optimal solution is not unique. Then, alternating between the optimal solutions maintains the same payoff but this optimal solution cannot be replicated as a recursive solution.

Corollary 1. *Suppose the assumptions of Theorem 1 hold. And let V, a, c, U be functions solving (4). Denote by w_0 one of the solutions of $\max_{w \in W, w \geq \bar{w}} V(w)$. Let $h^t \in H^t$ be the history at time t . Define*

$$\begin{aligned}
U_1(h^1, y_1) &= U(w_0, y_1), \\
U_t(h^t, y_t) &= U(U_{t-1}(h^{t-1}, y_{t-1}), y_t), \\
c_1(h^1, y_1) &= c(w_0, y_1), \\
c_t(h^t, y_t) &= c(U_{t-1}(h^{t-1}, y_{t-1}), y_t), \\
a_1(h^1) &= a(w_0), \\
a_t(h^t) &= a(U_{t-1}(h^{t-1}, y_{t-1})), \\
V_1(h^1) &= V(w_0), \\
V_t(h^t) &= V(U_{t-1}(h^{t-1}, y_{t-1})).
\end{aligned}$$

Then V_t is the value function and (c_t, a_t) is the optimal policy to (1,2,3) at time t .

2.3 Existence

We next examine the dynamic programming problem (4). For this purpose we first state two helpful mathematical results on upper semi-continuous bounded (real-valued) functions. Appendix A provides the proofs.

Lemma 3. *Let X be a compact metric space and $\mathcal{F}(X)$ the set of upper semi-continuous, bounded (real-valued) functions on X . Then $\mathcal{F}(X)$ equipped with the supremum norm is a complete metric space.*

Theorem 2. *Let X, Y be compact metric spaces and $\Gamma(x) : X \rightarrow 2^Y$ a closed, set-valued mapping. Further let $\mathcal{F}(X)$ be the set of upper semi-continuous bounded functions equipped with the supremum norm. Also let $F : X \times Y \rightarrow \mathbb{R}$ be an upper semi-continuous function. Define the operator T on the domain $\mathcal{F}(X)$ by*

$$(Tf)(x) = \sup_{y \in \Gamma(x)} \{F(x, y) + \beta f(y)\}.$$

Then $T : \mathcal{F}(X) \rightarrow \mathcal{F}(X)$ and this mapping has a unique fixed point.

We want to apply this theorem to the recursive formulation (4) of the dynamic principal-agent problem. For this purpose we need the following lemma:

Lemma 4. *The set-valued mapping $\Gamma : W \rightarrow 2^{A \times C \times W}$ defined by*

$$\Gamma(w) = \left\{ \begin{array}{l} (a, c(y), U(y)) \in \mathbb{R} \times \mathbb{R}^N \times \mathbb{R}^N | \\ w = \max_{\tilde{a} \in A} \mathbb{E} [u(\tilde{a}, c(\bar{Y})) + \beta U(\bar{Y}) | \tilde{a}], \\ w = \mathbb{E} [u(a, c(\bar{Y})) + \beta U(\bar{Y}) | a], \\ a \in A, c(y) \in W, U(y) \in C, y \in Y \end{array} \right\}$$

is closed and uniformly compact.

Proof. Note that the sum of upper semi-continuous functions is again upper semi-continuous. Since A is a fixed compact set, it follows, from Hogan (1973, Theorem 7), that all functions involved in the constraints are continuous. Thus by Hogan (1973, Theorem 10) the set-valued mapping is closed. Since its values are all contained in the compact set $A \times C \times W$, it is trivially uniformly compact. \square

Finally we can state and prove an existence theorem for the dynamic principal-agent model.

Theorem 3. *The unique value function $V : W \rightarrow \mathbb{R}$ in the recursive formulation (4) is an element of the complete metric space $\mathcal{F}(W)$ of upper semi-continuous bounded functions on the compact domain W .*

Proof. The sets W and $A \times C \times W$ are compact metric spaces. Lemma 4 states that the feasible region for the dynamic optimization problem is a closed, uniformly compact, set-valued mapping. Theorem 2 now implies that the dynamic programming problem (4) has a unique fixed point V ; the value function is upper semi-continuous and bounded on its compact domain W . \square

The interested reader may be wondering why we cannot simply apply Berge's Theorem of the Maximum; see Berge (1963). The reason why we need the more involved Theorem 2 is that the constraint mapping in the recursive formulation (4) is only upper semi-continuous and may not be continuous; in general, the constraint mapping will not be continuous because the openness condition fails. Unfortunately, sufficient conditions for openness require many more assumptions; see Hogan (1973, Theorems 12, 13).

Corollary 2. *Let V^* be the optimal value of the non-recursive problem (1,2,3), then there is a value function $V : W \rightarrow \mathbb{R}$ as in Theorem 1 solving the recursive problem with $V^* = \max_{w \geq \bar{w}} V(w)$.*

Proof. This is an immediate consequence of Theorems 1 and 3. \square

Corollary 2 not only characterizes the optimal value V^* of the original *dynamic principal-agent problem*, it also provides us with a theoretical foundation for a numerical solution approach. Theorem 1 states that if an upper semi-continuous value function V for the DP problem (4) exists, then $V^* = \max_{w \geq \bar{w}} V(w)$ is the optimal value for the original sequential dynamic principal-agent problem (1,2,3). Theorem 3 shows that the DP problem (4) possesses indeed a unique upper semi-continuous bounded value function. And not only that, the proof of this theorem via Theorem 2 implies that this unique value function is the fixed point of a contraction mapping. This contraction property of the Bellman operator suggests that we can apply value function iteration to numerically approximate the value function of the DP problem (4). We describe our numerical procedure in Section 2.6 below.

2.4 Lack of Commitment and Renegotiation

The description and solution of the dynamic principal-agent problem assume that neither the principal nor the agent can walk away from the contract once they have agreed upon it in the first period. We now consider two modifications of the original model. First, we weaken the assumption on the agent’s commitment to a contract. Second, we consider a model in which the principal is permitted to renegotiate the contract.

2.4.1 No Contractual Commitment

Constraint (3) in the dynamic principal-agent problem imposes the condition that the contract must provide the agent with a particular reservation utility \tilde{w} at time $t = 1$. There is no similar requirement in later time periods, and so the agent’s continuation utility w —see the dynamic programming problem (4)—may fall below this initial reservation utility \tilde{w} . Clearly, this model feature does not capture situations in which the agent continues to have an outside option, providing him with utility \tilde{w} , after the first period.

We now present a modification of the standard model and permit the agent to walk away at any time during the contract. Therefore, the principal now must provide a reservation utility at each time t . Thus, the state space now has the reservation utility \tilde{w} as its lower bound—that is, the state space becomes $\tilde{W} = [\tilde{w}, \bar{w}]$. We say that the agent is not committed to the contract.⁷ The new assumption results in the following variational problem, where

⁷Seen as a game, both players are still committed to the strategies chosen at time 1. In that sense, the lower bound in every period is analogous to the reservation utility. The bound simply makes all those low-utility cases infeasible, in which the agent would walk away.

$w \in \tilde{W}$,

$$\begin{aligned}
V(w) &= \max_{\hat{c}, \hat{U}, \hat{a}} \mathbb{E} \left[v(\bar{Y} - \hat{c}(\bar{Y})) + \beta V(\hat{U}(\bar{Y})) \mid \hat{a} \right], \\
w &= \mathbb{E} \left[u(\hat{a}, \hat{c}(\bar{Y})) + \beta \hat{U}(\bar{Y}) \mid \hat{a} \right], \\
\hat{a} &\in \arg \max_{a \in A} \mathbb{E} \left[u(a, \hat{c}(\bar{Y})) + \beta \hat{U}(\bar{Y}) \mid a \right] \\
\hat{c}(\bar{Y}) &\in C, \quad \hat{U}(\bar{Y}) \in \tilde{W}.
\end{aligned} \tag{5}$$

Observe that the only, but clearly important, difference between the modification (5) and the original DP problem (4) is the constraint $\hat{U}(\bar{Y}) \in \tilde{W}$. As we depict in the analysis of the unemployment insurance problem of Hopenhayn and Nicolini (1997) in Section 4, the modified formulation removes some outcomes we may regard as unreasonable. (In addition, if we want to allow for an outside option for the principal, then we can check the principal's commitment to the contract, just as in the static principal-agent model, after we have solved the problem by verifying that $V_t(h^t)$, as defined in Corollary 1, exceeds some exogenous lower bound for all possible histories.)

2.4.2 Renegotiation

We consider a second simple yet important modification of the original model. Suppose in some period t the principal must deliver a utility promise w , but, in fact, there is a $w_0 \in W$ with $w_0 > w$ and $V(w_0) > V(w)$. If the principal stays true to her past promise w , then she cannot just give the agent more utility w_0 . In essence, this condition means that if she previously threatened the agent with a lower utility then she has to carry out that threat now. We now weaken this assumption to obtain a different recursive formulation. Suppose that the agent acts as if the threats of the principal always remain credible, but that the principal now has the option to renegotiate the contract at time point t . So, in this situation, she offers the agent $w_0 > w$, and the agent accepts the new offer since it improves upon the original utility promise w . This modification of the model results in the following recursive problem,

$$\begin{aligned}
V(w) &= \max_{\hat{c}, \hat{U}, \hat{a}} \mathbb{E} \left[v(\bar{Y} - \hat{c}(\bar{Y})) + \beta V(\hat{U}(\bar{Y})) \mid \hat{a} \right], \\
w &\leq \mathbb{E} \left[u(\hat{a}, \hat{c}(\bar{Y})) + \beta \hat{U}(\bar{Y}) \mid \hat{a} \right], \\
\hat{a}(w) &\in \arg \max_{a \in A} \mathbb{E} \left[u(a, \hat{c}(\bar{Y})) + \beta \hat{U}(\bar{Y}) \mid a \right] \\
\hat{c}(\bar{Y}) &\in C, \quad \hat{U}(\bar{Y}) \in W.
\end{aligned} \tag{6}$$

Notice that the first constraint of the original DP problem (4) has been relaxed to an inequality; see the first constraint of the modified problem (6). Unlike in the standard model, the principal will offer the agent a higher utility if this increases her own payoff $V(w)$.

This change simplifies the problem immensely and results in the following theorem.

Theorem 4. *The solution $V : W \rightarrow \mathbb{R}$ of (6) is decreasing and therefore quasi-concave.*

Proof. Let \mathcal{F}_w denote the feasible set of the optimization problem (6) in state w . For any two $w_1, w_2 \in W$ with $w_1 > w_2$, the feasible regions satisfy $\mathcal{F}_{w_1} \subseteq \mathcal{F}_{w_2}$. Therefore, $V(w_1) \leq V(w_2)$. \square

2.5 Additional Properties of the Value Function?

Spear and Srivastava (1987) state that the value function V is both decreasing and concave. Neither of these claims is correct. In addition, their paper is rather imprecise in its definitions of the feasible sets. In the following we explain the underlying technical issues in more detail.


Spear and Srivastava (1987, p. 604) provide the following argument that the value function is decreasing: “We note that $V(w)$ is the Pareto frontier for this problem, and must be non-increasing. This follows because otherwise, the principal could simply offer the agent more and be better off.” However, their mathematical formulation of the recursive dynamic problem requires that “[...] w is the expected utility of the agent”; see condition (i) on page 603 in Spear and Srivastava (1987). But their Pareto argument requires that w is only a lower bound on the agent’s expected utility. Therefore, the Pareto argument only applies to the model with renegotiation but is incorrect for their model. We observe below that the value function in the social planning problem in Section 3 is not decreasing on the entire domain. Intuitively, it is costly to the principal to enforce a low continuation volatility on the agent. When she does just that, the agent’s effort is so low that the principal is worse off compared to a slightly higher utility promise.

Spear and Srivastava (1987) claim and prove that the function V is concave in the state variable w . But the concavity proof in Spear and Srivastava (1987, pp. 604–5) as stated in the paper is insufficient. The authors attempt to prove concavity by deriving a contradiction. They assume that there is an interval such that the function V is convex over that interval. From this assumption they derive a contradiction and thus conclude that no such interval can exist. But that argument does not prove the concavity of the value function. The continuous, decreasing, and piecewise strictly concave function f defined as

$$f(x) = \begin{cases} -e^x & x \leq 0 \\ -e^{x-3} - 1 + e^{-3} & x > 0 \end{cases}$$

is a counterexample to this argument. Figure 2.5 shows the graph of f on the interval $[-2, 2]$. Surprisingly, this counterexample appears to fall exactly into the appropriate class of functions for the value function V in the model extension with renegotiation; see Theorem 4 above.

Figure 1: Graph of the function f on the interval $[-2, 2]$.



The function f is decreasing and quasi-concave; it is strictly concave for $x \leq 0$ and for $x \geq 0$. But the function f is not concave.

We conjecture that the concavity statement of Spear and Srivastava (1987) is incorrect in general, even though we could not find an actual value function as a simple counterexample in the context of a dynamic principal–agent model satisfying the assumptions in the paper. In fact, we observed that during the value function iterations some iterates themselves were non-concave, even when starting from a concave guess. Of course, this observation does not conclusively disprove their claim but in combination with the false proof this example casts considerable doubt on the concavity result in Spear and Srivastava (1987).

Spear and Srivastava (1987, p. 602) state that a model in which the agent could quit would require an additional lower-bound constraint. Similarly, they argue that a model in which the principal could terminate the contract at any time, would also require an additional constraint. They then conclude, “Since these additional constraints complicate the analysis tremendously, we focus on the case with commitment.” Contrary to their conclusions, we have shown in Section 2.4 that both model features can be easily added to the model. In fact, Theorem 4 shows that the model with renegotiation allows the derivation of additional theoretical properties of the value function; and our computational analysis of the unemployment insurance problem of Hopenhayn and Nicolini (1997) in Section 4 shows that relaxing the commitment of the agent may greatly simplify the analysis because that assumption may considerably reduce the feasible region of the resulting nonlinear optimization problem.

The confusion in Spear and Srivastava (1987) regarding the difficulty of some natural modifications of their model may possibly stem from their rather imprecise definitions of some sets that are important for a precise problem formulation. While Spear and Srivastava (1987, p. 604) provides a general description of how to obtain a lower and an upper bound on the value set W , the paper does not explicitly account for the set W in the general problem formulation on pp. 603–4 and in the reformulation using the first-order approach on p. 606. Assumption (A.3) in Spear and Srivastava (1987, Section IV) imposes a zero lower bound on the agent’s compensation. This assumption implies that the set W has a closed lower bound. Moreover, while Spear and Srivastava (1987, p. 602) assume that the action set A is compact, the reformulation does not allow for boundary solutions and in fact assumes an interior solution. However, as we demonstrate in our analysis of the unemployment insurance problem of Hopenhayn and Nicolini (1997) in Section 4, the optimal effort at the lowest element of W may actually be a boundary element of A . As a result, the first-order conditions in Spear and Srivastava (1987, p. 606) cannot hold on the entire domain.

2.6 Computation

The dynamic programming problem (4) may appear rather daunting at first. However, we can numerically solve the problem by combining the standard dynamic programming technique of value function iteration—see, for example, Judd (1998, Chapter 12)—with the polynomial solution method of Renner and Schmedders (2015) for static principal–agent models.

We use a continuous approximation scheme for the value function $V : W \rightarrow \mathbb{R}$, even though Theorem 3 only states that V is semi-continuous. In our economic applications, the value function always turned out to be continuous. And, if a value function were truly discontinuous, then we would likely detect that property in the numerical approach via sizable numerical errors. Obviously, at that point we would need a different approximation scheme.

Let $\{w_1, \dots, w_M\}$ be a set of interpolation nodes, for example Chebyshev zeros, in the state space W . Suppose after iteration $k = 0, 1, 2, \dots$, function values $\{v_1^k, \dots, v_M^k\}$ for the approximate value function at the M interpolation nodes are given. Then let $V^k(w)$ denote the unique polynomial in the variable w of degree $M - 1$ interpolating the M points $\{(w_1, v_1^k), \dots, (w_M, v_M^k)\}$.

The contraction property in the Bellman equation in (4) readily suggests an application of value function iteration; to generate a sequence of approximations V^k for the (true) value function V , we simply apply the iteration $V^{k+1} = TV^k$, $k = 0, 1, 2, \dots$. Thus, in each value function iteration we have to solve the following M problems, given the current approximation

$V^k(w)$:

$$\begin{aligned}
v_i^{k+1} &= \max_{\hat{c}(y), \hat{U}(y), \hat{a}} \mathbb{E} \left[v(\bar{Y} - \hat{c}(\bar{Y})) + \beta V^k(\hat{U}(\bar{Y})) \mid \hat{a} \right], \\
w_i &= \mathbb{E} \left[u(\hat{a}, \hat{c}(\bar{Y})) + \beta \hat{U}(\bar{Y}) \mid \hat{a} \right], \\
\hat{a} &\in \arg \max_{a \in A} \mathbb{E} \left[u(a, \hat{c}(\bar{Y})) + \beta \hat{U}(\bar{Y}) \mid a \right], \\
\hat{U}(y) &\in W \quad \forall y \in Y, \\
\hat{c}(y) &\in C \quad \forall y \in Y.
\end{aligned}$$

Note that, as always in standard DP value function iteration, we do not need to approximate any function other than the value function. So, after we solve the M problems above, we only have to determine the next approximation $V^{k+1}(w)$ for the value function using the just computed values $\{v_1^{k+1}, \dots, v_M^{k+1}\}$.

For a fixed interpolation node w_i , the problem is rather similar to a static principal–agent problem with the caveat that the objective function is no longer linear, as is often assumed for static principal–agent problems. Instead, in iteration $k + 1$, we now face the nonlinear objective function $\mathbb{E} \left[v(\bar{Y} - \hat{c}(\bar{Y})) + \beta V^k(\hat{U}(\bar{Y})) \mid \hat{a} \right]$. The major complication of this optimization problem at each node w_i is, however, the incentive compatibility constraint, which, just as in static principal–agent models, makes the problem into a bi-level optimization problem. To solve this problem, we now apply the polynomial solution method of Renner and Schmedders (2015) for static principal–agent models. Renner and Schmedders (2015) apply modern methods from algebraic geometry—see Laurent (2009) and Lasserre (2010)—and the global optimization approach for rational functions of Jibeteau and De Klerk (2006) to solve static principal–agent models.

If the agent’s objective function, $\mathbb{E} \left[u(a, \hat{c}(\bar{Y})) + \beta \hat{U}(\bar{Y}) \mid a \right]$ is a polynomial or rational function in a then we can replace the agent’s problem by a set of inequalities and equations; see Renner and Schmedders (2015, Section 4.1). Otherwise we approximate the function by a polynomial or rational function and then perform the replacement; see the solution of the famous Mirrlees (1999) example in Renner and Schmedders (2015, Section 5.1). As a result, we obtain a standard optimization problem with the same objective function, $\mathbb{E} \left[v(\bar{Y} - \hat{c}(\bar{Y})) + \beta V^k(\hat{U}(\bar{Y})) \mid \hat{a} \right]$, as before. The final nonlinear optimization problem we solve with the large-scale NLP solver IPOPT; see Wächter and Biegler (2006).

3 A Social Planning Problem

We consider a continuous action version of the social planning problem from Phelan and Townsend (1991). There are two states, $Y = \{1, 2\}$. The agent’s action set is $A = [0, 0.6]$. We

approximate the probability function f , which in Phelan and Townsend (1991, Section III) is simply given by a table for their four possible action values, by the following pair of functions:

$$\left\{ \frac{0.9}{1+3a}, \frac{0.1+3a}{1+3a} \right\}.$$

The compensation set is $C = [0, 2.25]$. In the language of Phelan and Townsend (1991), the social contract provides the agent with some consumption amount $c \in C$. The principal's and the agent's utility functions v and u , respectively, are defined as follows:

$$\begin{aligned} v(x) &= x, \\ u(a, c) &= 2\sqrt{c} + 2\sqrt{1-a}. \end{aligned}$$

Phelan and Townsend (1991) define a contract design problem of finding a surplus-maximizing contract for an agent who prefers lower action and higher consumption. Surplus maximization leads to a risk-neutral principal (the social planner) in our description of the model. (In addition to the model with moral hazard, Phelan and Townsend (1991) also examine the full information version of their model. For completeness, Appendix B reports results for our version of the full information model as well.)

The agent's utility function and the consumption bounds lead to the set of utilities:

$$W = \left[\frac{u(0, 0)}{1-\beta}, \frac{u(0, 2.25)}{1-\beta} \right] = \left[\frac{2}{1-\beta}, \frac{5}{1-\beta} \right].$$

Clearly the agent's utility function is not polynomial in his action a . Therefore, we approximate this function with a degree 2 Chebyshev polynomial,

$$2\sqrt{c} + 1.99827 - 0.446702a^2 - 0.94914a.$$


We approximate the agent's value function U by a degree-39 polynomial and so use 40 Chebyshev nodes. For $\beta = 0.8$, the endogenous state space is $W = [10, 25]$, while for $\beta = 0.95$ it is $W = [40, 100]$. Due to the approximation of the agent's utility function we choose the Chebyshev nodes for $\beta = 0.8$ in the interval $[9.991, 24.991]$; and for $\beta = 0.95$ in the interval $[39.965, 99.965]$. Figure 2 displays the graphs of the value function for the two discount factors. Table 1 provides detailed results for $\beta = 0.8$. The value function for our modification shows the same properties as the surplus function reported in Phelan and Townsend (1991, Section IV). It is increasing for small values of w , decreasing after a global maximum, and concave. Also, the behavior of the optimal action is similar. It rises for small values of w and after a global maximum it decreases monotonically.

Table 1: Results for $\beta = 0.8$ with absolute L^1 error of 10^{-4} .

w	V	a	c_1	c_2	U_1	U_2
9.991	5.500	0.000	0.000	0.000	9.991	9.991
10.088	6.072	0.070	0.000	0.000	10.023	10.707
10.377	6.780	0.141	0.000	0.008	10.281	11.074
10.850	6.969	0.244	0.000	0.060	10.608	11.621
11.496	6.895	0.301	0.004	0.127	11.037	12.354
12.296	6.663	0.299	0.058	0.199	11.607	13.123
13.231	6.284	0.282	0.141	0.286	12.508	14.007
14.276	5.750	0.260	0.246	0.394	13.599	14.999
15.405	5.059	0.235	0.372	0.526	14.804	16.079
16.587	4.211	0.212	0.525	0.685	16.068	17.218
17.793	3.217	0.189	0.704	0.871	17.348	18.386
18.992	2.099	0.169	0.908	1.080	18.609	19.551
20.151	0.893	0.152	1.131	1.308	19.819	20.682
21.241	-0.354	0.136	1.364	1.546	20.950	21.747
22.235	-1.587	0.123	1.598	1.787	21.974	22.713
23.105	-2.745	0.111	1.826	2.044	22.863	23.531
23.830	-3.775	0.094	2.064	2.250	23.581	24.207
24.391	-4.645	0.059	2.250	2.250	24.161	24.806
24.773	-5.328	0.000	2.250	2.250	24.719	24.719
24.967	-5.703	0.000	2.250	2.250	24.961	24.961

The table provides detailed results for the social planning problem. The first column lists every second of the 40 Chebyshev nodes in the endogenous state space $[9.991, 24.991]$. The second column provides the values $V(w)$ of the social planner's value function at the grid points. The third column reports the agent's optimal action a , while the next two columns report the consumption values c_1 and c_2 of the social contract at the grid points. Finally, the last two columns present the agent's utility values U_1 and U_2 .

Figure 2: Plot of the value function V for $\beta = 0.8$ (left) and $\beta = 0.95$ (right).



The figure displays the degree-39 polynomial approximations of the social planner's value function V on the endogenous state space. The left graph shows the approximation of the value function for $\beta = 0.8$ on the domain $[9.991, 24.991]$. The right graph shows the value function for $\beta = 0.95$ on the domain $[39.965, 99.965]$.

4 Optimal Unemployment Insurance

In this section we examine the design problem of an unemployment insurance system presented in Hopenhayn and Nicolini (1997).⁸ The textbook of Ljungqvist and Sargent (2000, pp. 427–434) describes a slightly modified version of this problem and provides some helpful economic interpretation. We briefly introduce the recursive formulation of the economic model and refer to Hopenhayn and Nicolini (1997) for the initial sequential formulation and other details.

4.1 Recursive Formulation

At time zero, a risk-neutral principal offers a contract to a risk-averse agent who is unemployed. The contract offers the agent an expected discounted utility, w . For the agent to obtain w , the contract specifies current consumption c and current (job search) effort a as well as promised continuation values w^e and w^u for the agent in the subsequent period depending on whether he is employed or unemployed, respectively. The four characteristics c , a , w^e , and w^u of the contract must satisfy the following condition:

$$w = u(c) - a + \beta(p(a)w^e + (1 - p(a))w^u),$$

where u denotes the agent's Bernoulli utility of consumption, β his discount factor, and p the probability of the agent finding a job as a function of his search effort, a . The effort level a

⁸The notation in this section is self-contained. We largely adopt the notation of Hopenhayn and Nicolini (1997).

specified in the contract must be incentive compatible, so

$$a \in \arg \max_{\tilde{a}} u(c) - \tilde{a} + \beta (p(\tilde{a})w^e + (1 - p(\tilde{a}))w^u).$$

The conditions implicitly assume that the agent has no other source of income and cannot borrow or save. Put differently, the principal directly controls the agent's consumption c . Hopenhayn and Nicolini (1997) impose the following technical assumptions.

Assumptions [Hopenhayn and Nicolini (1997)].

1. *The discount factor satisfies $\beta < 1$.*
2. *The set $A \subset \mathbb{R}$ of effort levels is a closed interval and $0 \in A$.*
3. *The consumption set is $C = \mathbb{R}_+$.*
4. *The Bernoulli utility function $u : C \rightarrow \mathbb{R}$ is increasing and concave.*
5. *The probability function $p : A \rightarrow [0, 1]$ is increasing, strictly concave, twice differentiable, and satisfies an Inada condition.*

Once the agent finds a job, he receives a permanent and constant wage b . The principal has the power to tax an employed agent and so the agent's permanent and constant consumption (see Appendix D for a proof) when employed, c^e , may differ from his wage. The agent's resulting discounted utility is $w^e = u(c^e)/(1 - \beta)$. The net transfer per period from the principal to the agent is $c^e - b = u^{-1}((1 - \beta)w^e) - b$, which we denote by $G(w^e)$. Intuitively speaking, if the principal promises the agent a continuation utility of w^e in the employment state, then she must provide him with consumption c^e . The average discounted continuation cost to the principal is then $G(w^e) = c^e - b$. Note that for $c^e < b$ the continuation cost is negative and the principal levies a tax on the employed agent.

Finally, we can state the principal's recursive problem when the agent is unemployed. Suppose the principal needs to offer a contract that guarantees the agent an average discounted utility of w . The principal wants to minimize the cost, $C(w)$, of offering such a contract. Then the principal faces the following cost minimization problem:

$$\begin{aligned} C(w) &= \min_{a, c, w^e, w^u} c + \beta (p(a)G(w^e) + (1 - p(a))C(w^u)) & (7) \\ &w = u(c) - a + \beta (p(a)w^e + (1 - p(a))w^u) \\ &a \in \arg \max_{\tilde{a} \in A} u(c) - \tilde{a} + \beta (p(\tilde{a})w^e + (1 - p(\tilde{a}))w^u) \\ &a \in A, c \in C. \end{aligned}$$

We call C the cost function. Hopenhayn and Nicolini (1997) observe that for $w^e > w^u$ the agent's objective function is strictly concave in his decision variable a . Therefore, the first-order conditions for his optimization problem are necessary and sufficient. Hopenhayn and Nicolini (1997, p. 424, (15)) implicitly assume the existence of an interior solution and state the first-order condition as an equation,

$$\beta p'(a)(w^e - w^u) = 1. \quad (8)$$

Since the probability function p is strictly concave and satisfies an Inada condition, this implicit assumption implies that equation (8) has a unique positive solution whenever $w^e - w^u > 0$. This observation motivates Hopenhayn and Nicolini (1997) to replace the incentive compatibility constraint in the principal's recursive problem (7) by the first-order condition (8). However, such a replacement is only valid as long as the optimal solution to the principal's problem satisfies $w^e > w^u$. Moreover, the probability function in their calibrated model does not satisfy an Inada condition and so they should impose a positive lower bound on $w^e - w^u$; see Ljungqvist and Sargent (2000, p. 433).

4.2 Theoretical Properties

Under the assumptions of Hopenhayn and Nicolini (1997), the domains of both the principal and the agent are unbounded and, as a result, the dynamic principal-agent problem could be ill-defined. For a sound theoretical and numerical analysis of the model we need to impose somewhat different technical assumptions.

Assumption 2. *We impose the following assumptions:*

- (I) *The discount factor satisfies $\beta \in (0, 1)$.*
- (II) *The effort set is $A = \mathbb{R}_+$.*
- (III) *The consumption set is $C = \mathbb{R}_+$.*
- (IV) *The Bernoulli utility function $u : C \rightarrow \mathbb{R}$ is strictly increasing, strictly concave, twice differentiable, and satisfies $u(0) = 0$.*
- (V) *The probability function $p : A \rightarrow [0, 1]$ is increasing, strictly concave, twice differentiable, and satisfies $p(0) = 0$.*
- (VI) *The wage of an employed agent is non-negative, $b \geq 0$.*

In addition, we impose two economic assumptions:

(VII) *An unemployed agent has the outside option of autarchy.*

(VIII) *The principal may tax an employed agent but does not provide payments to him.*

Conditions (I)–(VII) of Assumption 2 are adapted from Ljungqvist and Sargent (2000). The technical conditions (I)–(III) are essentially identical to those of Hopenhayn and Nicolini (1997). We impose the three requirements (IV)–(VI), which deviate from those of Hopenhayn and Nicolini (1997), because they are satisfied by the functions and parameters of the calibrated economy in Hopenhayn and Nicolini (1997). Ljungqvist and Sargent (2000) impose the economic assumption (VII) to obtain a lower bound on the possible continuation values. Finally, our second economic assumption (VIII) deviates from both Hopenhayn and Nicolini (1997) and Ljungqvist and Sargent (2000). Hopenhayn and Nicolini (1997) allow the principal to make payments to an employed agent. Then, however, the state space may theoretically be unbounded. Ljungqvist and Sargent (2000) assume that the principal cannot tax an employed agent and so the agent just consumes his wage. Their assumption greatly simplifies the analysis since it determines the value of w^e .

Condition (VII) of Assumption 2 implies that an agent can achieve a minimal average utility level w^{aut} that solves the following Bellman equation:

$$w^{aut} = \max_a u(0) - a + \beta (p(a)w^{e,b} + (1 - p(a))w^{aut}), \quad (9)$$

with the discounted life-time utility of an employed agent in autarchy being $w^{e,b} = u(b)/(1 - \beta)$ since he consumes his entire wage b each period. Equation (9) implies $0 \leq w^{aut} \leq u(0) + \beta w^{e,b} < w^{e,b}$ and so the optimal time-invariant effort level of an agent in autarchy satisfies the first-order condition⁹

$$\beta p'(a^{aut})(w^{e,b} - w^{aut}) = 1. \quad (10)$$

The principal must offer the agent a utility of at least w^{aut} at the beginning of the contract, but she has no restrictions during the contract. Once the principal and the agent enter a contract at time 0, they cannot walk away from it. (Below, corresponding to the model without commitment in Section 2.4, we also analyze the set-up in which the agent can walk away from the contract at any time. In that set-up, we must impose the lower bound w^{aut} on the agent's continuation utility for the entire duration of the contract.) Finally, we observe that assumption (VIII) implies that the principal cannot pay an employed agent; consequently, she can only offer him a continuation utility of at most $w^{e,b} = u(b)/(1 - \beta)$ in the employment state. In sum, conditions (VII,VIII) of Assumption 2 imply $c^e \in [0, b]$, $w^e \in [u(0)/(1 - \beta), u(b)/(1 - \beta)]$, and $w^u \in [u(0)/(1 - \beta), u(b)/(1 - \beta)]$ for the model with commitment (and $w^u \in [w^{aut}, u(b)/(1 - \beta)]$ for the model without).

⁹This first-order condition relies on the implicit assumption that $\beta p'(0)w^{e,b} \geq 1$.

Proposition 1. *Under Assumption 2, the first-order condition (8) of Hopenhayn and Nicolini (1997) cannot hold on an entire compact domain $W = [\underline{w}, \bar{w}]$ of the cost function $C : W \rightarrow \mathbb{R}$.*

Proof. Suppose $W = [\underline{w}, \bar{w}]$. The upper bound on w leads to an upper bound \bar{c} on consumption via $u(c) \leq u(\bar{c}) = \bar{w}(1 - \beta)$. Recall that any future continuation utilities, including U^e, U^u , must be elements of W . As a result we obtain the following restriction:

$$u(c) - a + \beta(p(a)U^e + (1 - p(a))U^u) \leq u(c) - a + \beta\bar{w} < \bar{w} \quad \forall a \in A \setminus \{0\}.$$

If we evaluate the cost function C at the largest value of its domain—that is to say at \bar{w} —then it must hold that $a = 0$. Thus, $U^u = \bar{w}$ and so, since $U^e \in W$, we immediately obtain $U^e \leq U^u$, and so with $p'(a) > 0$ the first order conditions (8) of Hopenhayn and Nicolini (1997)

$$\beta p'(a)(U^e - U^u) = 1$$

cannot hold for $w = \bar{w}$. □

Proposition 1 shows that the first-order condition (8) of Hopenhayn and Nicolini (1997) cannot hold on a compact domain. Therefore, the numerical approach of Hopenhayn and Nicolini (1997) cannot work under Assumption 2 unless we make further simplifying assumptions on endogenous variables as done in Ljungqvist and Sargent (2000). Here we do not pursue this route and instead apply our method.

4.3 Quantitative Analysis

We now solve a slightly modified version of the model specification from Hopenhayn and Nicolini (1997, Section III.B).¹⁰ We examine both the case with and the case without commitment.

To reflect the properties of the probability function in Hopenhayn and Nicolini (1997), we choose

$$p(a) = 1 - \frac{1}{1 + r \cdot a},$$

with $a \in [0, \infty)$. We slightly alter the utility function so that $u(0) = 0$ and choose

$$u(c) = \frac{(c + 0.5)^{1-\sigma}}{1 - \sigma} - \frac{(0.5)^{1-\sigma}}{1 - \sigma},$$

where $\sigma \in \mathbb{N} \setminus \{1\}$ and the value 0.5 denotes the agent's home production. The remaining parameters are $\beta = 0.98$, $b = 1$, and $r = 1$. We adopt $\sigma = 0.5$ from Hopenhayn and Nicolini

¹⁰Hopenhayn and Nicolini (1997, Section III.B) assume $\beta = 0.999$ and $p(a) = 1 - e^{-ra}$ with an unreported value of r that needs to be $r = 0.000141$ for their hazard rate of 10% for reemployment of unemployed workers. Needless to say, these two parameter values make the problem numerically extremely unstable.

(1997) and so obtain $W = [0, 51.7638]$. The first-order condition (10) for autarchy yields $w^{aut} = 47.7736$ and $a^{aut} = 0.977482$.

We first analyze the model with commitment. We choose 200 interpolation nodes and a piecewise linear interpolation. This enables us to approximate the cost function C on the domain W with an ℓ^∞ error of 0.0002. Figure 3 shows the graph of the cost function. For the




Figure 3: Cost function. At $w^{aut} = 47.7735$ the cost is $C(w^{aut}) = -0.516485$.

principal to convince the agent to enter into the insurance contract, the initial utility promise must be at least w^{aut} . Since the cost function is increasing for $w \geq w^{aut}$, it is optimal for the principal to start with a utility promise of $w = w^{aut}$ at a cost of $C(w^{aut}) = -0.516485$.

Figure 4 depicts the continuation value $U^u(w)$ for an unemployed agent and the identity function of the utility promise w . The difference $U^u(w) - w$ is nonpositive on the entire domain and strictly negative in its interior. Therefore, an unemployed agent experiences a (weakly) decreasing utility from period to period as long as he remains unemployed. Figure 5 displays the accompanying consumption policy and Figure 6 shows the resulting action of the agent.

Iterating over the function values $U^u(w)$ we observe the equilibrium outcomes over time as long as the agent remains unemployed. Starting from w^{aut} his continuation utility decreases. His consumption over time is strictly decreasing until it hits zero. In fact, after 46 periods of perpetual unemployment his consumption becomes zero and remains so as long as he is unemployed. While the agent's effort level increases over time at first, it eventually decreases steeply and becomes zero at $w = \underline{w}$. As a result, $w = \underline{w}$ is an absorbing state and the agent remains unemployed forever. In the case of employment the agent pays the principal an in-




Figure 4: Continuation utility of the unemployed agent.




Figure 5: Consumption of unemployed agent as a function of $w \in W$.




Figure 6: Agent’s action as a function of $w \in W$.

creasing part of his wage in taxes; see Figure 7. Put differently, the longer the agent has been unemployed when he finally regains employment, the higher is his unemployment tax and the lower is his consumption in the employment state.

Now we consider the case without commitment; the principal must provide at least w^{aut} to the agent in each period and so the state space of continuation utilities shrinks to $\tilde{W} = [w^{aut}, \bar{w}] = [47.7736, 51.7638]$. In this case the cost function is much simpler; see Figure 8. As a result, we only require 50 Chebyshev nodes to obtain an ℓ^∞ error of less than 0.0002.

Unlike in the previous case with commitment, the cost function is strictly monotone on its domain. Not surprisingly, the cost function is minimized at the smallest element of \tilde{W} —that is, at $w = w^{aut}$ with an optimal value of $C(w^{aut}) = -0.306094$. Naturally, the principal offers the agent only the autarchy value as continuation value. As a result, contrary to the other case, all variables are constant over time as long as the agent is unemployed. The continuation utility is constant at $w_0 \equiv U^u(w^{aut}) = w^{aut}$ with the resulting consumption $c^u(w_0) = 0.405530$ and effort $a(w_0) = 0.683007$. When the agent becomes employed, his consumption is $c^e(w_0) = 0.973177$ —that is, he pays an unemployment tax of $b - c^e(w_0) = 0.026823$.

We complete the discussion of the model without commitment with some sensitivity analysis. Table 2 reports the equilibrium outcomes as a function of the agent’s utility parameter σ . The more risk averse the agent, the lower are the costs for the principal. As we would




Figure 7: Consumption of the agent when employed as a function of $w \in W$.




Figure 8: Cost function for $\tilde{W} = [w^{\text{aut}}, \bar{w}]$ (without commitment).

σ	$C(w_0)$	$w_0 = U^u(w_0)$	$c^u(w_0)$	$c^e(w_0)$	$a(w_0)$
0.5	-0.3061	47.77	0.4056	0.9732	0.6830
2	-3.092	62.10	0.5241	0.9006	0.4675
3	-6.971	83.53	0.4967	0.8165	0.4170
4	-12.48	121.7	0.4325	0.7035	0.4226
5	-18.53	188.7	0.3573	0.5841	0.4865
6	-23.87	306.4	0.2917	0.4811	0.6086

Table 2: Sensitivity analysis for $\tilde{W} = [w^{aut}, \bar{w}]$ (without commitment).

expect, the one-period ahead consumption volatility is decreasing in σ . For $\sigma \geq 3$, as his risk aversion increases further, the agent is willing to forgo consumption in both states, u and e , for a reduction in volatility. Moreover, then the agent's optimal effort level is increasing. He wants to increase the probability of being employed, which removes all consumption volatility.

5 Conclusion

The contributions of this paper are threefold. First, we have provided a recursive formulation for discrete time principal-agent problems that allows for a computational treatment. We did this by removing a redundant constraint from the previous description, reworking the original proof in Spear and Srivastava (1987). Second, we have shown that the Bellman operator for this problem has the contraction mapping property. This property shows the existence of a solution and justifies the use of value function iteration for these kinds of problems. Lastly, we have combined these results with the polynomial reformulation approach to static principal-agent problems, see Renner and Schmedders (2015), to solve for explicit solutions for two prominent social planning models from the economics literature.

The numerical experiments have shown us that small changes in the assumptions of the game can have serious consequences for the economic results. Hopenhayn and Nicolini (1997) emphasize a monotonically decreasing payoff schedule for their unemployment insurance problem. However, if we assume that the principal has to provide a reservation utility to the agent then the payoff scheme exhibits constant parts. These results show that we need to proceed carefully when dealing with incentive constraints in dynamic environments. Solving these problems without simplifying assumptions is essential if we want to obtain correct and economically meaningful solutions. Our analysis shows that using the sophisticated methods developed in real algebraic geometry can be of great help to economists working on dynamic incentive models.

Appendix

A Additional Proofs for Section 2

Proof of Lemma 3. $\mathcal{F}(X)$ is a subset of the set of bounded functions, which is a complete metric space with the supremum norm $\|\cdot\|$ defined by $\|f\| = \sup\{|f(x)| : x \in X\}$. Thus it suffices to show that $\mathcal{F}(X)$ is closed. For this let $\{f_k\}$ be a sequence of functions $f_k : X \rightarrow \mathbb{R}$ in $\mathcal{F}(X)$ with limit f (which is bounded). It remains to be shown that f is upper semi-continuous for any $x_0 \in X$.

Let $\varepsilon > 0$ and $x_0 \in X$ be given. Since the sequence $\{f_k\}$ converges to f , there is a sufficiently large k_0 such that $\|f - f_{k_0}\| < \varepsilon/3$. Since f_{k_0} is upper semi-continuous there is, by definition, a $\delta > 0$ such that $f_{k_0}(x) - f_{k_0}(x_0) \leq \varepsilon/3$ for all $\|x - x_0\| < \delta$. We obtain the following bound:

$$\begin{aligned} f(x) - f(x_0) &= f(x) - f(x_0) + f_{k_0}(x) - f_{k_0}(x) + f_{k_0}(x_0) - f_{k_0}(x_0) \\ &= f(x) - f_{k_0}(x) + f_{k_0}(x_0) - f(x_0) + f_{k_0}(x) - f_{k_0}(x_0) \\ &< \varepsilon \end{aligned}$$

for all $\|x - x_0\| < \delta$. The first and second differences are each less than $\varepsilon/3$ because of the convergence with the supremum norm in the complete metric space of bounded functions; the last difference is less than or equal to $\varepsilon/3$ because of the upper semi-continuity of f_{k_0} . Thus f is upper semi-continuous. This completes the proof that $\mathcal{F}(X)$ is closed. \square

Proof of Theorem 2. Lemma 3 implies that $\mathcal{F}(X)$ is a complete metric space. Hogan (1973, Theorem 5), or alternatively Berge (1963, Theorem 2, p. 116), ensures that the function Tf is indeed upper semi-continuous, and so the operator T maps $\mathcal{F}(X)$ to itself. The operator also satisfies, by construction, Blackwell's sufficient condition for a contraction mapping; see Stokey and Lucas (1989, Theorem 3.3). Therefore, the contraction mapping theorem, see Stokey and Lucas (1989, Theorem 3.2), implies the existence and uniqueness of a fixed point. \square

B Additional Results for Section 3

We report additional results for the model of Phelan and Townsend (1991) analyzed in Section 3. Phelan and Townsend (1991) also consider an infinitely repeated full information economy. Table 3 reports the principal's value function and the optimal values for the action a and the compensation (consumption) c for the full information economy for $\beta = 0.8$ (at 8 Chebyshev nodes) using the agent's exact utility function (instead of the polynomial approximation). Similar to the value function for the full information model in Phelan and Townsend


Table 3: Results for the full information model with $\beta = 0.8$.

w	10.144	11.264	13.333	16.037	18.963	21.667	23.736	24.856
V	7.728	7.346	6.416	4.708	2.171	-0.849	-3.611	-5.394
a	0.499	0.439	0.346	0.257	0.188	0.140	0.111	0.029
c	0.094	0.142	0.276	0.550	0.990	1.536	2.047	2.25

The table reports results for the infinitely repeated full information version of the continuous action model of the social planning problem in Phelan and Townsend (1991). The table shows the principal's value function V and the optimal values for the action a and the compensation (consumption) c for the full information economy for $\beta = 0.8$ (at 8 Chebyshev nodes) using the agent's exact utility function (instead of the polynomial approximation).

(1991, Section IV), the value function is strictly decreasing in the utility promise w to the agent. Figure 9 shows the graph of the value function V both for $\beta = 0.8$ and for $\beta = 0.95$. For $\beta = 0.95$, the optimal values for a and c are identical to those for $\beta = 0.8$, and the values of w and V are simply four times the corresponding values for $\beta = 0.8$.

Figure 9: Plot of the value function V for the full information model.



The left graph shows the value function V for the full information model on its domain $[10, 25]$ for $\beta = 0.8$. The right graph shows V on $[40, 100]$ for $\beta = 0.95$.

C Barycentric Lagrange Interpolation

This form of interpolation is ideal for the one-dimensional case. It allows a numerically stable and efficient approximation of functions with thousands of interpolation nodes. We only look at a special case of this approximation method, which uses Chebyshev nodes. We also restrict ourselves without loss of generality to the case of approximating a function on $[-1, 1]$. The details can be found in Berrut and Trefethen (2004).

We use the Chebyshev nodes x_j of the second kind, which are defined as

$$x_j = \cos \frac{j\pi}{n}, \quad j = 0, \dots, n.$$

These are the extreme points of the degree- n Chebyshev polynomial of the second kind. In particular these contain the two boundary points $-1, 1$. Also note that unlike the nodes of the first kind, these points allow us to reuse data from a coarser approximation, if the approximation order is chosen accordingly.

Given the function values f_j we get the following interpolation formula for the interpolant p .

$$p(x) = \frac{\sum_{j=0}^n \frac{w_j}{x-x_j} f_j}{\sum_{j=0}^n \frac{w_j}{x-x_j}},$$

where w_j are the barycentric weights. In general these weights can be difficult expressions however in the Chebyshev case they simplify to

$$w_j = (-1)^j \delta_j, \quad \delta_j = \begin{cases} 1/2, & j = 0 \text{ or } j = n, \\ 1, & \text{otherwise.} \end{cases}$$

It is also easy to see that the formula is invariant under a linear transformation. In particular it is independent of the interval being used.

To make the computation go smoothly the only thing needed is to check whether x is equal to one of the nodes. Then the formula is numerically stable and efficient; an evaluation costs $O(n)$ operations. For the details see the references in Berrut and Trefethen (2004).

D Additional Proof for Section 4.1

In this section we want to compute the terminal value function for the state of employment in the model by Hopenhayn and Nicolini (1997). We do this to help the reader understand where the description of the function is coming from.

Employment is an absorbing state. Here, however, the principal is allowed to extract taxes from the agent. Her value function in the case of employment is

$$G(w) = \min_{c \in C, U \in W^e} (1 - \beta)(c - \tau) + \beta G(U),$$

$$w = (1 - \beta)u(c) + \beta U,$$

where τ denotes the wage the agent receives and $C = [0, \tau]$. It is now easy to see that the state space $W^e = [u(0), u(\tau)]$. Since the function $(1 - \beta)u : C \rightarrow W^e$ is strictly monotone we can solve the constraint for c and therefore simplify the problem to

$$G(w) = \min_{U \in W^e} (1 - \beta) \left(u^{-1} \left(\frac{w - \beta U}{1 - \beta} \right) - \tau \right) + \beta G(U).$$

If u is twice continuously differentiable and strictly concave, then we know that the inverse is strictly convex. We claim that the following value function solves the above equation:

$$G(w) = u^{-1}(w) - \tau.$$

It is clear that for this function the right-hand side is convex. So the critical point can be determined by solving

$$(1 - \beta) \left((u^{-1})' \left(\frac{w - \beta U}{1 - \beta} \right) \left(-\frac{\beta}{1 - \beta} \right) \right) + \beta (u^{-1})'(U) = 0.$$

The derivative is strictly increasing and so the unique solution to this equation is w . In particular the constraints are not active here. Plugging this back into the Bellman equation we see that it solves the problem.

References

- ABRAHAM, A. AND N. PAVONI (2008): “Efficient Allocations with Moral Hazard and Hidden Borrowing and Lending: A Recursive Formulation,” *Review of Economic Dynamics*, 11, 781–803.
- ABREU, D., D. PEARCE, AND E. STACCHETTI (1990): “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, 58, 1041–1063.
- BERGE, C. (1963): *Topological Spaces*, Oliver & Boyd: Edinburgh and London.
- BERRUT, J. AND L. TREFETHEN (2004): “Barycentric Lagrange Interpolation,” *SIAM Review*, 46, 501–517.
- DOEPKE, M. AND R. M. TOWNSEND (2006): “Dynamic Mechanism Design with Hidden Income and Hidden Actions,” *Journal of Economic Theory*, 126, 235 – 285.
- HOGAN, W. W. (1973): “Point-to-Set Maps in Mathematical Programming,” *SIAM Review*, 15, 591–603.
- HOPENHAYN, H. A. AND J. P. NICOLINI (1997): “Optimal Unemployment Insurance,” *Journal of Political Economy*, 105, 412–438.
- JIBETEAN, D. AND E. DE KLERK (2006): “Global Optimization of Rational Functions: A Semidefinite Programming Approach,” *Mathematical Programming*, 106, 93–109.
- JUDD, K. (1998): *Numerical Methods in Economics*, MIT Press.
- KOCHERLAKOTA, N. R. (2004): “Figuring out the Impact of Hidden Savings on Optimal Unemployment Insurance,” *Review of Economic Dynamics*, 7, 541–554.

- LAMBERT, R. A. (1983): “Long-term Contracts and Moral Hazard,” *Bell Journal of Economics*, 14, 441–452.
- LASSERRE, J. (2010): *Moments, Positive Polynomials and Their Applications*, vol. 1 of *Imperial College Press Optimization Series*, Imperial College Press.
- LAURENT, M. (2009): *Sums of Squares, Moment Matrices and Optimization Over Polynomials*, Springer, vol. 149 of *IMA Volumes in Mathematics and its Applications*, 157–270.
- LJUNGQVIST, L. AND T. J. SARGENT (2000): *Recursive Macroeconomic Theory*, Cambridge, MA: MIT Press.
- MIRRLIERS, J. A. (1999): “The Theory of Moral Hazard and Unobservable Behaviour: Part I,” *The Review of Economic Studies*, 66, 3–21.
- PHELAN, C. AND R. M. TOWNSEND (1991): “Computing Multi-period, Information-Constrained Optima,” *Review of Economic Studies*, 58, 853–81.
- RENNER, P. AND K. SCHMEDDERS (2015): “A Polynomial Optimization Approach to Principal-Agent Problems,” *Econometrica*, 83, 729–769.
- ROGERSON, W. P. (1985): “Repeated Moral Hazard,” *Econometrica*, 53, 69–76.
- SPEAR, S. E. AND S. SRIVASTAVA (1987): “On Repeated Moral Hazard with Discounting,” *The Review of Economic Studies*, 54, 599–617.
- STOKEY, N. L. AND R. E. LUCAS (1989): *Recursive Methods in Economic Dynamics*, Harvard University Press.
- WÄCHTER, A. AND L. T. BIEGLER (2006): “On the Implementation of an Interior-Point Filter Line-Search Algorithm for Large-Scale Nonlinear Programming,” *Mathematical Programming*, 106, 25–57.
- WERNING, I. (2001): “Repeated Moral-hazard with Unmonitored Wealth: A Recursive First-order Approach,” *working paper*.