

Renner, Philipp; Schmedders, Karl

Article

Discrete-time dynamic principal-agent models: Contraction mapping theorem and computational treatment

Quantitative Economics

Provided in Cooperation with:

The Econometric Society

Suggested Citation: Renner, Philipp; Schmedders, Karl (2020) : Discrete-time dynamic principal-agent models: Contraction mapping theorem and computational treatment, Quantitative Economics, ISSN 1759-7331, The Econometric Society, New Haven, CT, Vol. 11, Iss. 4, pp. 1215-1251,
<https://doi.org/10.3982/QE960>

This Version is available at:

<https://hdl.handle.net/10419/253564>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc/4.0/>

Discrete-time dynamic principal–agent models: Contraction mapping theorem and computational treatment

PHILIPP RENNER

Department of Economics, Lancaster University

KARL SCHMEDDERS

IMD Lausanne

We consider discrete-time dynamic principal–agent problems with continuous choice sets and potentially multiple agents. We prove the existence of a unique solution for the principal’s value function only assuming continuity of the functions and compactness of the choice sets. We do this by a contraction mapping theorem and so also obtain a convergence result for the value function iteration. To numerically compute a solution for the problem, we have to solve a collection of static principal–agent problems at each iteration. As a result, in the discrete-time setting solving the static problem is the difficult step. If the agent’s expected utility is a rational function of his action, then we can transform the bi-level optimization problem into a standard nonlinear program. 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 CLASSIFICATION. C63, D80, D82.

1. INTRODUCTION

The economics literature on repeated principal–agent problems has seen many advances in recent years. [Spear and Srivastava \(1987\)](#) introduced the core idea for the analysis of such models—namely the inclusion of an artificial state variable in the description of the solution. Many of the recent advances have occurred in the continuous-time setting (see, e.g., [Sannikov \(2008\)](#)), where solving the problem is reduced to solving

Philipp Renner: p.renner@lancaster.ac.uk

Karl Schmedders: karl.schmedders@imd.org

We are indebted to two anonymous referees for thoughtful reviews of earlier versions of this paper. 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, the University of Lancaster, and the 2016 NSF/NBER/CEME Conference on Mathematical Economics at Johns Hopkins for helpful comments. We thank Dave Brooks for a thorough editorial review of an earlier version. Philipp Renner gratefully acknowledges financial support from the Swiss National Science Foundation (SNSF) under grant 148769.

an ordinary differential equation. However, there is no one-to-one relationship between discrete-time and continuous-time problems, and so the insights from continuous-time models do not simply extend to discrete-time settings. In addition, the closed-form solutions of continuous-time models come at the price of rather difficult numerical problems, because solving the ensuing differential equations is rather difficult, particularly as the number of dimensions increases. This property severely limits the applicability of such models. A better understanding of repeated principal–agent relationships requires, therefore, the study of such models in both discrete and continuous time.

The present paper focuses on discrete-time dynamic principal–agent problems with compact choice sets and continuous functions. The resulting mathematical problem is a standard dynamic programming problem, which allows for solutions of large-scale instances using modern methods such as sparse grids (Brumm and Scheidegger (2017)), or machine learning as in Scheidegger and Bilonis (2017).¹ This problem structure makes it possible to include multiple agents or principals as well as additional state variables. As a drawback, we lose the ability to compute closed-form solutions, and so the policies are more difficult to interpret.

We close a gap in the economics literature by proving the existence of a unique solution for the principal's value function, which solves the dynamic programming problem in the recursive formulation. For this result, we only require minimal assumptions on the problem specifications, such as continuity of functions and compactness of the decision spaces. The theorem's proof, using a contraction mapping result, provides insights into the critical issues in these type of problems. As a result, it enables modelers to use it as a template for their version of the problem. As a corollary to the main result, we obtain that the Bellman equation resulting from the first-order approach is also a contraction mapping.

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.

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

¹Note that algorithms for solving dynamic programming problems can be easily parallelized. On the contrary, partial differential equations require sequential algorithms, and so their tractable size is far smaller.

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.

Zhao (2007) considers a two-agent model with a resource constraint. He assumes separability of the utility functions in consumption and effort as well as concavity of the consumption function. These assumptions enable him to develop a recursive formulation omitting the constraint involving the policy. He also proves, under a compactness condition on the state space, that the Bellman operator of the recursive problem is a contraction mapping and its fixed point is an upper semicontinuous function.

We show that results similar to those in Zhao (2007) hold in much more generality. We use results from the mathematical programming literature, so these proofs can be easily extended to different settings with multiple agents and with additional constraints as in Zhao (2007), or to cases where we have different preferences, for example, those of the Epstein–Zin type. The resulting dynamic-programming (DP) reformulation, however, 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.

As a way around the limitations of the FOA, ex post validation has been proposed to ensure global optimality after candidate solutions have been computed. It has, however, been known since Mirrlees (1999, first circulated in 1975) that the actual solution may not be among the critical points of the simplified FOA problem. In such a case, numerical optimization routines cannot find the true solution but ex post verification might still certify optimality. This problem gets compounded in the dynamic model since we have to solve many different parameterizations of the agent's problem. This model feature increases the likelihood of us coming across this technical problem. In addition, the FOA also precludes the use of constraints for the agent's problem, since this would lead to an automatic failure of many standard constraint qualifications (see Proposition 1.1 in Ye, Zhu, and Zhu (1997)).

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 Chapter 12 in [Judd \(1998\)](#)) with the polynomial solution method of [Renner and Schmedders \(2015\)](#) for static principal–agent models. The polynomial method allows us, similar 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 literature on economic incentives. 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 when we use similar assumptions to those used in the paper. To achieve this goal, the principal threatens the agent with a very large punishment in the distant future. In fact, when the game reaches this punishment state the agent no longer receives an insurance payment and the premium eats up all his earnings once he finds a job. We use our approach to also examine a bounded version of this problem, where we assume limited commitment on the agent's part. For example, if the agent receives less utility than he would in a world without insurance, then he walks away from the contract. We find that in this situation the optimal policy may feature a region in which the policy is constant over time. The levels of insurance payment and premium are then determined by the agent's risk aversion. Our results show that the entire state space of the recursive problem can play a role in the solution. In particular, it is essential to solve the problem on the entirety of the state space. This observation makes our assumption of compactness for the continuation payoff space essential for any computational treatment.

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 simplifying assumption for two reasons. First, there are very influential 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 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 \(2001\)](#) proposed a first-order approach for repeated moral hazard problems with unmonitored wealth of the agent. [Kocherlakota \(2004\)](#), however, showed analytically that the FOA must fail for an enhanced version of the [Hopenhayn and Nicolini \(1997\)](#) unemployment insurance model with hidden savings and linear disutility of search effort. He also points out that the recursive formulation of [Spear and Srivastava \(1987\)](#) cannot easily be extended to the case where we face hidden states. This is also demonstrated by [Fernandes and Phelan \(1999\)](#). They show that we would need an incentive constraint and a state for each type of agent. For a continuous state variable, this would mean infinitely many types and so an infinite-dimensional state space. [Doepke and Townsend \(2006\)](#) presented a recursive solution method for dynamic principal–agent environments with hidden states. They cannot, however, 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 exists for infinitely repeated principal–agent relationships with continuous choice sets and hidden states. The development of such a method is a task for the future, but is beyond the scope of the present paper. Nevertheless, we believe that the present paper is a significant step toward the 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 \(1997\)](#). In Section 5, we present a two-agent model. Section 6 concludes. [Appendices](#) provide additional proofs and results. In addition, an online supplement ([Renner and Schmedders \(2020\)](#)) provides additional results and a brief overview of the interpolation approach used to approximate the value functions, as well as a brief discussion of continuous-time models.

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 \mathcal{A} . The principal cannot monitor the agent’s actions but only the resulting out-

put. 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. *The principal–agent model satisfies the following conditions.*

- (i) *The set A of actions is a nonempty, compact subset of a finite-dimensional Euclidean space, $A \subset \mathbb{R}^L$.*
- (ii) *The set Y of outputs is a finite set of cardinality N .²*
- (iii) *The set C of compensations is non-empty and compact, $C = [0, \bar{c}]$.*
- (iv) *For all $y \in Y$, $f(y|\cdot): A \rightarrow [0, 1]$ is a continuous function.*
- (v) *For all $a \in A$, $\sum_{y \in Y} f(y|a) = 1$.*
- (vi) *The Bernoulli utility functions $v: \mathbb{R} \rightarrow \mathbb{R}$ and $u: A \times C \rightarrow \mathbb{R}$ are continuous.*
- (vii) *The principal and the agent have a common discount factor $\beta \in (0, 1)$.³*

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 of (at least) since the work of [Lambert \(1983\)](#) and [Rogerson \(1985\)](#). [Spear and Srivastava \(1987\)](#) showed 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 the 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 $c = (c_1, c_2, \dots)$ and $a = (a_1, a_2, \dots)$, respectively.

²We can extend our analysis to infinite output sets. Doing so, however, 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.

³This last condition, while standard, is not necessary for the theoretical framework.

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,

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

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

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

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 a position to define the *dynamic principal–agent problem*. The principal faces the following infinite-horizon utility maximization problem:⁴

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

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

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

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

An optimal compensation contract for the dynamic principal–agent problem is a sequence of compensation functions c^* together with a sequence of action functions a^* such that (c^*, a^*) is an optimal solution to the optimization problem (1), (2), (3). At the optimal solution (c^*, 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 supgame 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.

⁴In line with the notation for the measure μ , we place the argument 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 a would be to write \mathbb{E}^a . Here, we follow the notation in [Spear and Srivastava \(1987\)](#).

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 reprove 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 bounds

$$\underline{w} = \frac{1}{1-\beta} \min_{c \in C} \max_{a \in A} u(a, c) \quad \text{and} \quad \bar{w} = \frac{1}{1-\beta} \max_{a \in A, c \in C} u(a, c), \quad (4)$$

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:

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

Also assume that the function V is upper semicontinuous. Then $\max_{w \in W, w \geq \tilde{w}} V(w)$ is the optimal value of problem (1), (2), (3).⁵

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

⁵Since 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.

the recursive formulation (5) 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 \(1990\)](#).

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 3.

COROLLARY 1. *Suppose the assumptions of Theorem 1 hold. And let V , a , c , U be functions solving (5). Denote by w_0 one of the solutions of $\max_{w \in W, w \geq \tilde{w}} V(w)$. 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

In this paper, we mainly focus on standard principal–agent models with a single agent. However, our approach does extend to models with multiple agents. To emphasize this feature, in this subsection we immediately prove existence for models with one or more agents. For the model with a single principal and multiple agents, we assume that the agents are working together on a common project. While they can observe each other's actions, the principal cannot (just as in the single-agent setup).⁶

Denote the finite number of agents by $N_a \geq 1$. For multiple agents, the set of feasible utilities is no longer as easily determined as in the one-agent case; recall the bounds in expression (4). We can, however, use the results from [Kitti \(2016, Proposition 2\)](#) to obtain a convenient utility set. Under the assumption that utilities (payoffs) form a compact set, which is satisfied for our setting, there is a nonempty compact set $W \subset \mathbb{R}^{N_a}$ of all feasible utilities for the N_a agents.

We define A_i to be the compact set of actions for player i and $u_i(a_i, a_{-i})$ his utility function, which depends on the (vector of) actions a_{-i} of the other players. We want to

⁶The resulting set problem is similar yet slightly different to the one in [Abreu \(1988\)](#).

show existence of a solution to the following dynamic programming problem:

$$\begin{aligned}
 V(w) &= \max_{\hat{c}, \hat{U}, \hat{a}} \mathbb{E} \left[v \left(\bar{Y} - \sum_{i=1}^{N_a} \hat{c}_i(\bar{Y}) \right) + \beta V(\hat{U}(\bar{Y})) \mid \hat{a} \right], \\
 w_i &= \mathbb{E} [u_i(\hat{a}, \hat{c}_i(\bar{Y})) + \beta \hat{U}_i(\bar{Y}) \mid \hat{a}], \\
 \hat{a}_i &\in \arg \max_{a_i \in A_i} \mathbb{E} [u(a_i, \hat{a}_{-i}, \hat{c}_i(\bar{Y})) + \beta \hat{U}_i(\bar{Y}) \mid a_i, \hat{a}_{-i}], \\
 \hat{c}_i(\bar{Y}) &\in C, \quad \hat{U}_i \in W, \quad i = 1, \dots, N_a.
 \end{aligned} \tag{6}$$

We follow the dynamic programming literature to show existence and uniqueness of the value function and derive its properties; see, for example, [Stokey and Lucas \(1989, Section 4.2\)](#). However, our model does not satisfy the assumption of a continuous feasible set mapping; see [Stokey and Lucas \(1989, Assumption 4.3\)](#). Instead, we need to weaken this assumption to upper semicontinuity (sometimes also called “upper hemicontinuity”; see [Border and Aliprantis \(2006\)](#)).⁷ For this purpose, we first state two helpful mathematical results on upper semicontinuous bounded (real-valued) functions. [Appendix A](#) provides the proofs.

LEMMA 1. *Let X be a compact metric space and $\mathcal{F}(X)$ the set of upper semicontinuous, 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 $\Gamma(x) : X \rightarrow 2^Y$ an upper semicontinuous set-valued mapping. Further, let $\mathcal{F}(X)$ be the set of upper semicontinuous bounded functions equipped with the supremum norm. Also, let $F : X \times Y \rightarrow \mathbb{R}$ be an upper semicontinuous 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 (6) of the dynamic principal–agent problem. For this purpose, we need the following lemma:

LEMMA 2. *Define $A = \prod_{i=1}^{N_a} A_i$, $\tilde{C} = \prod_{i=1}^{N_a N} C$ and $\tilde{W} = \prod_{i=1}^{N_a N} W$. The feasible set mapping $\Gamma : W \rightarrow 2^{A \times \tilde{C} \times \tilde{W}}$ for the dynamic programming problem defined by*

$$\begin{aligned}
 \Gamma(w) &= \left\{ (a, c, U) \in A \times \tilde{C} \times \tilde{W} \mid \right. \\
 w_i &= \mathbb{E} [u_i(a, c_i(\bar{Y})) + \beta U_i(\bar{Y}) \mid a],
 \end{aligned}$$

⁷Due to our assumptions, the feasible set is contained in a compact subset of \mathbb{R}^n . Therefore, the closed-graph theorem applies, which states that the graph of the feasible set mapping is closed if and only if the mapping is upper semicontinuous; see [Border and Aliprantis \(2006, Theorem 17.11\)](#). Thus, these mappings are also called *closed*.

$$a_i \in \arg \max_{\hat{a}_i \in A_i} \mathbb{E}[u(\hat{a}_i, a_{-i}, \hat{c}_i(\bar{Y})) + \beta \hat{U}_i(\bar{Y}) \mid a_i, a_{-i}],$$

$$\hat{c}_i(\bar{Y}) \in C, \hat{U}_i \in W, i = 1, \dots, N_a \}$$

is upper semicontinuous and uniformly compact.

PROOF. We first rewrite Γ . Define the agents' value functions as

$$\Psi_i(a_{-i}, c_i, U_i) = \max_{\hat{a}_i \in A_i} \mathbb{E}[u(\hat{a}_i, a_{-i}, c_i(\bar{Y})) + \beta U_i(\bar{Y}) \mid a_i, a_{-i}].$$

Thus we have

$$\Gamma(w) = \left\{ (a, c(y), U(y)) \in A \times \tilde{C} \times \tilde{W} \mid \right.$$

$$w_i = \mathbb{E}[u_i(a, c_i(\bar{Y})) + \beta U_i(\bar{Y}) \mid a],$$

$$w_i = \Psi_i(a_{-i}, c_i, U_i),$$

$$\left. c_i(\bar{Y}) \in C, U_i \in W, a_i \in A_i, i = 1, \dots, N_a \right\}.$$

Since any a_i that solves the above conditions attains the same value for the value function Ψ_i , a_i is a maximizer.

Note that as A_i is a fixed compact set, it follows, from [Hogan \(1973, Theorem 7\)](#) or [Border and Aliprantis \(2006, Theorem 17.31\)](#), that all functions Ψ_i , $i = 1, \dots, N_a$, are continuous. By our assumption, so is the expectation operator, which is just a sum and product of continuous functions. Thus, by [Hogan \(1973, Theorem 10\)](#), the set-valued mapping is upper semicontinuous. 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 a solution of the dynamic principal-agent model.

THEOREM 3. *The unique value function $V : W \rightarrow \mathbb{R}$ in the recursive formulation (6) is an element of the complete metric space $\mathcal{F}(W)$ of upper semicontinuous bounded functions on the compact domain W .*

PROOF. The sets W and $A \times C \times W$ are compact metric spaces. Lemma 2 states that the feasible region for the dynamic optimization problem is an upper semicontinuous, uniformly compact, set-valued mapping. Theorem 2 now implies that the dynamic programming problem (6) has a unique fixed point V ; the value function is upper semicontinuous 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 (6) is only upper semicontinuous and may not be continuous; in general, the constraint mapping will

not be continuous because the lower semicontinuity condition fails. Unfortunately, sufficient conditions for lower semicontinuity 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 $V^* = \max_{w \geq \tilde{w}} V(w)$.*

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 semicontinuous value function V for the DP problem (6) exists, then $V^* = \max_{w \geq \tilde{w}} V(w)$ is the optimal value for the original sequential dynamic principal-agent problem (1), (2), (3). Theorem 3 shows that the DP problem (6) possesses indeed a unique upper semicontinuous 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 (6). We describe our numerical procedure in Section 2.6, below.

Even though the so-called first-order approach for solving static principal-agent problems may lead to incorrect solutions in the absence of stringent assumptions (see, e.g., Mirrlees (1999) and Renner and Schmedders (2015)), it remains very popular in economics applications. Because of its popularity, we also briefly address this approach here for the repeated principal-agent model. Note that the standard DP assumptions for this model are also violated if we replace the agents' utility maximization problems in the original formulation (6) by their respective first-order conditions. For simplicity, we consider only the case in which the agents are unconstrained. Under the assumption that the agents' objective functions are continuously differentiable, we obtain the following result.

COROLLARY 3. *Consider the first-order relaxation of the principal-agent problem,*

$$\begin{aligned} V(w) &= \max_{\hat{c}, \hat{U}, \hat{a}} \mathbb{E} \left[v \left(\bar{Y} - \sum_{i=1}^{N_a} \hat{c}_i(\bar{Y}) \right) + \beta V(\hat{U}(\bar{Y})) \mid \hat{a} \right], \\ w_i &= \mathbb{E} [u_i(\hat{a}, \hat{c}_i(\bar{Y})) + \beta \hat{U}_i(\bar{Y}) \mid \hat{a}], \\ \underline{0} &= \nabla_{a_i} \mathbb{E} [u(a_i, \hat{a}_{-i}, \hat{c}_i(\bar{Y})) + \beta \hat{U}_i(\bar{Y}) \mid a_i, \hat{a}_{-i}] \Big|_{a_i = \hat{a}_i} \\ \hat{c}_i(\bar{Y}) &\in C, \quad \hat{U}_i \in W, \quad i = 1, \dots, N_a. \end{aligned} \tag{7}$$

There exists a unique value function $V : W \rightarrow \mathbb{R}$ solving the recursive formulation (7). The value function is an element of the complete metric space $\mathcal{F}(W)$ of upper semicontinuous bounded functions on the compact domain W .

2.4 Outside option

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 a modifications of the original model. We weaken the assumption on the agent's commitment to a contract.

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 (5)), 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.⁸ The new assumption results in the following variational problem with $w \in \tilde{W}$:

$$\begin{aligned} V(w) &= \max_{\hat{c}, \hat{U}, \hat{a}} \mathbb{E}[v(\bar{Y} - \hat{c}(\bar{Y})) + \beta V(\hat{U}(\bar{Y}) \mid \hat{a})], \\ w &= \mathbb{E}[u(\hat{a}, \hat{c}(\bar{Y})) + \beta \hat{U}(\bar{Y}) \mid \hat{a}], \\ \hat{a} &\in \arg \max_{a \in A} \mathbb{E}[u(a, \hat{c}(\bar{Y})) + \beta \hat{U}(\bar{Y}) \mid a], \\ \hat{c}(\bar{Y}) &\in C, \quad \hat{U}(\bar{Y}) \in \tilde{W}. \end{aligned} \tag{8}$$

Observe that the only, but clearly important, difference between the modification (8) and the original DP problem (5) 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, 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.)

2.5 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.

⁸Seen as a game, both players are still committed to the strategies chosen at time $t = 1$. 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.

Spear and Srivastava (1987, p. 604) provided 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 without commitment for the principal 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) claimed and proved that the function V is concave in the state variable w . But the concavity proof in Spear and Srivastava (1987, pp. 604–605) 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 1 shows the graph of f on the interval $[-2, 2]$. 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




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.

counterexample in the context of a dynamic principal–agent model satisfying the assumptions in their paper. In fact, we observed that during the value function iterations some iterates themselves were nonconcave, 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 features can be easily added to the model. In fact, 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 implementing certain natural modifications to 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 \(1987, p. 604\)](#) provided 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 pages 603–604 and in the reformulation using the first-order approach on page 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 (5) 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, e.g., [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 semicontinuous. In our economics applications, the value function always turned out to be continuous. And, if a value function were truly discontinuous, 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 (5) 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 $V^k(w)$:

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

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 $\mathbb{E}[v(\bar{Y} - \hat{c}(\bar{Y})) + \beta V^k(\hat{U}(\bar{Y})) \mid \hat{a}]$. 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 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) applied modern methods from algebraic geometry (see Laurent (2009) and Lasserre (2010)) and the global optimization approach for rational functions of Jibetean and De Klerk (2006) to solve static principal–agent models.

If the agent's objective function, $\mathbb{E}[u(a, \hat{c}(\bar{Y})) + \beta \hat{U}(\bar{Y}) \mid a]$, is a polynomial or rational function in a , 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, $\mathbb{E}[v(\bar{Y} - \hat{c}(\bar{Y})) + \beta V^k(\hat{U}(\bar{Y})) \mid \hat{a}]$, 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:

$$v(x) = x, \quad u(a, c) = 2\sqrt{c} + 2\sqrt{1-a}.$$

[Phelan and Townsend \(1991\)](#) defined 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.

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 the function u with a degree 2 Chebyshev polynomial,

$$2\sqrt{c} + 1.998 - 0.4467a^2 - 0.9491a.$$

The agents expected utility now appears as follows:

$$\begin{aligned} & \frac{0.9}{1+3a} (2\sqrt{c_1} + 1.998 - 0.4467a^2 - 0.9491a + \beta \hat{U}_1) \\ & + \frac{0.1+3a}{1+3a} (2\sqrt{c_2} + 1.998 - 0.4467a^2 - 0.9491a + \beta \hat{U}_2) \\ & = \frac{1.998 + 5.046a - 3.294a^2 - 1.340a^3 + 1.8\sqrt{c_1} + 0.2\sqrt{c_2} + 6a\sqrt{c_2} + 0.9\hat{U}_1\beta + 0.1\hat{U}_2\beta + 3a\hat{U}_2\beta}{1+3a}, \end{aligned}$$

where \hat{U}_i is the future utility promise if state i happens. Let p the numerator of this function, q the denominator, and the agent's constraint $g(a) = (0.6 - a)a \geq 0$.

Now we can apply our polynomial reformulation approach to this function. We need the following equation to hold:

$$p(a) - \rho q(a) = \sigma_0(a) + \sigma_1(a)g(a), \quad \forall a, \quad (9)$$

where ρ is a constant and σ_i are sum of squares polynomials. To ensure that σ_i is a sum of squares, we represent it by a symmetric matrix

$$\sigma_0(a) = \begin{pmatrix} 1 \\ a \\ a^2 \end{pmatrix} \begin{pmatrix} m_{0,0}^{(0)} & m_{1,0}^{(0)} & m_{2,0}^{(0)} \\ m_{1,0}^{(0)} & m_{1,1}^{(0)} & m_{2,1}^{(0)} \\ m_{2,0}^{(0)} & m_{2,1}^{(0)} & m_{2,2}^{(0)} \end{pmatrix} \begin{pmatrix} 1 \\ a \\ a^2 \end{pmatrix}, \quad \sigma_1(a) = \begin{pmatrix} 1 \\ a \end{pmatrix} \begin{pmatrix} m_{0,0}^{(1)} & m_{1,0}^{(1)} \\ m_{1,0}^{(1)} & m_{1,1}^{(1)} \end{pmatrix} \begin{pmatrix} 1 \\ a \end{pmatrix}.$$

The matrices formed by $m_{k,l}^{(i)}$ have to be positive semidefinite. Which is equivalent to their minors being nonnegative. Recall that two polynomials are equal if and only if their coefficients are equal. Therefore, equation (9) reduces to the following five equations:

$$\begin{aligned} 1.998 + 1.8\sqrt{c_1} + 0.2\sqrt{c_2} + 0.9\hat{U}_1\beta + 0.1\hat{U}_2\beta - \rho(1) &= m_{0,0}^{(0)}, \\ 5.046 + 6\sqrt{c_2} + 3\hat{U}_2\beta - \rho(3) &= 2m_{1,0}^{(0)} + 0.6m_{0,0}^{(1)}, \\ -3.294 &= 0.6m_{1,0}^{(0)} + m_{1,1}^{(0)} + 2m_{2,0}^{(0)} - m_{0,0}^{(1)} + 0.6m_{1,0}^{(1)}, \\ -1.340 &= -m_{1,0}^{(0)} + 2m_{2,1}^{(0)} - m_{1,0}^{(1)} + 0.6m_{1,1}^{(1)} = m_{2,2}^{(0)} - m_{1,1}^{(1)}. \end{aligned} \tag{10}$$

We now can state the optimization problem:

$$\begin{aligned} V(w) &= \max_{\hat{c}, \hat{U}, \hat{a}, m_{k,l}^{(i)}, \rho} \mathbb{E}[\bar{Y} - \hat{c}(\bar{Y}) + \beta V(\hat{U}(\bar{Y})) \mid \hat{a}], \\ w &= \mathbb{E}[u(\hat{a}, \hat{c}(\bar{Y})) + \beta \hat{U}(\bar{Y}) \mid \hat{a}], \\ \text{equations (10) hold,} \\ M^{(0)}, M^{(1)} &\text{ positive semidefinite,} \\ w = \rho, \quad \hat{c}(\bar{Y}) &\in C, \quad \hat{U}(\bar{Y}) \in W, \quad \bar{Y} \in \{1, 2\}. \end{aligned}$$

We approximate the principal's value function V 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. (In an online supplement, we provide a long table with 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 decreases monotonically.




FIGURE 2. The social planner's 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 \(1997\)](#).⁹ 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 the reader 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 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 \(1997\)](#) imposed the following technical assumptions.

⁹The notation in this section is self-contained. We largely adopt the notation of [Hopenhayn and Nicolini \(1997\)](#).

ASSUMPTIONS (Hopenhayn and Nicolini (1997)). *The unemployment insurance model satisfies the following conditions.*

- (i) *The discount factor satisfies $\beta < 1$.*
- (ii) *The set $A \subset \mathbb{R}$ of effort levels is a closed interval and $0 \in A$.*
- (iii) *The consumption set is $S = \mathbb{R}_+$.*
- (iv) *The Bernoulli utility function $u : S \rightarrow \mathbb{R}$ is increasing and concave.*
- (v) *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 B 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, $\tilde{C}(w)$, of offering such a contract. Then the principal faces the following cost minimization problem:

$$\begin{aligned}\tilde{C}(w) &= \min_{a, c, w^e, w^u} c + \beta(p(a)G(w^e) + (1 - p(a))\tilde{C}(w^u)), \\ w &= u(c) - a + \beta(p(a)w^e + (1 - p(a))w^u), \\ a &\in \arg \max_{\tilde{a} \in A} u(c) - \tilde{a} + \beta(p(\tilde{a})w^e + (1 - p(\tilde{a}))w^u), \\ a &\in A, \quad c \in C.\end{aligned}\tag{11}$$

We call \tilde{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,

$$\beta p'(a)(w^e - w^u) = 1.\tag{12}$$

Since the probability function p is strictly concave and satisfies an Inada condition, this implicit assumption implies that equation (12) 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 (11) by the first-

order condition (12). 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 $S = \mathbb{R}_+$.*
- (iv) *The Bernoulli utility function $u : S \rightarrow \mathbb{R}$ 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 nonnegative, $b \geq 0$.*

In addition, we impose two economic assumptions:

- (vii) *An unemployed agent has the outside option of autarky.*
- (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:

$$w^{\text{aut}} = \max_a u(0) - a + \beta(p(a)w^{e,b} + (1 - p(a))w^{\text{aut}}), \quad (13)$$

with the discounted lifetime utility of an employed agent in autarky being $w^{e,b} = u(b)/(1 - \beta)$ since he consumes his entire wage b each period. Equation (13) implies $0 \leq w^{\text{aut}} \leq u(0) + \beta w^{e,b} < w^{e,b}$ and so the optimal time-invariant effort level of an agent in autarky satisfies the first-order condition¹⁰

$$\beta p'(a^{\text{aut}})(w^{e,b} - w^{\text{aut}}) = 1. \quad (14)$$

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 into 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 setup in which the agent can walk away from the contract at any time. In that setup, 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 $w^u \in [w^{\text{aut}}, u(b)/(1 - \beta)]$ for the model without).

PROPOSITION 1. *Under Assumption 2, the first-order condition (12) of Hopenhayn and Nicolini (1997) cannot hold on an entire compact domain $W = [\underline{w}, \bar{w}]$ of the cost function $\tilde{C}: W \rightarrow \mathbb{R}$.*

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 \tilde{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 (12) 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 (12) 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.

¹⁰This first-order condition relies on the implicit assumption that $\beta p'(0)w^{e,b} \geq 1$.

4.3 Quantitative analysis

We now solve a slightly modified version of the model specification from [Hopenhayn and Nicolini \(1997, Section III.B\)](#).¹¹ We examine both the case with full commitment and the case with only one-sided commitment. A similar model with one-sided commitment where principals compete over agents but where the agents' actions do not influence the outcome can be found in [Krueger and Uhlig \(2006\)](#).

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 \(1997\)](#) and so obtain $W = [0, 51.7638]$. We also limit the consumption transfer to the agent to the wage he can earn. The first-order condition (14) for autarky yields $w^{\text{aut}} = 47.7736$ and $a^{\text{aut}} = 0.977482$.

We first analyze the model with full commitment. We choose 200 interpolation nodes and a piecewise linear interpolation. This enables us to approximate the cost function \tilde{C} on the domain W with an ℓ^∞ error of 0.0002. Figure 3 shows the graph of the cost function. For the principal to convince the agent to enter into the insurance contract,




FIGURE 3. The principal's cost \tilde{C} as a function of $w \in W$. The figure shows the graph of the cost function \tilde{C} of the principal. The horizontal axis shows the expected discounted utility w promised to the agent. At $w^{\text{aut}} = 47.7735$, the cost is $\tilde{C}(w^{\text{aut}}) = -0.516485$.

¹¹[Hopenhayn 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.




FIGURE 4. Continuation utility of the unemployed agent as a function of $w \in W$. The figure shows the 45-degree line and the graph of the continuation utility of the employed agent as a function of the expected discounted utility promised to the agent, w .

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

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 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 re-




FIGURE 5. Consumption of the unemployed agent and his action as a function of $w \in W$. The left graph in the figure shows the consumption of the unemployed agent as a function of w —the expected discounted utility promised to the agent. The right graph displays the agent's action as a function of w . The autarky utility level is $w^{\text{aut}} = 47.7736$.




FIGURE 6. Consumption of the agent when employed as a function of $w \in W$. The figure shows the graph of the consumption of the employed agent as a function of w —the expected discounted utility promised to the agent. The autarky utility level is $w^{\text{aut}} = 47.7736$.

mains 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 increasing part of his wage in taxes; see Figure 6. 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 with one-sided commitment; the principal must provide at least autarky w^{aut} to the agent in each period because otherwise the agent would be just as well off without the insurance. So the state space of continuation utilities shrinks to $\tilde{W} = [w^{\text{aut}}, \bar{w}] = [47.7736, 51.7638]$. In this case, the cost function is much simpler; see Figure 7. 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 $w = w^{\text{aut}}$ with an optimal value of $\tilde{C}(w^{\text{aut}}) = -0.306094$. Naturally, the principal offers the agent only the autarky value as the 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 $w_0 \equiv U^u(w^{\text{aut}}) = w^{\text{aut}}$ 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 $b - c^e(w_0) = 0.026823$.

Krueger and Uhlig (2006) show that if the agent can choose between a large number of insurance providers, the principal has return 0. Therefore, we choose a starting promise $w_0 > w^{\text{aut}}$ such that $\tilde{C}(w_0) = 0$ and start the simulation from this point. We find that $w_0 = 48.0853$ and the optimal policy is to transfer $c^u = 0.622557$ in consumption to the agent inducing an effort level of $a = 0.749938$. The consumption promise for the




FIGURE 7. Cost function for $\tilde{W} = [w^{\text{aut}}, \bar{w}]$ (outside option). The figure shows the graph of the cost function \tilde{C} of the principal for the model with one-sided commitment. The horizontal axis shows the expected discounted utility w promised to the agent. The domain for the model with the outside option is $\tilde{W} = [w^{\text{aut}}, \bar{w}] = [47.7736, 51.7638]$.

employed agent is $c^e = 0.978875$. In the next period, the contract is again the stationary solution at the autarky level. In contrast, [Mueller, Storesletten, and Zilibotti \(2019\)](#) obtain in a similar setting that effort is hump-shaped and promised utilities are decreasing over time. Our result is similar thus far in that the promised utility is decreasing, too. But since the stationary solution is reached after a single step, effort is not hump-shaped.

To obtain comparable results to those of [Mueller, Storesletten, and Zilibotti \(2019\)](#), we adjust the utility the agent can obtain from working. We assume the agent has to implement a fixed amount of effort a_{fix} and set this to 0.5. Therefore, the bounds on the utility promises change as well, to $[23.847, 51.764]$. The new value function is now below the original one in Figure 7 but the overall shape stays the same. The graph now intersects the axis closer to the upper bound than to the lower bound, namely at $w_0 = 48.2742$. We simulate the optimal contract and Figure 8 shows the effort level over time. Just as in [Mueller, Storesletten, and Zilibotti \(2019\)](#), the effort is hump-shaped and utility promises are decreasing.

We complete the discussion of the model with one-sided commitment with some sensitivity analysis. Table 1, below, 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 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. A TWO-AGENT MODEL

Recall that our existence proof in Section 2.3 applies to models with multiple agents. In this section, we describe the solution of a model with two agents. The agents collaborate




FIGURE 8. Simulation of effort if $a_{\text{fix}} = 0.5$ starting at $\tilde{C}(w_0) = 0$. The figure shows the effort level over time during a simulation of the model with $a_{\text{fix}} = 0.5$ starting at $\tilde{C}(w_0) = 0$.

on a common project for the principal. The agents' joint effort influences the output y and they can observe each other's effort level. On the contrary, the principal cannot observe the agents' actions but only the output. Therefore, the principal faces the following recursive problem:

$$\begin{aligned}
 V(w_1, w_2, y) &= \max_{a_i, c_i, \tilde{w}_1(y), \tilde{w}_2(y)} y - c_1 - c_2 + \beta \sum_{\tilde{y}} p(\tilde{y} | y, a_1, a_2) V(\tilde{w}_1(\tilde{y}), \tilde{w}_2(\tilde{y}), \tilde{y}), \\
 w_1 &= u(c_1, a_1) + \beta \sum_{\tilde{y}} p(\tilde{y} | y, a_1, a_2) \tilde{w}_1(\tilde{y}), \\
 w_2 &= u(c_2, a_2) + \beta \sum_{\tilde{y}} p(\tilde{y} | y, a_1, a_2) \tilde{w}_2(\tilde{y}), \\
 a_1 &\in \arg \max_{\tilde{a}_1 \in A} u(c_1, \tilde{a}_1) + \beta \sum_{\tilde{y}} p(\tilde{y} | y, \tilde{a}_1, a_2) \tilde{w}_1(\tilde{y}), \\
 a_2 &\in \arg \max_{\tilde{a}_2 \in A} u(c_2, \tilde{a}_2) + \beta \sum_{\tilde{y}} p(\tilde{y} | y, a_1, \tilde{a}_2) \tilde{w}_2(\tilde{y}), \\
 a_i &\in A, \quad c_i \in C, \quad w_i(\tilde{y}) \in W, \quad i \in 1, 2.
 \end{aligned}$$

TABLE 1. Sensitivity analysis for $\tilde{W} = [w^{\text{aut}}, \bar{w}]$ (outside option).

σ	$\tilde{C}(w_0)$	$w_0 = U^u(w_0)$	$c^u(w_0)$	$c^e(w_0)$	$a(w_0)$
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

Note: The table reports the equilibrium outcomes as a function of the agent's utility parameter σ for the model with one-sided commitment.

Note that—as long as we assume that the principal has sufficiently deep pockets—the set of feasible utility promises W will just be a box defined by the highest and lowest consumption values, \bar{c} , \underline{c} , and the agents' preferences; see conditions (4).

We use the principal–agent model from [Araujo and Moreira \(2001\)](#) and modify it for a dynamic setting. The first-order approach is known to fail in the original model. The set of possible outcomes is $Y = \{l, h\}$, where l (h) denotes a low (high) outcome. The probability of a high outcome is $a_1^3 + a_2^3$ with $0 \leq a_i \leq 0.9$, $a_1^3 + a_2^3 \leq 1$ and is independent of the current state. Note that the probability of the high (low) outcome is increasing (decreasing) in the effort a_i , $i = 1, 2$, of each agent. The two agents have identical utility functions defined by

$$u(c_i, a_i) = \sqrt{c_i} - a_i^2.$$

For the numerical exercise below, we choose the following parameter values: $l = 0$, $h = 2$, $\beta = 0.98$, $\bar{c} = 2$.

For solving the principal's dynamic programming problem—that is, for the approximation of her value function on the domain $W \times W \times Y$ —we use a state-of-the-art numerical solution procedure; namely the method of adaptive sparse grids by [Brumm and Scheidegger \(2017\)](#). Note that sparse grids are the multivariate version of a piecewise linear interpolation. At the time of writing, sparse grids can be used to approximate up to 50-dimensional functions and are also applicable to nondifferential functions. However, the kinks in the piecewise linear approximation lead to some “jumpiness” when we try to simulate the policy functions.

Figure 9 shows a plot of the value function for the low-outcome state, $y = l$. The principal–agent problem with two agents is degenerate on the boundary of the domain. If agent i 's promise is $w_i = 0$, then he shows zero effort, $a_i = 0$, receives zero consumption, $c_i = 0$, and all future promises are zero as well. Similar to the unemployment insurance model of [Hopenhayn and Nicolini \(1997\)](#) with a single agent, this model with two agents has an absorbing state at $(w_1, w_2) = (0, 0)$. The principal's value function attains its maximum at $(w_1, w_2) = (0.1382, 0.1382)$ for both discrete states, l and h .

To obtain the optimal policy, we simulate starting from the point in the state space where we found the maximum. The simulation results can be found in Appendix C. We observe that the strategy for both the low and the high state is to try and stay in the symmetric state. That is, both agents work and exert the same amount of effort and they receive the same promised utility. We also see modest fluctuations in the promise away from the symmetric solution, but since we are close the boundary where the problem degenerates, these variations are most likely due to the nature of the sparse-grid interpolation.

Lastly, we consider the case of one-sided commitment, in which the agents have an outside option of getting 0.1 in consumption. This option increases the lower bound on W to 15.8114. The resulting value function is plotted in Figure 10. The simulation results are in Appendix C.

The value function attains its maximum at the corner $(w_1, w_2) = (15.8114, 15.8114)$. In the low state, the principal implements some effort and transfers a small value consumption to the agents. We stay at the same point in the state space while we are




FIGURE 9. The principal's value function in state l for the two-agent problem. The figure shows the graph of the principal's value function in the low-outcome state, $y = l$. The function values $V(w_1, w_2, l)$ depend on the utility promises w_1 and w_2 to the two agents.




FIGURE 10. The principal's value function in state l for the model with one-sided commitment. The figure shows the graph of the principal's value function in the low-outcome state, $y = l$ for the model with one-sided commitment. The function values $V(w_1, w_2, l)$ depend on the utility promises w_1 and w_2 to the two agents.

in the low state. In the high state, the principal shifts to transfer more consumption over time until a reversal point is reached at which she returns to the starting point (15.8114, 15.8114) in the state space.

6. CONCLUSION

The contributions of this paper are threefold. First, we have provided a recursive formulation for discrete-time dynamic principal–agent problems that allows for a computational treatment. We did this by removing a redundant constraint from the previous description and 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 implies the existence of a solution and justifies the use of value-function iteration for these kinds of problems. Finally, we have combined these results with the polynomial reformulation approach to static principal–agent problems of [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 research on dynamic incentive models.

APPENDIX A: ADDITIONAL PROOFS FOR SECTION 2

A.1 *Proof of Theorem 1*

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:

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

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.¹² 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.

¹²If we assumed Y to be uncountably infinite, we would have to show that the expectation exists; that is, that U is indeed integrable over Y . However, for a finite sum such a proof is not necessary.

LEMMA 3. 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, $\forall w \in W$:

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

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 $c(w, y) \in C^N$ and $U(w, y) \in W^N$ 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 4. 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 \mathbf{a} we get the following:

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

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

PROOF OF THEOREM 1. Lemmas 3 and 4 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 3.¹³ 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 than 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 semicontinuous and the set W is compact. \square

A.2 Existence proof

PROOF OF LEMMA 1. $\mathcal{F}(X)$ 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

¹³The converse is not necessarily true, that is, an optimal strategy might not correspond to a recursive one. There will, however, 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.

$f_k : X \rightarrow \mathbb{R}$ 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 semicontinuous 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 semicontinuity of f_{k_0} . Thus f is upper semicontinuous. This completes the proof that $\mathcal{F}(X)$ is closed. \square

PROOF OF THEOREM 2. Lemma 1 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 semicontinuous, 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

APPENDIX B: 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

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

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.

APPENDIX C: NUMERICAL RESULTS

In this Appendix (Tables 2, 3, 4), we present results from simulations of the two-agent problem.

TABLE 2. Simulation of the full-commitment case starting in state l .

Iter.	obj	w_1	w_2	a_1	a_2	c_1	c_2	State
1	76.6977	0.1382	0.1382	0.807675	0.779216	0.442385	0.366243	l
2	78.6538	0.138105	0.153179	0.791192	0.796193	0.407931	0.364439	h
3	78.6302	0.138078	0.197277	0.809775	0.776947	0.446075	0.456331	h
4	78.7069	0.138622	0.13809	0.807746	0.779141	0.44306	0.384029	h
5	78.7074	0.138135	0.13813	0.804981	0.78209	0.436498	0.389775	h
6	78.7074	0.138136	0.138159	0.798169	0.789181	0.422234	0.403125	h
7	78.7055	0.138099	0.13874	0.769248	0.816732	0.365296	0.462877	h
8	78.6939	0.138124	0.13813	0.804405	0.782699	0.43525	0.36443	h
9	78.6778	0.138144	0.16019	0.831374	0.752065	0.495411	0.360408	h
10	78.7024	0.138164	0.138131	0.832144	0.751122	0.495403	0.33264	h

Note: The table reports results from a simulation of the two-agent problem with full commitment starting in the low-output state. Due to the approximation errors in the sparse-grid solution, the agents' effort levels oscillate around $a_1 = a_2 = 0.5^{1/3} \approx 0.79370$.

TABLE 3. Simulation of the full-commitment case starting in state h .

Iter.	obj	w_1	w_2	a_1	a_2	c_1	c_2	State
1	78.7058	0.1382	0.1382	0.810337	0.776336	0.448103	0.376827	h
2	78.6889	0.138128	0.139749	0.70257	0.867662	0.256231	0.588513	h
3	78.7075	0.138204	0.138138	0.793434	0.793967	0.412551	0.413217	h
4	78.7057	0.138125	0.138392	0.819989	0.765546	0.469314	0.358542	h
5	78.7066	0.138131	0.138351	0.810851	0.775775	0.449099	0.377841	h
6	78.7052	0.13815	0.138172	0.781425	0.805608	0.386705	0.437456	h
7	78.7052	0.139631	0.138465	0.798926	0.788406	0.425693	0.402677	h
8	78.7059	0.138149	0.138164	0.767199	0.818542	0.361472	0.465976	h
9	78.7059	0.138198	0.138242	0.818639	0.767089	0.466498	0.361484	h
10	78.7078	0.138047	0.13811	0.793484	0.793917	0.412453	0.41341	h

Note: The table reports results from a simulation of the two-agent problem with full commitment starting in the high-output state. Due to the approximation errors in the sparse-grid solution, the agents' effort levels oscillate around $a_1 = a_2 = 0.5^{1/3} \approx 0.79370$.

TABLE 4. Simulation of the one-sided-commitment case starting in state l .

Iter.	obj	w_1	w_2	a_1	a_2	c_1	c_2	State
1	20.101	15.3193	15.3193	0.898774	0.649488	0.424758	0.417135	l
2	21.3176	15.8012	15.4133	0.864486	0.707363	0.459117	0.440269	h
3	20.5984	16.205	15.5715	0.9	0	0.453048	0.323198	h
4	20.3303	16.6711	15.3193	0.004804	0.9	0.471182	0.462086	h
5	20.1752	16.3395	15.7645	0	0.9	0.462571	0.471136	h
6	20.0212	16.0798	16.134	0.9	0	0.491135	0.477957	h
7	19.8508	16.5775	15.7628	0	0.9	0.492467	0.484222	h
8	17.7522	16.1943	16.2139	0.9	0.00102	0.484784	0.474759	l
9	19.5443	16.8104	15.7475	0	0.9	0.497411	0.49546	h
10	19.3608	16.5101	16.1787	0	0.9	0.492651	0.507926	h
11	19.2776	16.1474	16.5985	0.9	0	0.504975	0.503631	h
12	17.085	17.0627	15.8197	0	0.9	0.507806	0.506081	l
13	18.9193	16.8735	16.1246	0	0.9	0.518599	0.518198	h
14	18.9293	16.2978	16.694	0.9	0	0.504579	0.516139	h
15	18.7479	16.9319	16.1855	0.9	0	0.511694	0.508099	h
16	18.661	17.381	15.7984	0.000463	0.9	0.514784	0.519786	h
17	16.5773	17.0156	16.2204	0.9	0.007549	0.512306	0.520451	l
18	18.493	17.471	15.8245	0.002284	0.9	0.526668	0.525187	h
19	18.4185	17.0892	16.2555	0.008616	0.9	0.529253	0.531935	h
20	16.3798	16.7307	16.6426	0.9	0	0.530686	0.520405	l

Note: The table reports results from a simulation of the two-agent problem with one-sided commitment starting in the low-output state l .

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 (4), 781–803. [1219]
- Abreu, D. (1988), "On the theory of infinitely repeated games with discounting." *Econometrica*, 56 (2), 383–396. [1223]

- Abreu, D., D. Pearce, and E. Stacchetti (1990), "Toward a theory of discounted repeated games with imperfect monitoring." *Econometrica*, 58 (5), 1041–1063. [1221, 1222, 1223]
- Araujo, A. and H. Moreira (2001), "A general Lagrangian approach for non-concave moral hazard problems." *Journal of Mathematical Economics*, 35 (1), 17–39. [1242]
- Berge, C. (1963), *Topological Spaces*. Oliver & Boyd, Edinburgh. [1225, 1245, 1247]
- Border, K. and C. D. Aliprantis (2006), *Infinite Dimensional Analysis—A Hitchhiker's Guide*, third edition. Springer, Berlin. [1224, 1225]
- Brumm, J. and S. Scheidegger (2017), "Using adaptive sparse grids to solve high-dimensional dynamic models." *Econometrica*, 85 (5), 1575–1612. [1216, 1242]
- Doepke, M. and R. M. Townsend (2006), "Dynamic mechanism design with hidden income and hidden actions." *Journal of Economic Theory*, 126 (1), 235–285. [1219]
- Fernandes, A. and C. Phelan (1999), "A recursive formulation for repeated agency with history dependence." Staff Report 259, Federal Reserve Bank of Minneapolis. [1219]
- Hogan, W. W. (1973), "Point-to-set maps in mathematical programming." *SIAM Review*, 15 (3), 591–603. [1225, 1226, 1247]
- Hopenhayn, H. A. and J. P. Nicolini (1997), "Optimal unemployment insurance." *Journal of Political Economy*, 105 (2), 412–438. [1218, 1219, 1227, 1229, 1233, 1234, 1235, 1236, 1237, 1242, 1244, 1247]
- Jibeteau, D. and E. De Klerk (2006), "Global optimization of rational functions: A semidefinite programming approach." *Mathematical Programming*, 106, 93–109. [1230]
- Judd, K. (1998), *Numerical Methods in Economics*. MIT Press, Cambridge, MA. [1218, 1229]
- Kitti, M. (2016), "Subgame perfect equilibria in discounted stochastic games." *Journal of Mathematical Analysis and Applications*, 435 (1), 253–266. [1223]
- Kocherlakota, N. R. (2004), "Figuring out the impact of hidden savings on optimal unemployment insurance." *Review of Economic Dynamics*, 7 (3), 541–554. [1219]
- Krueger, D. and H. Uhlig (2006), "Competitive risk sharing contracts with one-sided commitment." *Journal of Monetary Economics*, 53 (7), 1661–1691. [1237, 1239]
- Lambert, R. A. (1983), "Long-term contracts and moral hazard." *Bell Journal of Economics*, 14 (2), 441–452. [1216, 1220]
- Lasserre, J. (2010), *Moments, Positive Polynomials and Their Applications*. Imperial College Press Optimization Series, Vol. 1. Imperial College Press, London. [1230]
- Laurent, M. (2009), "Sums of squares, moment matrices and optimization over polynomials." In *Emerging Applications of Algebraic Geometry*, 157–270, Springer, New York, NY. [1230]
- Ljungqvist, L. and T. J. Sargent (2000), *Recursive Macroeconomic Theory*. MIT Press, Cambridge, MA. [1233, 1235, 1236]

- Mirrlees, J. A. (1999), “The theory of moral hazard and unobservable behaviour: Part I.” *Review of Economic Studies*, 66 (1), 3–21. [1217, 1226, 1230]
- Mueller, A., K. Storesletten, and F. Zilibotti (2019), “Sovereign debt and structural reforms.” *American Economic Review*, 109 (12), 4220–4259. [1240]
- Phelan, C. and R. M. Townsend (1991), “Computing multi-period, information-constrained optima.” *Review of Economic Studies*, 58 (5), 853–881. [1218, 1219, 1231, 1232]
- Renner, P. and K. Schmedders (2015), “A polynomial optimization approach to principal–agent problems.” *Econometrica*, 83 (2), 729–769. [1216, 1218, 1226, 1229, 1230, 1244]
- Renner, P. and K. Schmedders (2020), “Supplement to ‘Discrete-time dynamic principal–agent models: Contraction mapping theorem and computational treatment’.” *Quantitative Economics Supplemental Material*, 11, <https://doi.org/10.3982/QE960>. [1219]
- Rogerson, W. P. (1985), “Repeated moral hazard.” *Econometrica*, 53 (1), 69–76. [1216, 1220]
- Sannikov, Y. (2008), “A continuous-time version of the principal–agent problem.” *Review of Economic Studies*, 75 (3), 957–984. [1215]
- Scheidegger, S. and I. Bilonis (2017), “Machine learning for high-dimensional dynamic stochastic economies.” Preprint. Available at <https://ssrn.com/abstract=2927400>. [1216]
- Spear, S. E. and S. Srivastava (1987), “On repeated moral hazard with discounting.” *Review of Economic Studies*, 54 (4), 599–617. [1215, 1216, 1217, 1219, 1220, 1221, 1222, 1227, 1228, 1229, 1244]
- Stokey, N. L. and R. E. Lucas (1989), *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge, MA. [1224, 1247]
- 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 (1), 25–57. [1230]
- Werning, I. (2001), “Repeated moral-hazard with unmonitored wealth: A recursive first-order approach.” Working paper. [1219]
- Ye, J., D. Zhu, and Q. Zhu (1997), “Exact penalization and necessary optimality conditions for generalized bilevel programming problems.” *SIAM Journal on Optimization*, 7 (2), 481–507. [1217]
- Zhao, R. R. (2007), “Dynamic risk-sharing with two-sided moral hazard.” *Journal of Economic Theory*, 136 (1), 601–640. [1217]

Co-editor Kjetil Storesletten handled this manuscript.

Manuscript received 4 September, 2017; final version accepted 13 May, 2020; available online 10 June, 2020.