ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

An International Society for the Advancement of Economic Theory in its Relation to Statistics and Mathematics

http://www.econometricsociety.org/

Econometrica, Vol. 83, No. 2 (March, 2015), 729–769

A POLYNOMIAL OPTIMIZATION APPROACH TO PRINCIPAL–AGENT PROBLEMS

PHILIPP RENNER
Hoover Institution, Stanford University, Stanford, CA 94305, U.S.A.

KARL SCHMEDDERS
University of Zurich and Swiss Finance Institute, 8044 Zurich, Switzerland

Econometrica, Vol. 83, No. 2 (March, 2015), 729–769

A POLYNOMIAL OPTIMIZATION APPROACH TO PRINCIPAL–AGENT PROBLEMS

BY PHILIPP RENNER AND KARL SCHMEDDERS[1]

This paper presents a new method for the analysis of moral hazard principal–agent problems. The new approach avoids the stringent assumptions on the distribution of outcomes made by the classical first-order approach and instead only requires the agent's expected utility to be a rational function of the action. This assumption allows for a reformulation of the agent's utility maximization problem as an equivalent system of equations and inequalities. This reformulation in turn transforms the principal's utility maximization problem into a nonlinear program. Under the additional assumptions that the principal's expected utility is a polynomial and the agent's expected utility is rational in the wage, the final nonlinear program can be solved to global optimality. The paper also shows how to first approximate expected utility functions that are not rational by polynomials, so that the polynomial optimization approach can be applied to compute an approximate solution to nonpolynomial problems. Finally, the paper demonstrates that the polynomial optimization approach extends to principal–agent models with multidimensional action sets.

KEYWORDS: Principal–agent model, moral hazard, polynomial optimization, first-order approach.

1. INTRODUCTION

IN MORAL HAZARD PRINCIPAL–AGENT PROBLEMS, the principal maximizes her expected utility subject to two constraints involving the agent's utility function, a participation constraint and an incentive-compatibility constraint. The participation constraint is rather straightforward; it just imposes a lower bound on the agent's expected utility. The incentive constraint, on the other hand, involves the agent's utility maximization problem. As a consequence, principal–agent problems are a type of bilevel optimization problems,[2] a class of optimization problems that are notoriously difficult. The most popular solution

[1]We thank the co-editor and three anonymous referees for insightful reports. We are indebted to Eleftherios Couzoudis, Johannes Horner, Ken Judd, Diethard Klatte, Felix Kubler, Rida Laraki, George Mailath, Steve Matthews, Walt Pohl, Andy Postlewaite, Gregor Reich, Che-Lin Su, and Rakesh Vohra for helpful discussions on the subject. We thank seminar audiences at the University of Zurich, the 2012 Cowles Summer Conference on Economic Theory, the 2012 ICE Conference at the Becker Friedman Institute, the University of Mannheim, the University of Trier, and the 2014 Initiative for Computational Economics at the Hoover Institution for comments. We are very grateful to Janos Mayer for detailed comments on earlier versions. Philipp Renner and Karl Schmedders gratefully acknowledge financial support from SNSF Grant 148769 and the Swiss Finance Institute, respectively.

[2]The major feature of bilevel optimization problems is that they include two mathematical programs in a single optimization problem. One of the mathematical programs is part of the constraints of the other one. This hierarchical relationship is expressed by calling the two programs the lower-level and the upper-level problem, respectively. In the principal–agent problem, the agent's problem is the lower-level and the principal's problem is the upper-level problem.

 DOI: 10.3982/ECTA11351

approach to principal–agent problems is the first-order approach, which replaces the agent's maximization problem by the corresponding first-order condition and leads to an optimization problem for the principal that is more tractable. Unfortunately, this approach requires very restrictive assumptions on the probability distribution of outcomes, which fail to hold in many economic applications.[3] A more widely applicable solution approach for principal–agent problems is obviously desirable.

In this paper, we present a new method for the analysis of moral hazard principal–agent problems. The new approach avoids the stringent assumptions on the distribution of outcomes made by the classical first-order approach and instead only requires the agent's expected utility to be a rational function of the action. This assumption enables us to apply ideas from polynomial optimization and, similarly to the first-order approach, to transform the principal's utility maximization problem from a bilevel optimization problem to a nonlinear program. For the special but standard case of univariate effort, we obtain an equivalent reformulation. For nonpolynomial problems with one-dimensional effort, we show how to apply the polynomial optimization approach by first approximating nonpolynomial functions with Chebyshev polynomials. Finally, we show how to develop a relaxed reformulation for the model with multidimensional effort and demonstrate that the objective function value of the relaxation converges to the optimal value and that its solution converges to an optimal solution of the original problem.

For principal–agent problems with a one-dimensional effort set for the agent, our assumption that the agent's expected utility function be rational in effort allows us to employ the global optimization approach for rational functions of Jibetean and De Klerk (2006). We transform the agent's expected utility maximization approach into an equivalent semidefinite programming (SDP) problem via a sum of squares representation of the agent's utility function. Semidefinite programs are a special class of convex programming problems which can be solved efficiently both in theory and in practice; see Vandenberghe and Boyd (1996) and Boyd and Vandenberghe (2004). We can further reformulate the constraints of the SDP into a set of inequalities and equations, thereby transforming the principal's bilevel optimization problem into a "normal" nonlinear program. The additional assumptions that all objective functions and constraints are rational, that the action set is an interval, and that the set of wages is compact, imply that the resulting problem is a polynomial optimization problem, which is globally solvable. We can then use the methods implemented in GloptiPoly (see Henrion, Lasserre, and Löfberg (2009)) to find a globally optimal solution to the principal–agent problem. That is, we can obtain a numerical certificate of global optimality.

[3]In economic applications, the first-order approach is then often just assumed to be applicable. In that case, of course, the resulting conclusions may or may not be valid. Needless to say, this custom is rather unsatisfactory.

The first-order approach, a widely used solution method for principal–agent problems, replaces the incentive-compatibility constraint that the agent chooses a utility-maximizing action, by the first-order condition for the agent's utility maximization problem. Mirrlees (1999) (originally circulated in 1975) was the first to show that this approach is invalid in general (even though it had frequently been applied in the literature). Under two conditions on the probability function of outcomes, the monotone likelihood-ratio condition (MLRC) and the convexity of distribution function condition (CDFC), Rogerson (1985) proved the validity of the first-order approach. Mirrlees (1979) had previously surmised that these two assumptions would be sufficient for a valid first-order approach and so these conditions are also known as the Mirrlees–Rogerson conditions. The CDFC is a rather unattractive restriction. Rogerson (1985) pointed out that the CDFC generally does not hold in the economically intuitive case of a stochastic production function with diminishing returns to scale generating the output. In addition, Jewitt (1988) observed that most of the standard textbook probability distributions do not satisfy the CFDC.[4] Jewitt (1988) provided a set of sufficient technical conditions avoiding the CDFC and two sets of conditions for principal–agent models with multiple signals on the agent's effort. Sinclair-Desgagné (1994) introduced a generalization of the CDFC for an extension of the Mirrlees–Rogerson conditions to a first-order approach for multi-signal principal–agent problems. Finally, Conlon (2009) clarified the relationship between the different sets of sufficient conditions and presented multi-signal generalizations of both the Mirrlees–Rogerson and the Jewitt sufficient conditions for the first-order approach. Despite this progress,[5] all of these sufficient sets of conditions are regarded as highly restrictive; see Conlon (2009) and Kadan, Reny, and Swinkels (2011).

Principal–agent models in which the agent's action set is one-dimensional dominate both the literature on the first-order approach as well as the applied and computational literature; see, for example, Araujo and Moreira (2001), Judd and Su (2005), and Armstrong, Larcker, and Su (2010). However, the analysis of linear multi-task principal–agent models in Holmström and Milgrom (1991) demonstrates that multivariate agent problems exhibit some fundamental differences in comparison to the common one-dimensional models. The theoretical literature that allows the set of actions to be multidimensional, for example, Grossman and Hart (1983), Kadan, Reny, and Swinkels (2011), and Kadan and Swinkels (2012), focuses on the existence and properties of

[4] LiCalzi and Spaeter (2003) described two special classes of distributions that satisfy the CDFC.

[5] Araujo and Moreira (2001) introduced a Lagrangian approach different from Mirrlees (1999). Instead of imposing conditions on the outcome distribution, they included more information in the Lagrangian, namely a second-order condition as well as the behavior of the utility function on the boundary in order to account for possible nonconcave objective functions. A number of additional technical assumptions considerably limits the applicability of this approach as well.

equilibria. To the best of our knowledge, the first-order approach has not been extended to models with multiple decision variables affecting the outcome probabilities.

We show how to extend our polynomial optimization approach to principal–agent models in which the agent has more than one decision variable. When we apply the multivariate optimization approach of Jibetean and De Klerk (2006), we encounter a theoretical difficulty. Unlike univariate nonnegative polynomials, multivariate nonnegative polynomials are not necessarily sums of squares of fixed degree. This fact has the consequence that we can no longer provide an exact reformulation of the agent's utility maximization problem but only a relaxation depending on the degree of the involved polynomials. The relaxed problem yields an upper bound on the agent's maximal utility. We then use this relaxation to replace the agent's optimization problems by equations and inequalities including a constraint that requires the upper utility bound not to deviate from the true maximal utility by more than some prespecified tolerance level. We then prove that as the tolerance level converges to zero, the optimal solutions of the sequence of nonlinear programs involving the relaxation converge; and, in fact, the limit points yield optimal solutions to the original principal–agent problem.

Although our main results are of theoretical nature, our paper also contributes to the computational literature on principal–agent problems. Due to the strong assumptions of the first-order approach, the computational literature has shied away from it. Prescott (1999, 2004) approximated the action and compensation sets by finite grids and then allowed for action and compensation lotteries. The resulting optimization problem is linear and thus can be solved with efficient large-scale linear programming algorithms. Judd and Su (2005) avoided the compensation lotteries and only approximated the action set by a finite grid. This approximation results in a mathematical program with equilibrium constraints (MPEC). Contrary to the LP approach, the MPEC approach may face difficulties finding global solutions, since the standard MPEC algorithms only search for locally optimal solutions. Despite this shortcoming, MPEC approaches have recently received a lot of attention in economics; see, for instance, Su and Judd (2012) and Dubé, Fox, and Su (2012). Our polynomial optimization approach does not need lotteries and instead allows us to solve principal–agent problems with continuous action and compensation sets.

The remainder of this paper is organized as follows. In Section 2, we describe the principal–agent model and the classical first-order approach. Section 3 gives a short introduction to polynomial optimization and Section 4 states and proves our main result. In Section 5, we consider two nonpolynomial applications and show how to find approximate solutions with the polynomial optimization approach. We extend the polynomial approach to models with multidimensional action sets in Section 6. Section 7 concludes. Lastly, we also provide some online Supplemental Material (Renner and Schmedders (2015)).

2. THE PRINCIPAL–AGENT MODEL

In this section, we briefly describe the principal–agent model under consideration. Next we review the first-order approach. We complete our initial discussion of principal–agent problems by proving the existence of a global optimal solution.

2.1. *The Principal–Agent Problem*

The agent chooses an action ("effort level") $\mathbf{a}$ from a set $A \subset \mathbb{R}^L$. The signal s (on the outcome, e.g., "output" or "gross profit") received by the principal is an element of the sample space S of a probability space. In many cases, S is chosen to be a discrete subset of the reals. Let $\mu(\bullet|\mathbf{a})$ be a parameterized probability measure on the set of signals S. Then for any $\tilde{S} \subset S$, $\mu(\tilde{S}|\bullet)$ is a function mapping A into $\mathbb{R}$. Of course, $\int_S \mu(ds|\mathbf{a}) = 1$ for all $\mathbf{a} \in A$.

The principal cannot monitor the agent's action but only the signal. Thus, the principal will pay the agent conditional on the observed signal. Let W denote a subset of the set of functions over S. Call an element $w \in W$ a contract. If S is a discrete subset of $\mathbb{R}$ with cardinality N, then the contract ("compensation scheme") between the principal and the agent is a vector $\mathbf{w} = (w_1, w_2, \ldots, w_N) \in W \subset \mathbb{R}^N$. The principal has a Bernoulli utility function over income, $u: I \times S \to \mathbb{R}$, with $I = (\underline{I}, \infty) \subset \mathbb{R}$ for some $\underline{I} \in \mathbb{R} \cup \{-\infty\}$. For example, if the principal receives the signal s and pays the wage $w(s)$, then she receives utility $u(w(s), s)$. The agent has a Bernoulli utility function over income and actions given by $v: J \times A \to \mathbb{R}$, with $J = (\underline{J}, \infty) \subset \mathbb{R}$ for some $\underline{J} \in \mathbb{R} \cup \{-\infty\}$. Both the principal and the agent have von Neumann–Morgenstern utilities. W and μ are now chosen in such a manner that the expected utility functions are well-defined.

The expected utility functions of the principal and agent are

$$U(w, \mathbf{a}) = \int_S u\big(w(s), s\big)\mu(ds|\mathbf{a}) \quad \text{and}$$

$$V(w, \mathbf{a}) = \int_S v\big(w(s), \mathbf{a}\big)\mu(ds|\mathbf{a}),$$

respectively. We are now in the position to state the principal–agent problem:

$$\begin{aligned} (1) \qquad & \max_{w \in W, \mathbf{a} \in A} U(w, \mathbf{a}) \\ & \text{s.t.} \quad \mathbf{a} \in \arg\max_{\mathbf{b} \in A} V(w, \mathbf{b}), \\ & \qquad\; V(w, \mathbf{a}) \geq \underline{V}. \end{aligned}$$

The objective of this optimization problem is to maximize the principal's expected utility. The first constraint,

$$\mathbf{a} \in \arg\max_{\mathbf{b}\in A} V(w, \mathbf{b}), \tag{2}$$

is the incentive-compatibility constraint for the agent; he will only take actions that maximize his own expected utility. We assume implicitly that the agent does not work against the principal; that is, if he is indifferent between several different actions, then he will choose the action most beneficial to the principal. The second constraint is the participation constraint for the agent. He has an outside option and will accept a contract only if he receives at least the expected utility $\underline{V}$ of that outside opportunity.

REMARK—A Comment on the General Problem Assumptions: We briefly provide the reader with some guidance for the general theoretical approach and the subsequent implementation on a computer which we pursue in this paper. Generally speaking, our approach considers problems of the following form:

$$\begin{aligned} &\max_{(w,a)\in X} U(w, a) \\ &\quad \text{s.t.} \quad a \in \arg\max_{b\in A} V(w, b). \end{aligned}$$

We do not impose standard assumptions such as, for example, concave utility functions. Instead, the mathematically essential assumptions for the solution approach are that

- the set A is either a one-dimensional interval, or it is a compact set given by polynomial inequalities; and
- the function V is either a rational function in the variable b or it can be approximated by such on the feasible region.

Moreover, in this paper we restrict attention to deterministic contracts and payoffs. But this assumption is not essential. We forgo any mixing, since its purpose is usually to "convexify" the lower-level problem. This convexification often increases the number of variables significantly. Therefore, in such a case it is more efficient to rely on methods from bilevel optimization.

For the computer implementation of the solution approach to be tractable, we need additional standard assumptions. The function U must have a sufficient order of differentiability; the set X must be finite-dimensional, compact, and defined by a finite number of differentiable functions which satisfy some constraint qualification. Furthermore, a numerical representation on a computer will usually require an appropriate discretization of the problem in order to reduce it to finite dimensions. Therefore, for numerical necessities, it suffices to consider models with a finite set $S = \{s_1, \dots, s_N\}$ and with $W \subset \mathbb{R}^m$ for some m. (There do exist rare exceptions to these last requirements. For example, the normal distribution has the nice property that all its moments exist and are polynomial in the mean and standard deviation. Thus, if we integrate a polynomial function over a normal distribution, the result is a polynomial and no approximation is necessary.)

The principal cannot observe the agent's actions but knows his utility function. Thus, the described principal–agent model exhibits pure moral hazard and no hidden information. The first-order approach for models of this type has been examined by Mirrlees (1999), Rogerson (1985), Jewitt (1988), Sinclair-Desgagné (1994), Alvi (1997), Jewitt, Kadan, and Swinkels (2008), Conlon (2009), and others.

2.2. *The First-Order Approach*

In general, it is very difficult to find a global optimal solution to the principal–agent problem (1). For the model with a one-dimensional action set, $A = [\underline{a}, \bar{a}]$ with $\underline{a} \in \mathbb{R}$ and $\bar{a} \in \mathbb{R} \cup \{\infty\}$, the popular first-order approach replaces the incentive-compatibility constraint (2) by a stationarity condition. If the set A is sufficiently large so that the optimal solution to the agent's expected utility maximization problem has an interior solution, then, for $S = \{s_1, \dots, s_N\}$, the necessary first-order condition is

$$\frac{\partial}{\partial a} V(\mathbf{w}, a) = \sum_{i=1}^{N} \left(\frac{\partial}{\partial a} v(w_i, a) \mu(s_i|a) + v(w_i, a) \frac{\partial}{\partial a} \mu(s_i|a) \right) = 0. \tag{3}$$

For an application of the first-order approach, standard monotonicity, curvature, and differentiability assumptions are imposed. Rogerson (1985) introduced the following assumptions (in addition to some other minor technical conditions):

(1) The function $\mu(s|\bullet) : A \to [0, 1]$ is twice continuously differentiable for all $s \in S$.

(2) The principal's Bernoulli utility function $u : I \times S \to \mathbb{R}$ is strictly increasing, concave, and twice continuously differentiable in its first argument.

(3) The agent's Bernoulli utility function $v : J \times A \to \mathbb{R}$ satisfies $v(w, a) = \psi(w) - a$. The function $\psi : J \to \mathbb{R}$ is strictly increasing, concave, and twice continuously differentiable.

These three assumptions alone are not sufficient for the first-order approach to be valid, since the probabilities $\mu(s_i|a)$ depend on the action a and thus affect the monotonicity and curvature of the expected utility functions. Rogerson (1985) proved the validity of the first-order approach under two additional assumptions on the probability function; see also Mirrlees (1979). We define the function $F_j(a) = \sum_{i=1}^{j} \mu(s_i|a)$. For $\mu(s_i|a) > 0$ for all $a \in A$ and all i, the conditions of Mirrlees (1979) and Rogerson (1985) are as follows:

(MLRC) (monotone likelihood-ratio condition[6]) The measure μ has the property that, for $a_1 \leq a_2$, the ratio $\frac{\mu(s_i|a_1)}{\mu(s_i|a_2)}$ is decreasing in i.

(CDFC) (convexity of the distribution function condition) The function F has the property that $F_i''(a) \geq 0$ for all $i = 1, 2, \dots, N$ and $a \in A$.

According to Conlon (2009), these assumptions are the most popular conditions in economics, even though other sufficient conditions exist; see Jewitt (1988). Sinclair-Desgagné (1994) generalized the conditions of Mirrlees (1979) and Rogerson (1985) for the multi-signal principal–agent problem. Conlon (2009), in turn, presented multi-signal generalizations of both the Mirrlees–Rogerson and the Jewitt sufficient conditions for the first-order approach. Despite this progress, all of these conditions are regarded as highly restrictive; see Conlon (2009) and Kadan, Reny, and Swinkels (2011).

2.3. *Existence of a Global Optimal Solution*

For the sake of completeness, we show the existence of a global optimal solution to the principal–agent problem (1) without assumptions on the differentiability, monotonicity, and curvature of the utility and probability functions. For this purpose, we introduce the following three assumptions.

ASSUMPTION 1—Feasibility: *There exists a contract $w \in W$ such that the agent is willing to participate, that is, $V(w, \mathbf{a}) \geq \underline{V}$ for some $\mathbf{a} \in A$.*

ASSUMPTION 2—Compactness: *Both decision variables are chosen from compact domains.*

(1) *The set A of actions is a nonempty, compact subset of a finite-dimensional Euclidean space, $A \subset \mathbb{R}^L$.*

(2) *The set W of possible contracts is a nonempty, compact metric space.*

ASSUMPTION 3—Continuity: *All functions in the model are well-defined and continuous. In particular, both integrals $U(w, a)$ and $V(w, a)$ exist for any pair (w, a).*

(1) *The principal's expected utility function $U : W \times A \to \mathbb{R}$ is continuous.*

(2) *The agent's expected utility function $V : W \times A \to \mathbb{R}$ is continuous.*

Under the stated assumptions, a global optimal solution to the optimization problem (1) exists.

PROPOSITION 1: *If Assumptions 1–3 hold, then the principal–agent problem (1) has a global optimal solution.*

[6]The MLRC implies a stochastic dominance condition, $F_i'(a) \leq 0$ for all $i = 1, 2, \dots, N$ and $a \in A$.

PROOF: Consider the optimal value function $\Psi : W \to \mathbb{R}$ for the agent defined by $\Psi(w) = \max\{V(\mathbf{w}, \mathbf{a}) \mid a \in A\}$. By Assumptions 2 and 3, the expected utility function V is continuous on the compact domain $W \times A$.[7] Thus, (a special case of) Berge's Maximum Theorem (Berge (1963)) implies that Ψ is continuous on its domain W. Using the function Ψ, we can state the feasible region $\mathcal{F}$ of the principal–agent problem (1),

$$\mathcal{F} = \{(w, \mathbf{a}) \in W \times A \mid V(w, \mathbf{a}) = \Psi(w), V(w, \mathbf{a}) \geq \underline{V}\}.$$

The feasible region $\mathcal{F}$ is nonempty by Assumption 1. As a subset of $W \times A$, it is clearly bounded. Since both V and Ψ are continuous functions and the constraints involve only an equation and a weak inequality, the set $\mathcal{F}$ is also closed. And so the optimization problem (1) requires the maximization of the continuous function U on the nonempty, compact feasible region $\mathcal{F}$. Now the proposition follows from the extreme value theorem of Weierstrass. *Q.E.D.*

3. UNIVARIATE POLYNOMIAL AND RATIONAL OPTIMIZATION

In this section, we briefly describe the underlying mathematics of our reformulation approach. Our approach relies on a classical result from real algebraic geometry to reformulate the agent's problem into a convex optimization problem. Our brief review is based upon the survey by Laurent (2009) and the book by Lasserre (2010). In Section 3.1, we describe so-called sums of squares for multidimensional polynomials. Subsequently, in Sections 3.2–3.4, we restrict attention to univariate polynomials and optimization. We provide a summary of the mathematical background for multidimensional problems in Appendix A.2.1.

3.1. *Polynomials and Sums of Squares Over* $\mathbb{R}^n$

For the study of polynomial optimization, it is necessary to first review a few fundamental concepts from real algebraic geometry.

The expression $\mathbb{R}[x_1, \dots, x_n]$ denotes the ring of polynomials (Greuel and Pfister (2002)) in n variables over the real numbers. Whenever possible, we use the abbreviation $\mathbb{R}[\mathbf{x}]$ with $\mathbf{x} = (x_1, \dots, x_n)^T$. For clarification, we denote the set of nonnegative integers by $\mathbb{N}$. For a vector $\boldsymbol{\alpha} \in \mathbb{N}^n$, we denote the monomial $x_1^{\alpha_1} \cdots x_n^{\alpha_n}$ by $\mathbf{x}^{\boldsymbol{\alpha}}$. The degree of this monomial is $|\boldsymbol{\alpha}| = \sum_{i=1}^n \alpha_i$. A polynomial $p \in \mathbb{R}[\mathbf{x}]$, $p = \sum_{\boldsymbol{\alpha}} a_{\boldsymbol{\alpha}} \mathbf{x}^{\boldsymbol{\alpha}}$ is a sum of terms $a_{\boldsymbol{\alpha}} \mathbf{x}^{\boldsymbol{\alpha}}$ with finitely many nonzero $a_{\boldsymbol{\alpha}} \in \mathbb{R}$. The degree of p is $\deg(p) = \max_{\{\boldsymbol{\alpha} \mid a_{\boldsymbol{\alpha}} \neq 0\}} |\boldsymbol{\alpha}|$.

Let $g_1, \dots, g_m \in \mathbb{R}[\mathbf{x}]$. Then the set

$$K = \{\mathbf{x} \in \mathbb{R}^n \mid g_i(\mathbf{x}) \geq 0, \forall i = 1, \dots, m\}$$

[7] $W \times A$ is compact due to Tychonoff's product theorem, and also metrizable. Therefore, $W \times A$ is a compact metric space and Berge's theorem can be applied.

is called a basic semi-algebraic set.

A central concept of polynomial optimization is the notion of a sum of squares.

DEFINITION 1: A polynomial $\sigma \in \mathbb{R}[\mathbf{x}]$ is called a *sum of squares* if there exist finitely many polynomials $p_1, \dots, p_m \in \mathbb{R}[\mathbf{x}]$ such that $\sigma = \sum_{i=1}^m p_i^2$. The expression $\Sigma[\mathbf{x}] \subset \mathbb{R}[\mathbf{x}]$ denotes the set of sums of squares. And $\Sigma_d[\mathbf{x}] \subset \mathbb{R}[\mathbf{x}]$ denotes the set of sums of squares up to degree d.

A sum of squares σ is always a nonnegative function. The converse, however, is not always true; that is, not every nonnegative polynomial is a sum of squares. Also it is clear that a polynomial can only be a sum of squares if it has even degree. Moreover, the degree of each polynomial p_i in the sum is bounded above by half the degree of σ. To see the link to positive semidefinite matrices, we consider the vector

$$\mathbf{v}_d(\mathbf{x}) = (\mathbf{x}^{\boldsymbol{\alpha}})_{|\boldsymbol{\alpha}| \le d} = (1, x_1, \dots, x_n, x_1^2, x_1 x_2, \dots, x_{n-1}x_n, x_n^2, \dots, x_n^d)^T$$

of all monomials $\mathbf{x}^{\boldsymbol{\alpha}}$ of degree at most d. This vector is of dimension $\binom{n+d}{d}$. There is a strong connection between sums of squares, the vector $\mathbf{v}_d(\mathbf{x})$, and positive semidefinite matrices.

LEMMA 1—Lasserre (2010, Proposition 2.1): *A polynomial $\sigma \in \mathbb{R}[\mathbf{x}]$ of degree $2d$ is a sum of squares if and only if there exists a symmetric positive semidefinite $\binom{n+d}{d} \times \binom{n+d}{d}$ matrix Q such that $\sigma = \mathbf{v}_d(\mathbf{x})^T Q \mathbf{v}_d(\mathbf{x})$, where $\mathbf{v}_d(\mathbf{x})$ is the vector of monomials in $\mathbf{x}$ of degree at most d.*

This lemma completes the description of multidimensional sums of squares. In the remainder of this brief survey, we restrict attention to the special case of univariate polynomials.

3.2. *Sum of Squares Representations Over $\mathbb{R}$*

Recall that a symmetric matrix $M \in \mathbb{R}^{q \times q}$ is called positive semidefinite if and only if $\mathbf{v}^T M \mathbf{v} \ge 0$ for all $\mathbf{v} \in \mathbb{R}^q$. We denote this property of a matrix M by $M \succcurlyeq 0$. The set of all symmetric positive semidefinite $q \times q$ matrices is a closed convex cone.

Now we illustrate the relationship between finding sum of squares representations and semidefinite matrices for the univariate case. For $n = 1$,

$$\mathbf{v}_d(x) = (1, x, x^2, \dots, x^d)^T.$$

We can identify a polynomial $p_i(x) = \sum_{j=0}^d a_{ij} x^j$ with its vector of coefficients $\mathbf{a}_i = (a_{i0}, a_{i1}, \dots, a_{id})^T$ and write $p_i(x) = \mathbf{a}_i^T \mathbf{v}_d(x)$. Next we aggregate m such polynomials in a matrix-vector product

$$\begin{bmatrix} p_1(x) \\ p_2(x) \\ \vdots \\ p_m(x) \end{bmatrix} = \begin{bmatrix} a_{10} & a_{11} & \dots & a_{1d} \\ a_{20} & a_{21} & \dots & a_{2d} \\ \vdots & \vdots & \vdots & \vdots \\ a_{m0} & a_{m1} & \dots & a_{md} \end{bmatrix} \begin{bmatrix} 1 \\ x \\ \vdots \\ x^d \end{bmatrix}.$$

Denoting the $(m \times (d+1))$ coefficient matrix on the right-hand side by V, we can write a sum of squares as

$$\sigma(x) = \sum_{i=1}^{m} p_i^2(x) = \big(V\mathbf{v}_d(x)\big)^T \big(V\mathbf{v}_d(x)\big) = \mathbf{v}_d(x)^T Q \mathbf{v}_d(x)$$

for $Q = V^T V$. By construction, the matrix Q is symmetric, positive semidefinite, and has at most rank m. Note that if we start indexing Q with 0, then Q_{ij} with $i + j = h$ contributes to the term of σ with degree h.

Observe that finding a sum of squares representation for the polynomial $\sigma(x)$ requires finding a symmetric positive semidefinite matrix Q such that the polynomials on the left-hand and right-hand side are identical. But that condition just requires the polynomials to have identical coefficients for all monomials. If σ has degree $2d$, then the coefficient conditions are $2d + 1$ linear equations in the $(d+1)(d+2)/2$ unknown elements of Q. This set of linear equations together with the requirement that Q is symmetric positive semidefinite describe a convex set. And so finding a sum of squares representation of a univariate polynomial σ is equivalent to a convex feasibility problem.

For polynomials in a single variable x, the set of nonnegative polynomials and the set $\Sigma[x]$ of sums of squares are identical.

LEMMA 2—Laurent (2009, Lemma 3.5): *Any nonnegative univariate polynomial is a sum of* (*at most*) *two squares*.

We next consider nonnegative univariate polynomials on closed intervals. For a general treatment, it suffices to examine two cases, $[-1, 1]$ and $[0, \infty)$. The next proposition states that nonnegative polynomials on these intervals can be expressed via two sums of squares and a polynomial that describes the respective interval via a semi-algebraic set. Note that $[-1, 1] = \{x \in \mathbb{R} \mid 1 - x^2 \geq 0\}$ and $[0, \infty) = \{x \in \mathbb{R} \mid x \geq 0\}$.

PROPOSITION 2—Lasserre (2010, Theorems 2.6, 2.7), Laurent (2009, Theorems 3.21, 3.23): *Let* $p \in \mathbb{R}[x]$ *be of degree* d.

(1) $p \geq 0$ *on* $[-1, 1]$ *if and only if*

$$p = \sigma_0 + \sigma_1 \cdot \big(1 - x^2\big), \quad \sigma_0, \sigma_1 \in \Sigma[x]$$

with $\deg(\sigma_0), \deg(\sigma_1 \cdot (1-x^2)) \leq d$ *if* d *is even and* $\deg(\sigma_0), \deg(\sigma_1 \cdot (1-x^2)) \leq d+1$ *if* d *is odd*.

(2) $p \geq 0$ *on* $[0, \infty)$ *if and only if*

$$p = \sigma_0 + \sigma_1 x, \quad \sigma_0, \sigma_1 \in \Sigma[x]$$

with $\deg(\sigma_0), \deg(x\sigma_1) \leq d$.

These results depend critically on the specific description of the intervals via the polynomials $1 - x^2$ and x, respectively. Other descriptions lead to weaker results with representations involving higher degree sum of squares polynomials.

Proposition 2 can also be used to show more general cases. The univariate polynomial $f(x)$ is nonnegative on $K = [a, \infty)$, $K = (-\infty, b]$, and $K = [a, b]$ if and only if

$$\begin{aligned}
&p(x) = f(x+a) \geq 0 \quad \forall x \in [0, \infty), \\
&p(x) = f(b-x) \geq 0 \quad \forall x \in [0, \infty), \\
&p(x) = f\big(\big(x(b-a) + (a+b)\big)/2\big) \geq 0 \quad \forall x \in [-1, 1],
\end{aligned}$$

respectively.

Next we describe the application of the representation results for nonnegative univariate polynomials to polynomial optimization.

3.3. *Polynomial Optimization in* $\mathbb{R}$

For a polynomial $p \in \mathbb{R}[x]$ and a nonempty semi-algebraic set $K \subset \mathbb{R}$, consider the constrained polynomial optimization problem,

$$p_{\min} = \inf_{x \in K} p(x). \tag{4}$$

We can rewrite problem (4) as follows:

$$\begin{aligned}
&\sup_{\rho} \rho \\
&\quad \text{s.t.} \quad p(x) - \rho \geq 0 \quad \forall x \in K.
\end{aligned} \tag{5}$$

For any feasible $\rho \in \mathbb{R}$, the following inequality holds:

$$\rho \leq p_{\min}. \tag{6}$$

Note that the constraints of the rewritten problem state that the polynomial $p - \rho$ must be nonnegative on the set K. Now consider the domain $K = [-1, 1] = \{x \mid 1 - x^2 \geq 0\}$. In this case, applying part (1) of Proposition 2 enables us to rewrite the infinitely many constraints of problem (5). With the polynomial g defined by $g(x) = 1 - x^2$, we obtain the following optimization problem:

$$\begin{aligned}
(7) \qquad & \sup_{\rho,\sigma_0,\sigma_1} \ \rho \\
& \text{s.t.} \quad p - \rho = \sigma_0 + \sigma_1 g, \\
& \qquad\quad \sigma_0, \sigma_0 \in \Sigma[x].
\end{aligned}$$

Note that the equality constraint here signifies equality as polynomials. Lemma 1 enables us to rewrite the optimization problem once more by replacing the unknown sums of squares σ_0 and σ_1 by positive semidefinite matrices. We define[8] the number $d_p = \lceil \frac{\deg(p)}{2} \rceil$ for a polynomial $p \in \mathbb{R}[x]$. According to Proposition 2 the number d_p is an upper bound for the degrees of σ_0 and σ_1. And so we can rewrite the optimization problem:

$$\begin{aligned}
(8) \qquad & \sup_{\rho,Q^{(0)},Q^{(1)}} \ \rho \\
& \text{s.t.} \quad p - \rho = v_{d_p}^T Q^{(0)} v_{d_p} + g v_{d_p-1}^T Q^{(1)} v_{d_p-1}, \\
& \qquad\quad Q^{(0)}, Q^{(1)} \succcurlyeq 0, \\
& \qquad\quad Q^{(0)} \in \mathbb{R}^{(d_p+1)\times(d_p+1)}, \quad Q^{(1)} \in \mathbb{R}^{d_p \times d_p}, \\
& \qquad\quad v_{d_p} = \left(1, x, \dots, x^{d_p}\right)^T, \quad v_{d_p-1} = \left(1, x, \dots, x^{d_p-1}\right)^T.
\end{aligned}$$

Note that the first functional constraint holds if and only if all coefficients (of identical monomials on the left- and right-hand side) are identical. Thus this functional constraint reduces to a set of linear constraints which only involve the coefficients of the terms. Let $p = \sum_{l=0}^{\deg(p)} c_l x^l$ and write $Q_{ij}^{(0)}$, $i, j = 0, 1, \dots, d_p$, for the (i, j)th entry of the matrix $Q^{(0)}$ (similarly for $Q^{(1)}$). Then we can rewrite the first constraint of problem (8),

$$\begin{aligned}
(9) \qquad & c_0 - \rho = Q_{0,0}^{(0)} + Q_{0,0}^{(1)}, \\
& c_l = \sum_{i+j=l} Q_{ij}^{(0)} + \sum_{i+j=l} Q_{ij}^{(1)} - \sum_{i+j=l-2} Q_{ij}^{(1)}, \quad l = 1, \dots, d.
\end{aligned}$$

This set of constraints is just a set of linear equations in the unknowns ρ and $Q_{ij}^{(m)}$. In particular, we observe that the final optimization problem is a semidefinite optimization problem (SDP). Note that the positive semidefinite constraint for the matrices $Q^{(0)}$ and $Q^{(1)}$ can be interpreted as polynomial inequality constraints. This fact follows from Proposition A.1 in the Supplemental Material.

[8] Recall the notation $\lceil x \rceil$ for the smallest integer not less than x.

The following proposition summarizes the relationship between the original problem and the reformulation.

PROPOSITION 3—Lasserre (2010, Theorem 5.8): *If* $p(x) = \sum_i c_i x^i$ *and* $K = \{x \in \mathbb{R} \mid 1 - x^2 \geq 0\} = [-1, 1]$, *then problem* (8) *is equivalent to* $\inf_{x\in[-1,1]} p(x)$ *and both problems have an optimal solution*.

The optimal solutions satisfy $\rho = p_{\min}$. In sum, the constrained optimization problem of minimizing a univariate polynomial on an interval of $\mathbb{R}$ reduces to an SDP, a convex optimization problem. Note that the optimization is over a compact domain and so we can replace the supremum by a maximum.

3.4. *Rational Objective Function*

Jibetean and De Klerk (2006) proved an analogous result for the case of rational objective functions. Let $p(x), q(x)$ be two polynomials defined on a subset $K \subset \mathbb{R}$. Consider the following optimization problem:

$$p_{\min} = \inf_{x\in K, q(x)\neq 0} \frac{p(x)}{q(x)}. \tag{10}$$

We can rewrite this problem in polynomial form.

PROPOSITION 4—Jibetean and De Klerk (2006, Theorem 2): *If p and q have no common factor and K is an open connected set or a* (*partial*) *closure of such a set*, *then*

(1) *if q changes sign on* K, *then* $p_{\min} = -\infty$,

(2) *if q is nonnegative on* K, *problem* (10) *is equivalent to*

$$p_{\min} = \sup\{\rho \mid p(x) - \rho q(x) \geq 0, \forall \mathbf{x} \in K\}.$$

Let $p, q \in \mathbb{R}[x]$ and set $d = \max(d_p, d_q)$. For $K = [-1, 1]$ and $g(x) = 1 - x^2$, we can again use Proposition 2 and reformulate problem (10),

$$\begin{aligned} &\sup_{\rho,\sigma_0,\sigma_1} \ \rho \\ &\quad \text{s.t.} \quad p - \rho q = \sigma_0 + g\sigma_1, \\ &\qquad\qquad \sigma_0 \in \Sigma_{2d}, \quad \sigma_1 \in \Sigma_{2(d-1)}. \end{aligned} \tag{11}$$

And so we can solve the constrained optimization problem (10) also as an SDP.

4. THE POLYNOMIAL OPTIMIZATION APPROACH FOR $A \subset \mathbb{R}$

In this section, we state our main result, Theorem 1, and illustrate it by a numerical example. Subsequently, we prove the theorem, and finally, we provide a discussion of the key assumptions and the resulting limitations of the polynomial optimization approach.

Throughout this section, we assume that the set S is finite. This assumption is not particularly restrictive, since, for computational purposes, it is always necessary to approximate the integrals in the expectation operators over continuous domains by finite sums and also to limit the compensation scheme to a finite-dimensional set. In the example in Section 5.2, we demonstrate how to approximate a continuous problem.

4.1. *The Main Theorem*

We introduce the following assumption on the agent's expected utility function.

ASSUMPTION 4—Rational Expected Utility Function: *The parameterized probability functions $\mu(s|\bullet): A \to [0,1]$ and the agent's Bernoulli utility function $v: J \times A \to \mathbb{R}$ are such that the agent's expected utility function is a rational function of the form*[9]

$$V(\mathbf{w}, a) = \sum_{j=1}^{N} v(w_j, a)\mu(s_j|a) = -\frac{\sum_{i=0}^{d} c_i(\mathbf{w})a^i}{\sum_{i=0}^{d} f_i(\mathbf{w})a^i}$$

for functions $c_i, f_i: W \to \mathbb{R}$ with $\sum_{i=0}^{d} f_i(\mathbf{w})a^i > 0$ for all $(\mathbf{w}, a) \in W \times A$.[10] *Moreover, the two polynomials in the variable a, $\sum_{i=0}^{d} c_i(\mathbf{w})a^i$ and $\sum_{i=0}^{d} f_i(\mathbf{w})a^i$, have no common factors and $d \in \mathbb{N}$ is maximal such that $c_d(\mathbf{w}) \neq 0$ or $f_d(\mathbf{w}) \neq 0$.*

In light of Proposition 2, it suffices to consider the set of actions $A = [-1, 1] = \{a \in \mathbb{R} \mid 1 - a^2 \geq 0\}$. (The unbounded case can be handled similarly.) We define the number $D = \lceil \frac{d}{2} \rceil$. The following theorem[11] provides us with an optimization problem that is equivalent to the principal–agent problem (1).

[9]The minus sign in the rational expression simplifies the application of mathematical methods from the literature.

[10]The positivity condition for the denominator is necessary, since a change in sign would lead to division by zero.

[11]Note that the row and column indexing of the two matrices in the theorem starts at 0.

THEOREM 1: *Let $A = [-1, 1]$ and suppose Assumption 4 holds. Then $(\mathbf{w}^*, a^*)$ solves the principal–agent problem (1) if and only if there exist $\rho^* \in \mathbb{R}$ as well as matrices $Q^{(0)*} \in \mathbb{R}^{(D+1)\times(D+1)}$ and $Q^{(1)*} \in \mathbb{R}^{D\times D}$ such that $(\mathbf{w}^*, a^*, \rho^*, Q^{(0)*}, Q^{(1)*})$ solves the following optimization problem*:

$$\max_{\mathbf{w},a,\rho,Q^{(0)},Q^{(1)}} U(\mathbf{w}, a) \tag{12}$$

$$\text{s.t.} \quad c_0(\mathbf{w}) - \rho f_0(\mathbf{w}) = Q^{(0)}_{0,0} + Q^{(1)}_{0,0}, \tag{12.a}$$

$$c_l(\mathbf{w}) - \rho f_l(\mathbf{w}) = \sum_{i+j=l} Q^{(0)}_{ij} + \sum_{i+j=l} Q^{(1)}_{ij} - \sum_{i+j=l-2} Q^{(1)}_{ij}, \tag{12.b}$$

$$l = 1, \dots, d,$$

$$Q^{(0)}, Q^{(1)} \succcurlyeq 0, \tag{12.c}$$

$$\rho\left(\sum_{i=0}^{d} f_i(\mathbf{w}) a^i\right) = \sum_{i=0}^{d} c_i(\mathbf{w}) a^i, \tag{12.d}$$

$$\sum_{i=0}^{d} c_i(\mathbf{w}) a^i \leq -\underline{V}\left(\sum_{i=0}^{d} f_i(\mathbf{w}) a^i\right), \tag{12.e}$$

$$-a^2 + 1 \geq 0, \tag{12.f}$$

$$\mathbf{w} \in W. \tag{12.g}$$

The new optimization problem (12) has the same objective function as the original principal–agent problem (1). Unlike the original problem, the new problem (12) is not a bilevel optimization problem. Instead, the constraint involving the agent's expected utility maximization problem has been replaced by inequalities and equations. Problem (12) has the additional decision variables $\rho \in \mathbb{R}$, $Q^{(0)} \in \mathbb{R}^{(D+1)\times(D+1)}$, and $Q^{(1)} \in \mathbb{R}^{D\times D}$. The optimal value ρ^* of the variable ρ in problem (12) will be $-V(\mathbf{w}^*, a^*)$, the negative of the agent's maximal expected utility. Constraints (12.a)–(12.c) use a sum of squares representation of nonnegative polynomials to ensure that, for a contract $\mathbf{w}$ chosen by the principal, $-V(\mathbf{w}, a) \geq \rho$ for all $a \in A$. That is, $-\rho$ is an upper bound on all possible utility levels for the agent. Note that equations (12.a) and (12.b) are linear in ρ and the elements of the matrices $Q^{(0)} \in \mathbb{R}^{(D+1)\times(D+1)}$ and $Q^{(1)} \in \mathbb{R}^{D\times D}$. Constraint (12.c) requires that these two matrices are symmetric positive semidefinite. (Proposition A.1 in the Supplemental Material summarizes properties of positive semidefinite matrices which show that constraint (12.c) can be written as a set of polynomial inequalities.) Next, constraint (12.d) ensures that the variable $-\rho$ is actually equal to the agent's utility for effort a and contract $\mathbf{w}$. Therefore, this constraint together with the constraints (12.a)–(12.c) forces any value of a satisfying the equation to be the agent's optimal effort choice as well as the value of ρ to be the corresponding maximal expected utility value. Put differently, for a given contract **w**, the first four constraints ensure an optimal effort choice by the agent. The last three constraints are straightforward. Constraint (12.e) is the transformed participation constraint for the agent's rational expected utility function. Constraint (12.f) is a polynomial representation of the feasible action set and constraint (12.g) is just the constraint on the compensation scheme from the original principal–agent problem (1).

4.2. *An Illustrative Example*

Before we prove the theorem, it is helpful to illustrate it by a simple example.

EXAMPLE 1: Let $A = [0, 1]$ and $\mathcal{W} = \mathbb{R}_+$. There are $N = 3$ possible outcomes $s_1 < s_2 < s_3$ which occur with the probabilities

$$\mu(s_1|a) = \binom{2}{0} a^0(1-a)^2, \quad \mu(s_2|a) = \binom{2}{1} a(1-a),$$
$$\mu(s_3|a) = \binom{2}{2} a^2(1-a)^0.$$

The principal is risk-neutral with Bernoulli utility $u(w, s) = s - w$. The agent is risk-averse and has utility

$$v(w, a) = \frac{w^{1-\eta} - 1}{1-\eta} - \kappa a^2,$$

where $\eta \neq 1$, $\eta \geq 0$, and $\kappa > 0$. The agent's expected utility is

$$V(w_1, w_2, w_3, a) = (1-a)^2 \frac{w_1^{1-\eta} - 1}{1-\eta} + 2(1-a)a \frac{w_2^{1-\eta} - 1}{1-\eta} + a^2 \frac{w_3^{1-\eta} - 1}{1-\eta} - \kappa a^2.$$

The second-order partial derivative of V with respect to a,

$$\frac{\partial^2 V}{\partial a^2} = \frac{2w_1^{1-\eta}}{1-\eta} - \frac{4w_2^{1-\eta}}{1-\eta} + \frac{2w_3^{1-\eta}}{1-\eta} - 2\kappa,$$

changes sign on $W \times A$. Thus, this function is not concave and so the classical first-order approach does not apply. We apply Theorem 1 to solve this principal–agent problem. For simplicity, we consider a specific problem with $\eta = \frac{1}{2}$, $\underline{V} = v(1, 0) = 0$, $\kappa = 2$, and $(s_1, s_2, s_3) = (0, 2, 4)$.

First we transform the set of actions $A = [0, 1]$ into the interval $A = [-1, 1]$ via the variable transformation $a \mapsto \frac{a+1}{2}$. The resulting expected utility functions are

$$\begin{aligned}
U(\mathbf{w}, a) &= 2 + 2a - \frac{w_1}{4} + \frac{aw_1}{2} - \frac{a^2 w_1}{4} - \frac{w_2}{2} \\
&\quad + \frac{a^2 w_2}{2} - \frac{w_3}{4} - \frac{aw_3}{2} - \frac{a^2 w_3}{4}, \\
V(\mathbf{w}, a) &= -\frac{5}{2} + \frac{\sqrt{w_1}}{2} + \sqrt{w_2} + \frac{\sqrt{w_3}}{2} - a - a\sqrt{w_1} \\
&\quad + a\sqrt{w_3} - a^2\sqrt{w_2} + \frac{a^2\sqrt{w_1}}{2} - \frac{a^2}{2} + \frac{a^2\sqrt{w_3}}{2}.
\end{aligned}$$

We observe that $V(\mathbf{w}, a)$ is a quadratic polynomial in a and so, referring to Assumption 4, $d = 2$ and $D = 1$. The representation of $-V(\mathbf{w}, a)$ according to that assumption has the nonzero coefficients $f_0(\mathbf{w}) = 1$ and $c_0(\mathbf{w}) = \frac{5}{2} - \frac{\sqrt{w_1}}{2} - \sqrt{w_2} - \frac{\sqrt{w_3}}{2}$, $c_1(\mathbf{w}) = 1 + \sqrt{w_1} - \sqrt{w_3}$, and $c_2(\mathbf{w}) = \frac{1}{2} - \frac{\sqrt{w_1}}{2} + \sqrt{w_2} - \frac{\sqrt{w_3}}{2}$. According to Theorem 1, the matrix $Q^{(0)}$ is a 2×2 matrix and $Q^{(1)}$ is just a single number. With

$$Q^{(0)} = \begin{pmatrix} n_{00} & n_{01} \\ n_{01} & n_{11} \end{pmatrix} \quad \text{and} \quad Q^{(1)} = m_{00},$$

we can rewrite the principal–agent problem following the theorem:

$$\begin{aligned}
\max_{w_1, w_2, w_3, a, \rho, n_{00}, n_{01}, n_{11}, m} \quad & U(w_1, w_2, w_3, a) \\
\text{s.t.} \quad & \frac{5}{2} - \frac{\sqrt{w_1}}{2} - \sqrt{w_2} - \frac{\sqrt{w_3}}{2} - \rho = n_{00} + m_{00}, \\
& 1 + \sqrt{w_1} - \sqrt{w_3} = 2n_{01}, \\
& \frac{1}{2} - \frac{\sqrt{w_1}}{2} + \sqrt{w_2} - \frac{\sqrt{w_3}}{2} = n_{11} - m_{00}, \\
& \rho = -V(w_1, w_2, w_3, a), \\
& n_{00} \geq 0, \quad n_{11} \geq 0, \quad n_{00}n_{11} - n_{01}^2 \geq 0, \quad m_{00} \geq 0, \\
& V(w_1, w_2, w_3, a) \geq 0, \\
& -a^2 + 1 \geq 0, \\
& w_1, w_2, w_3 \geq 0.
\end{aligned}$$

TABLE I

NUMERICAL SOLUTIONS TO THE PRINCIPAL–AGENT PROBLEM AS A FUNCTION OF η

η	$U(w_1^*, w_2^*, w_3^*, a^*)$	a^*	w_1^*	w_2^*	w_3^*	U_{FB}	a_{FB}	w_{FB}
0	1	1	$[0, 1)$	$[0, 1]$	3	1	1	3
$\frac{1}{4}$	0.6760	0.8260	0.2777	1.177	3.344	0.7471	0.7993	2.450
$\frac{1}{3}$	0.5723	0.7637	0.2879	1.273	3.441	0.6850	0.7541	2.332
$\frac{1}{2}$	0.3844	0.6446	0.3417	1.511	3.511	0.5814	0.6823	2.148
$\frac{4}{5}$	0.1292	0.4881	0.5314	1.798	3.296	0.4410	0.5918	1.926
2	−0.3444	0.2413	0.8749	1.817	2.416	0.1349	0.4196	1.544
4	−0.6102	0.1277	0.9657	1.597	1.866	−0.09165	0.3117	1.338

This optimization problem is not yet polynomial. We replace $\sqrt{w_i}$ by a new variable $\hat{w}_i$. This change of variable results in a polynomial problem. Then we can solve this nonlinear optimization problem with GloptiPoly (see Henrion, Lasserre, and Löfberg (2009)) and obtain the globally optimal contract $\mathbf{w}^* = (0.3417, 1.511, 3.511)$ and the resulting optimal effort $a^* = 0.6446$. Table I reports solutions for different levels of the agent's risk aversion η. For completion, the table also reports the corresponding first-best solutions[12] indexed by FB.

For $\eta = 0$, when the agent is risk-neutral, a continuum of contracts exists. However, the intervals of values for w_1 and w_2 are economically irrelevant since, for $w_3 = 1$, the optimal effort of $a^* = 1$ results in zero probability of outcomes 1 and 2 and the first-best solution.

The purpose of this example is to illustrate the statement of Theorem 1. As such, the example was deliberately constructed to be simple. And even though the principal–agent problem in the example does not satisfy the sufficient conditions of the first-order approach, that method does deliver the same solution as the polynomial optimization approach. In Section 5, we analyze much more difficult problems from the economic literature.

4.3. *Proofing of Theorem 1*

PROOF OF THEOREM 1: Comparing the original principal–agent problem (1) and the new problem (12), we observe that the upper-level problem has not been altered. In particular, we still maximize the same function U. Thus to show that these problems are indeed equivalent, it suffices to see that any feasible point for (12) corresponds to a feasible point for (1) and vice versa.

[12] Omitting the incentive compatibility constraint and maximizing the principal's expected utility only subject to the participation constraint leads to the first-best solution.

Let $(\hat{\mathbf{w}}, \hat{a}, \hat{\rho}, \hat{Q}^{(0)}, \hat{Q}^{(1)})$ be a feasible point for problem (12). Then by inequality (6) in Section 3.3, we have that

$$\hat{\rho} \leq \min_{a\in[-1,1]} -V(\hat{\mathbf{w}}, a) = -\max_{a\in[-1,1]} V(\hat{\mathbf{w}}, a) \leq -V(\hat{\mathbf{w}}, a)$$

for any $a \in [-1, 1]$. Thus by the equality condition $-V(\hat{\mathbf{w}}, \hat{a}) = \hat{\rho}$, we have that $V(\hat{\mathbf{w}}, \hat{a}) = \max_{a\in[-1,1]} V(\hat{\mathbf{w}}, a)$. Therefore, $\hat{a} \in \arg\max_{a\in[-1,1]} V(\hat{\mathbf{w}}, a)$ and $V(\hat{\mathbf{w}}, \hat{a}) \geq \underline{V}$. Hence, $(\hat{\mathbf{w}}, \hat{a})$ is a feasible point for (1).

Now let $(\hat{\mathbf{w}}, \hat{a})$ be a feasible point for (1). So $\hat{a} \in \arg\max_{a\in[-1,1]} V(\hat{\mathbf{w}}, a)$. By Proposition 4 from Section 3.4, there exist $\hat{Q}^{(0)}, \hat{Q}^{(1)} \succcurlyeq 0$ and a maximal $\hat{\rho}$ such that the following system of equations is satisfied:

$$\begin{aligned} c_0(\hat{\mathbf{w}}) - \hat{\rho} f_0(\hat{\mathbf{w}}) &= \hat{Q}^{(0)}_{0,0} + \hat{Q}^{(1)}_{0,0}, \\ c_l(\hat{\mathbf{w}}) - \hat{\rho} f_l(\hat{\mathbf{w}}) &= \sum_{i+j=l} \hat{Q}^{(0)}_{ij} + \sum_{i+j=l} \hat{Q}^{(1)}_{ij} - \sum_{i+j=l-2} \hat{Q}^{(1)}_{ij}, \quad l = 1, \ldots, d. \end{aligned}$$

Then $\hat{\rho} = \min_{a\in[-1,1]} -V(\hat{\mathbf{w}}, a) = -V(\hat{\mathbf{w}}, \hat{a})$ and therefore $(\hat{\mathbf{w}}, \hat{a}, \hat{\rho}, \hat{Q}^{(0)}, \hat{Q}^{(1)})$ is feasible for (12). *Q.E.D.*

The proof establishes that the feasible region of the original principal–agent problem (1) is a projection of the feasible region of the optimization problem (12). The first four constraints of problem (12) capture the agent's expected utility maximization problem. The constraints (12.a)–(12.d) force any value of a in a feasible solution to be the agent's optimal effort choice as well as the value of ρ to be the corresponding maximal expected utility value. Put differently, for a given contract $\mathbf{w}$, the first four constraints ensure an optimal effort choice by the agent.

With some additional assumptions, we can solve the optimization problem (12) to global optimality.

COROLLARY 1: *Suppose Assumption 4 holds and that the functions $c_i, f_i : W \to \mathbb{R}$ (in Assumption 4) are polynomials in $\mathbf{w} \in W$. Moreover, assume that U is a polynomial, $A = [-1, 1]$, and W is a basic semi-algebraic set. Then (12) is a polynomial optimization problem over a basic semi-algebraic set.*

PROOF: The only problematic constraints are the positive semidefiniteness constraints for the matrix. However, the positive semidefiniteness condition on the $Q^{(i)}$ is equivalent to the condition that the principal minors, that are themselves polynomials, are nonnegative. Thus the set of constraints defines a semi-algebraic set. *Q.E.D.*

If the conditions of the corollary are satisfied, we can use the methods employed in GloptiPoly (see Henrion, Lasserre, and Löfberg (2009)) to find a globally optimal solution to the principal–agent problem. That is, we can obtain a numerical certificate of global optimality. We use such an approach in Example 1 to ensure global optimality.

4.4. *Discussion of the Polynomial Approach's Assumptions and Limitations*

Theorem 1 rests on two key assumptions, namely, that the agent's choice set is a compact interval and his expected utility function is rational in effort. The review of the mathematical background and the derivation of the theorem show that we can easily dispense with the compactness assumption and replace it by an unbounded interval such as $[0, \infty)$. While the second assumption limits the direct applicability of the theorem, it does include the special case of agents' utility functions that are separable in wage and effort and feature a linear cost of effort (together with a rational probability distribution of outcomes).

Corollary 1 imposes additional assumptions on the utility functions and the set of wages; the principal's expected utility is polynomial and the agent's expected utility is rational in wages; the set of wages is a basic semi-algebraic set. The assumption on the set of wages appears to be innocuous. The assumptions on the utility functions rule out many standard utility functions such as exponential or logarithmic utility functions. Moreover, the principal's utility cannot exhibit constant risk aversion. Although the assumption on the principal's utility function is rather strong, it includes the special case of a risk-neutral principal and a polynomial probability distribution. If the assumptions of Corollary 1 do not hold, we can still attempt to solve the final NLP with standard nonlinear optimization routines.

Moreover, we show in Section 5 that we can apply the polynomial optimization approach even to nonpolynomial problems. Before we can do so, we first approximate the involved nonpolynomial functions with polynomial ones. In Appendix B, we briefly describe an excellent approximation result for Chebyshev interpolation. Therefore, an application of the polynomial optimization approach to the polynomial approximation of a principal–agent model has the potential to deliver an excellent approximation of the original optimal solution. As a result, even the assumptions on the expected utility functions in both the theorem and its corollary are not as limiting as they may appear at first.

The most serious limitation of our polynomial optimization approach is that it is not suited for a subsequent traditional theoretical analysis of the principal–agent model. A central topic of the economic literature on moral hazard problems has been the study of the nature of the optimal contract and its comparative statics properties. Studies invoking the first-order approach rely on the KKT conditions for the relaxed principal's problem to perform such an analysis. For example, Rogerson (1985) considered the case of a separable utility function with linear cost of effort; using our notation, we can write (slightly abusing notation) $v(w_i, a) = v(w_i) + a$. Rogerson (1985) stated the KKT conditions for the relaxed principal's problem, part of which are the equations

$$\frac{u'(s_i - w_i)}{v'(w_i)} = \lambda + \delta \frac{\mu'(s_i|a)}{\mu(s_i|a)} \tag{13}$$

for $i = 1, 2, \ldots, N$ with Lagrange multipliers λ and δ. Rogerson (1985) then used these equations not only to prove the validity of the first-order approach but also to show that the optimal wage contract is increasing in the output. An analogous approach to the analysis of the optimal contract has been used in many studies; see, for example, Holmström (1979), Jewitt (1988), and Jewitt, Kadan, and Swinkels (2008). The KKT conditions for the relaxed principal's problem are rather simple since that problem has only two constraints, the participation constraint and the first-order condition for the agent's problem. The optimization problem (12) stated in Theorem 1, however, has many more constraints. In addition, the constraints characterizing the agent's optimal effort choice are not intuitive. As a result, we cannot follow the traditional approach for theoretically deriving additional properties of the principal's problem by simply using the reformulated optimization problem (12).

Since we cannot follow the traditional theoretical route, we would instead have to rely on numerical solutions of many instances of problem (12) for a further analysis of the properties of the optimal contract. While at first such a numerical analysis may look rather unattractive compared to the theoretical analysis based on the first-order approach, it also offers some advantages. The first-order approach requires very strong assumptions and so applies only to a small set of principal–agent problems. A numerical analysis based on our polynomial optimization approach can examine many other problems that fall outside the classical first-order approach.

5. EXAMPLES FROM THE ECONOMIC LITERATURE

In this section, we demonstrate the versatility of the proposed polynomial optimization approach to principal–agent models. For this purpose, we first revisit a counterexample to the classical first-order approach presented by Mirrlees (1999) (originally circulated in 1975) and show that we can solve this example with the polynomial approach. Subsequently, we investigate an economic application of principal–agent models, namely, the executive compensation problem in Armstrong, Larcker, and Su (2010). Both examples have in common that they involve exponential functions and, therefore, do not satisfy the assumptions of Theorem 1. We first approximate the involved functions by Chebyshev polynomials and then apply the polynomial optimization approach to the resulting polynomial problem. Appendix B provides a brief introduction to Chebyshev interpolation.

5.1. *A Classical Example*

Mirrlees (1999, Example 1) considered the following optimization problem[13]:

$$\max_{z,b} -(b-2)^2-(z-1)^2 \quad \text{s.t.} \quad z \in \arg\max_x V(b,x) = be^{-(x+1)^2}+e^{-(x-1)^2} \tag{14}$$

and showed that the first-order approach fails to solve this problem. Not a single of the solutions delivered by the first-order approach matches the actual optimal solution which Mirrlees (1999) reported as $b=1$ and $z=0.957$.

Observe that for $b=0$, the unique optimal solution to the agent's problem is $z=1$ and so the objective function value for the principal is -4. For $b<0$, the principal's value will be less than -4. Thus, it suffices to consider the case $b \geq 0$. Furthermore, observe that the function V is a weighted sum of two functions describing bell curves, the first with nonnegative weight b centered at -1 and the second with weight 1 centered at 1. As a result, V is strictly monotonically increasing for $x<-1$ and strictly decreasing for $x>1$ and so all globally optimal solutions to the lower-level problem lie in $[-1,1]$ for all $b \geq 0$. Moreover,

$$V(b,x)-V(b,-x) = (b-1)\big(e^{-(x+1)^2}-e^{-(x-1)^2}\big),$$

and so for $b<1$, the value $V(b,x)$ is greater for positive values of x than for negative values; conversely, for $b>1$, the value $V(b,x)$ is greater for negative values of x than for positive values. Thus, the optimal solution for b lies in $[0,1]$. In sum, without loss of generality, we can impose the constraints $b \in [0,1]$ and $z \in [-1,1]$ and consider the following equivalent problem:

$$\begin{aligned} &\max_{z,b} -(b-2)^2-(z-1)^2 \\ &\text{s.t.} \quad z \in \arg\max_{x\in[-1,1]} V(b,x) = be^{-(x+1)^2}+e^{-(x-1)^2}, \\ &\qquad b \in [0,1]. \end{aligned} \tag{15}$$

Next, we approximate the function V by Chebyshev polynomials; see Appendix B. As a starting point, consider the Chebyshev polynomials of up to degree 2,

$$T_0(x)=1, \quad T_1(x)=x, \quad T_2(x)=-1+2x^2,$$

[13]To avoid a collision of notation, we denote the variable a in Mirrlees (1999) by b. Note that this example lacks a participation constraint.

with the interpolation nodes $\frac{\sqrt{3}}{2}, 0, -\frac{\sqrt{3}}{2}$. At these interpolation nodes, the agent's utility function attains the values

$$V\left(b, \frac{\sqrt{3}}{2}\right) = e^{-7/4-\sqrt{3}}\left(b + e^{2\sqrt{3}}\right),$$

$$V(b, 0) = \frac{1+b}{e},$$

$$V\left(b, -\frac{\sqrt{3}}{2}\right) = e^{-7/4-\sqrt{3}}\left(1 + be^{2\sqrt{3}}\right).$$

Computing the Chebyshev interpolant yields

$$\begin{aligned} p_2(b, x) = {} & \frac{\left(e^{7/2} + e^{11/4-\sqrt{3}} + e^{11/4+\sqrt{3}}\right)(b+1)}{3e^{9/2}} T_0(x) \\ & - \frac{e^{-7/4-\sqrt{3}}\left(e^{2\sqrt{3}} - 1\right)(b-1)}{\sqrt{3}} T_1(x) \\ & + \frac{1}{3} e^{-5/2-\sqrt{3}}\left(e^{3/4} - 2e^{3/2+\sqrt{3}} + e^{3/4+2\sqrt{3}}\right)(b+1) T_2(x). \end{aligned}$$

The maximal absolute approximation error satisfies

$$\max_{b \in [0,1], x \in [-1,1]} \left|V(b, x) - p_2(b, x)\right| \leq 0.102007.$$

The plot in Figure 1 shows the approximation error $|V(b, x) - p_2(b, x)|$ for $b = 0$. The global maximum of the approximation error is attained at the point $(b, z) = (0, 1)$. Replacing the agent's objective function V by its Chebyshev approximation, we obtain the following optimization problem:

$$\begin{aligned} (16) \qquad & \max_{z,b} -(b-2)^2 - (z-1)^2 \\ & \text{s.t.} \quad z \in \arg\max_{x \in [-1,1]} p_2(b, x), \\ & \qquad\quad b \in [0, 1], \\ & \qquad\quad z \in [-1, 1]. \end{aligned}$$

Now we apply the reformulation of Theorem 1 to this problem.[14] Since p_2 has degree $d = 2$, we have $D = 1$ and so we need to define

$$Q^{(0)} = \begin{pmatrix} q_{11} & q_{12} \\ q_{12} & q_{22} \end{pmatrix} \quad \text{and} \quad Q^{(1)} = u.$$

[14]Note that in this example, the variables z and b play the roles of a and $\mathbf{w}$, respectively, in Theorem 1.

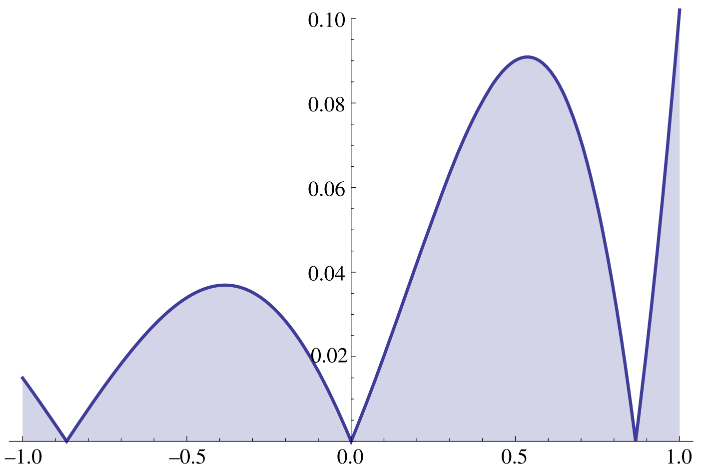

FIGURE 1.—The absolute error of the degree-2 Chebyshev approximation, $|V(0,x) - p_2(0,x)|$.

Thus, we have to compare the coefficients in the following expression:

$$-p_2(b,x) - \rho = q_{11} + 2xq_{12} + x^2 q_{22} + u - ux^2.$$

We obtain the following equalities:

$$\frac{b+1}{e} + \rho + q_{11} + u = 0,$$

$$-\frac{e^{-7/4-\sqrt{3}}\left(e^{2\sqrt{3}} - 1\right)(b-1)}{\sqrt{3}} + 2q_{12} = 0,$$

$$\frac{2}{3}e^{-5/2-\sqrt{3}}\left(e^{3/4} - 2e^{3/2+\sqrt{3}} + e^{3/4+2\sqrt{3}}\right)(b+1) + q_{22} - u = 0.$$

The positive semidefiniteness conditions $Q^{(0)}, Q^{(1)} \succcurlyeq 0$ yield the inequalities

$$q_{11}q_{22} - q_{12}^2 \geq 0,$$

$$q_{11}, q_{22}, u \geq 0.$$

The final constraints are

$$-x^2 + 1 \geq 0, \quad b \geq 0, \quad b \leq 1.$$

Solving this relaxed problem, we obtain the solution

$$z = 1, \quad b = 1.$$

This solution is already fairly close to the correct solution, despite the rather crude degree-2 Chebyshev approximation of the agent's utility function.

Of course, to obtain a better solution, we need to reduce the approximation error. For this reason, we increase the degree of the approximation to 12. Then we obtain an approximation parameterized by b, with an approximation error of less than $5.5 \cdot 10^{-8}$. Applying the polynomial approach, we obtain the following solution:

$$z = 0.9575, \quad b = 1,$$

which matches the optimal solution reported in Mirrlees (1999).

The Chebyshev approximation for the function V works extremely well in this example. The exponential function is a so-called *entire* (analytic) function, since it can be globally expressed by a power series. The convergence results for the Chebyshev approximation for entire functions are even better than those reported in Appendix B for ν-times differentiable functions; see Trefethen (2013). However, stating these result requires more mathematical background.

5.2. *Application: Executive Compensation Contracts*

At least as early as in Haubrich (1994), principal–agent models have been applied to the study of executive compensation contracts. Armstrong, Larcker, and Su (2010) presented a sophisticated study of optimal compensation contracts that emerge from principal–agent models with pure moral hazard, pure adverse selection, or a combination thereof. They pointed out that "the first-order approach typically fails when realistic contracting features (e.g., nonlinear compensation contracts and nonnormal probability distributions) are incorporated in the model." In order to avoid these difficulties, Armstrong, Larcker, and Su (2010) restricted the agent to choose his action from a finite set and allowed for mixed strategies. The resulting optimization problem is an MPEC, a mathematical program with equilibrium constraints. Since standard solvers are not guaranteed to find globally optimal solutions of MPECs, Armstrong, Larcker, and Su (2010) used different solvers with many different starting points to solve problems hundreds of times.

We now consider a slightly modified version of the model with pure moral hazard from Armstrong, Larcker, and Su (2010). A company hires an agent (a CEO or other high-level executive). The agent chooses an action $a \in [0, 1]$ (contrary to Armstrong, Larcker, and Su (2010), who imposed a finite set of actions). This action influences the stock price p of the company, which is log-normally distributed with the density

$$\mu(p|a) = \frac{1}{p\sigma\sqrt{2\pi}} \exp\left(-\frac{\left(\ln(p) - \rho(a)\right)^2}{2\sigma^2}\right),$$

where $\rho(a) = a^{\phi}$ with the productivity parameter $\phi > 0$. The principal offers the agent a contract with a fixed salary w, restricted stock β_S in the company, and stock options β_O with exercise price K. The monetary payoff to the agent of a contract (w, β_S, β_O, K) at a stock price p is

$$s(w, \beta_S, \beta_O, K; p) = w + p\beta_S + \beta_O \max(p - K, 0).$$

The agent has the utility function

$$v(s, a) = \frac{s^{1-\delta}}{1-\delta} - ca^2.$$

The principal is assumed to be risk-neutral and maximizes the expected net difference between the stock price p and the payment to the agent $s(w, \beta_S, \beta_O, K; p)$.

This model does not satisfy the assumptions of the polynomial optimization approach, since the stock price has a lognormal distribution.[15] To deal with this issue, we first perform a change of variable, $\wp = \ln p$, with the new variable $\wp \in (-\infty, \infty)$. As a result, the expectation in the two expected utility functions is now an integral on the set $\mathbb{R}$. Thus, we can use Gauss–Hermite quadrature to approximate the integral. We denote the nodes by $\wp_i$, $i = 1, 2, \ldots, N$, and the quadrature weights by q_i, $i = 1, 2, \ldots, N$. Moreover, replacing p by $\wp$ in the payoff function s and the density μ leads to the reparameterized functions $\tilde{s}$ and $\tilde{\mu}$, respectively. The resulting expected utility functions are

$$U(w, \beta_S, \beta_O, K, a) = \sum_{i=1}^{N} q_i \tilde{\mu}(\wp_i | a)\big(e^{\wp_i} - \tilde{s}(w, \beta_S, \beta_O, K; \wp_i)\big),$$

$$V\big(s(w, \beta_S, \beta_O, K, \wp), a\big) = \sum_{i=1}^{N} q_i \tilde{\mu}(\wp_i | a) v\big(\tilde{s}(w, \beta_S, \beta_O, K, \wp_i), a\big).$$

Following Armstrong, Larcker, and Su (2010), we set the productivity parameter $\phi = \frac{1}{2}$. After another change of variable, $a = \alpha^{1/\phi}$, the resulting functions in α are entire analytic. Thus, the Chebyshev approximation converges with the order $O(C^{-n})$ to V, for some constant $C > 1$. We set the agent's reservation utility to $\underline{V} = v(1/2, 0)$, his cost coefficient to $c = \frac{3}{4}$, and the scale parameter in the lognormal distribution to $\sigma = \frac{1}{2}$. The number of nodes for the Gauss–Hermite quadrature is 10 and the degree of the Chebyshev approximation is 13. Following Armstrong, Larcker, and Su (2010), we set K exogenously

[15]In fact, as in our general model setup in Section 2, this model has a continuous probability distribution of outcomes. The polynomial optimization approach can handle such models as well. Note that the payment term $s(w, \beta_S, \beta_O, K; p)$ now plays the role of the simple wage w in the discrete model.

TABLE II

NUMERICAL RESULTS FOR THE EXECUTIVE COMPENSATION PROBLEM

δ	U^*	w^*	β_S^*	β_O^*	V^*	a^*
$\frac{1}{8}$	1.227	0	0.067	0.829	1.143	0.847
$\frac{1}{4}$	1.151	0	0.230	0.584	1.333	0.734
$\frac{1}{2}$	1.045	0	0.441	0.240	2.000	0.587
1	0.924	0.122	0.526	0	0.000	0.443
2	0.800	0.413	0.391	0	−1.000	0.319
4	0.688	0.624	0.279	0	−0.333	0.225
8	0.593	0.762	0.196	0	−0.143	0.158

since stock options given to managers typically have a strike price set to the current stock price. For the given lognormal distribution, we set it to the expected value for zero effort, $a = 0$, from the agent, that is, $K = e^{\sigma^2/2}$. To illustrate the functions that are involved in the computation, Appendix C reports the approximations of both utility functions.

Table II displays numerical solutions to the executive compensation problem for different values of the executive's coefficient of risk aversion, δ. We denote the principal's and the agent's maximal expected utilities from the optimal contract by U^* and V^*, respectively. For all values of the parameter δ, the participation constraint is binding, $V^* = v(1/2, 0)$. The principal's payoff and the executive's effort are both decreasing in δ. Unsurprisingly, the number of options, β_O^*, in the optimal contract is decreasing in the risk aversion parameter, while conversely, the fixed wage w^* is increasing in δ. The number of shares, β_S^*, in the optimal contract is increasing in δ as long as options are a part of the contract and the fixed wage is zero; shares replace options in the optimal contract. For sufficiently large values of δ, options cease to be part of the optimal contract and instead a fixed wage becomes part of the contract. As δ increases, the wage increases and the number of shares decreases in the optimal contract; now the fixed wage replaces stock holdings.

6. THE POLYNOMIAL OPTIMIZATION APPROACH FOR $A \subset \mathbb{R}^L$

Principal–agent models in which the agent's action set is one-dimensional dominate not only the literature on the first-order approach but also the applied and computational literature; see, for example, Araujo and Moreira (2001), Judd and Su (2005), and Armstrong, Larcker, and Su (2010). However, the analysis of linear multi-task principal–agent models in Holmström and Milgrom (1991) demonstrates that multivariate agent problems exhibit some fundamental differences in comparison to the common one-dimensional models. For example, the compensation paid to the agent does not only serve the dual purpose of incentive for hard work and risk-sharing but, in addition, influences the agent's attention among his various tasks. The theoretical literature that allows the set of actions to be multidimensional, for example, Grossman and Hart (1983), Kadan, Reny, and Swinkels (2011), and Kadan and Swinkels (2012), focuses on the existence and properties of equilibria. To the best of our knowledge, the first-order approach has not received much attention for models with multidimensional action sets.[16]

We now extend our polynomial optimization approach to principal–agent models in which the agent has more than one decision variable, so $\mathbf{a} \in A \subset \mathbb{R}^L$. For this purpose, we state and prove a generalization of Theorem 1 and subsequently illustrate the multidimensional approach by a numerical example. We refer to Appendix A.2 for some mathematical background on the optimization of multivariate polynomials.

6.1. *The Multivariate Polynomial Optimization Approach*

Consider the principal–agent problem with a multidimensional set of actions, $A \subset \mathbb{R}^L$, and a finite set S of cardinality N. We impose the following assumption.

ASSUMPTION 5—Set of Actions: *The set of actions*, $A = \{\mathbf{a} \in \mathbb{R}^L \mid g_1(\mathbf{a}) \geq 0, \dots, g_m(\mathbf{a}) \geq 0\}$, *is a compact semi-algebraic set with a nonempty interior.*

A multidimensional version of Assumption 4, the assumption that the agent has a rational expected utility function, imposes

$$-V(\mathbf{w}, \mathbf{a}) = -\sum_{i=1}^{N} v(w_i, \mathbf{a})\mu(s_i|\mathbf{a}) = \frac{\sum_{\boldsymbol{\alpha}} c_{\boldsymbol{\alpha}}(\mathbf{w})\mathbf{a}^{\boldsymbol{\alpha}}}{\sum_{\boldsymbol{\alpha}} f_{\boldsymbol{\alpha}}(\mathbf{w})\mathbf{a}^{\boldsymbol{\alpha}}}.$$

Applying the rational generalization of the relaxation (25) from Appendix A.2.1 to the agent's expected utility optimization problem, we obtain the following relaxation for that problem:

$$\begin{aligned} &\sup_{\rho, Q^{(0)}, Q^{(1)}, \dots, Q^{(m)}} \quad \rho \\ &\text{s.t.} \quad \sum_{\boldsymbol{\alpha}} c_{\boldsymbol{\alpha}}(\mathbf{w})\mathbf{b}^{\boldsymbol{\alpha}} - \rho \sum_{\boldsymbol{\alpha}} f_{\boldsymbol{\alpha}}(\mathbf{w})\mathbf{b}^{\boldsymbol{\alpha}} \\ &\qquad = v_d^T Q^{(0)} v_d + \sum_{i=1}^{m} g_i v_{d-d_{g_i}}^T Q^{(i)} v_{d-d_{g_i}}, \end{aligned} \tag{17}$$

[16] Abraham, Koehne, and Pavoni (2011) derived sufficient conditions for the application of the first-order approach in a model in which the agent has two decision variables, effort and savings. Only the chosen effort level affects the probability distribution of outcomes.

$$Q^{(0)} \succcurlyeq 0, \quad Q^{(i)} \succcurlyeq 0 \quad \forall i = 1, 2, \dots, m,$$
$$Q^{(0)} \in \mathbb{R}^{\binom{n+d}{d} \times \binom{n+d}{d}},$$
$$Q^{(i)} \in \mathbb{R}^{\binom{n+d-d_{g_i}}{d-d_{g_i}} \times \binom{n+d-d_{g_i}}{d-d_{g_i}}} \quad \forall i = 1, 2, \dots, m,$$
$$v_d \text{ vector of monomials } \mathbf{b}^{\boldsymbol{\alpha}} \text{ up to degree } d,$$
$$v_{d-d_{g_i}} \text{ vector of monomials } \mathbf{b}^{\boldsymbol{\alpha}} \text{ up to degree } d - d_{g_i}.$$

The equality in the first constraint signifies an equality of the polynomials on the left-hand and right-hand side in the variables $\mathbf{b}$. So, once again, we need to equate the coefficients of two polynomials. These equations, in turn, are polynomials in the matrix elements $Q_{ij}^{(l)}$, $l = 0, 1, \dots, m$, and the variable ρ. Next we use Proposition A.1 from the Supplemental Material and replace the positive semidefinite matrices $Q^{(i)}$ by $L_{(i)}(L_{(i)})^T$, where $L_{(i)}$ are lower triangular matrices (with a nonnegative diagonal). This transformation allows us to drop the explicit constraints on positive semidefiniteness.

For a reformulation of the original principal–agent problem from a bilevel problem to a nonlinear program, we need to characterize the optimal choice of the agent via equations or inequalities. In the case of one-dimensional effort, this reformulation is (12.d), the generalization of which for multidimensional effort would be

$$\sum_{i=0}^{d} c_i(\mathbf{w})\mathbf{a}^i - \rho\left(\sum_{i=0}^{d} f_i(\mathbf{w})\mathbf{a}^i\right) = 0.$$

Unfortunately, since the relaxation of the agent's problem gives us a lower bound, we cannot, in general, impose this constraint. The resulting nonlinear program would most likely be infeasible. Instead, we use an idea of Couzoudis and Renner (2013), who allowed for solutions of optimization problems to be only approximately optimal; we do not force the left-hand side to be zero but instead only impose a small positive upper bound.

Now we are in the position to state and prove our second theorem, a multivariate extension of Theorem 1.

THEOREM 2: *Suppose the agent's expected utility maximization problem satisfies Assumption 5 and the multidimensional version of Assumption 4. Let* $\mathbf{v}_k$ *be the vector of monomials in* $b_1, \dots, b_L$ *up to degree* k. *Let* $d \in \mathbb{N}$ *and* $\varepsilon > 0$. *Including* $\rho \in \mathbb{R}$ *and lower triangular matrices* $L_{(0)} \in \mathbb{R}^{\binom{n+d}{d} \times \binom{n+d}{d}}$ *and* $L_{(i)} \in \mathbb{R}^{\binom{n+d-d_{g_i}}{d-d_{g_i}} \times \binom{n+d-d_{g_i}}{d-d_{g_i}}}$ *for* $i = 1, \dots, m$, *as additional decision variables, define* *the following relaxation of the principal–agent problem* (1):

$$\max_{\mathbf{w},\mathbf{a},\rho,L_{(0)},\dots,L_{(m)}} U(\mathbf{w},a) \tag{18}$$

$$\text{s.t.} \quad \sum_{\boldsymbol{\alpha}} c_{\boldsymbol{\alpha}}(\mathbf{w})\mathbf{b}^{\boldsymbol{\alpha}} - \rho \sum_{\boldsymbol{\alpha}} f_{\boldsymbol{\alpha}}(\mathbf{w})\mathbf{b}^{\boldsymbol{\alpha}} = \mathbf{v}_d^T L_{(0)} L_{(0)}^T \mathbf{v}_d + \sum_{i=1}^{m} g_i \mathbf{v}_{d-d_{g_i}}^T L_{(i)} L_{(i)}^T \mathbf{v}_{d-d_{g_i}}, \tag{18.a}$$

$$\varepsilon \sum_{\boldsymbol{\alpha}} f_{\boldsymbol{\alpha}}(\mathbf{w})\mathbf{a}^{\boldsymbol{\alpha}} \geq \sum_{\boldsymbol{\alpha}} c_{\boldsymbol{\alpha}}(\mathbf{w})\mathbf{a}^{\boldsymbol{\alpha}} - \rho \sum_{\boldsymbol{\alpha}} f_{\boldsymbol{\alpha}}(\mathbf{w})\mathbf{a}^{\boldsymbol{\alpha}}, \tag{18.b}$$

$$\sum_{i=0}^{d} c_i(\mathbf{w})\mathbf{a}^i \leq -\underline{V}\left(\sum_{i=0}^{d} f_i(\mathbf{w})\mathbf{a}^i\right), \tag{18.c}$$

$$g_i(\mathbf{a}) \geq 0 \quad \forall i = 1, 2, \dots, m, \tag{18.d}$$

$$\mathbf{w} \in W. \tag{18.e}$$

This optimization problem has the following properties:

(1) *Any feasible point*, $(\hat{\mathbf{w}}, \hat{\mathbf{a}}, \hat{\rho}, \hat{L}_{(0)}, \dots, \hat{L}_{(m)})$, *satisfies the inequality*

$$\max_{\mathbf{a}\in A} V(\hat{\mathbf{w}}, \mathbf{a}) - V(\hat{\mathbf{w}}, \hat{\mathbf{a}}) \leq \varepsilon. \tag{19}$$

(2) *Let* $(\bar{\mathbf{w}}, \bar{\mathbf{a}})$ *be a feasible solution of the principal–agent problem* (1). *Then*, *for any* $\varepsilon > 0$, *there exist* $d(\varepsilon) \in \mathbb{N}$ *and* $\bar{\rho}, \bar{L}_{(0)}, \dots, \bar{L}_{(m)}$, *such that* $(\bar{\mathbf{w}}, \bar{\mathbf{a}}, \bar{\rho}, \bar{L}_{(0)}, \dots, \bar{L}_{(m)})$ *is feasible for the relaxation* (18) *for* $d = d(\varepsilon)$.

(3) *Let* $(\bar{\mathbf{w}}, \bar{\mathbf{a}})$ *be an optimal solution to* (1). *For any* ε, *let* $d(\varepsilon)$ *be as in point* (2). *Denote by* $u(\varepsilon)$ *the optimal objective value of the relaxation* (18) *for given* ε *and* $d = d_\varepsilon$. *Then* $\lim_{\varepsilon\to 0^+} u(\varepsilon) = U(\bar{\mathbf{w}}, \bar{\mathbf{a}})$.

(4) *Again*, *let* $(\bar{\mathbf{w}}, \bar{\mathbf{a}})$ *be an optimal solution to* (1), *and for any* ε, *let* $d(\varepsilon)$ *be as in point* (2). *Then*, *the set of limit points for* $\varepsilon \to 0^+$ *of any sequence of optimal solutions to* (18), $(\mathbf{w}^*(\varepsilon), \mathbf{a}^*(\varepsilon), \rho^*(\varepsilon), L_{(0)}^*(\varepsilon), \dots, L_{(m)}^*(\varepsilon))$, *projected onto* $W \times A$, *is contained in the set of optimal solutions to the original principal–agent problem* (1).

Before we prove the theorem, we briefly describe the optimization problem (18). This problem has the same objective function as the original principal–agent problem (1). Constraint (18.a) uses a sum of squares representation of positive polynomials to ensure that for a contract $\mathbf{w}$ chosen by the principal, $-V(\mathbf{w}, \mathbf{a}) \geq \rho$ for all $\mathbf{a} \in A$. It is important to emphasize that this equation holds not only for the optimal choice but in fact for all possible $\mathbf{a} \in A$. Therefore, for the purpose of this constraint, we need to duplicate the effort vector $\mathbf{a}$; in the functional equation (18.a), we denote effort by $\mathbf{b}$. Thus again $\mathbf{b}$ is not a variable in the optimization problem. We obtain the equations by comparing the coefficients of the polynomials in $\mathbf{b}$. The positive semidefinite matrices in the relaxation of the agent's problem (17) are represented via products of lower triangular matrices. Proposition A.1 in the Supplemental Material shows that any positive semidefinite matrix can be represented in this fashion (even having the property that all diagonal elements are nonnegative). Put differently, constraint (18.a) ensures that $-\rho$ is an upper bound on the agent's possible expected utility levels. Next, constraint (18.b) imposes a lower bound on the agent's expected utility level, namely, $V(\mathbf{w}, \mathbf{a}) + \varepsilon \geq -\rho$. Therefore, the constraints (18.a) and (18.b) force the value of a in any feasible solution to result in a utility for the agent satisfying $-\rho - \varepsilon \leq V(\mathbf{w}, \mathbf{a}) \leq -\rho$. That is, for a given contract $\mathbf{w}$, the first two constraints ensure an effort choice by the agent that is within ε of being optimal. The last three constraints are straightforward. Constraint (18.c) is the transformed participation constraint for the agent's rational expected utility function. Constraint (18.d) defines the set of the feasible actions and constraint (18.e) is just the constraint on the compensation scheme from the original principal–agent problem (1).

PROOF OF THEOREM 2: Under the assumptions of the theorem, the agent's constraints satisfy the conditions of Putinar's Positivstellensatz and so we obtain the sums-of-squares representation for the agent's problem. For fixed d, we then restrict the degree of the sum of squares coefficients as is done in the relaxation.

(1) Every feasible point $(\hat{\mathbf{w}}, \hat{\mathbf{a}}, \hat{\rho}, L_{(0)}, \ldots, L_{(m)})$ provides an upper bound $-\hat{\rho}$ on the maximal value of $V(\hat{\mathbf{w}}, \mathbf{a}) = -\frac{\sum_\alpha c_\alpha(\hat{\mathbf{w}})\mathbf{a}^\alpha}{\sum_\alpha f_\alpha(\hat{\mathbf{w}})\mathbf{a}^\alpha}$, since (18.a) implies that

$$\sum_{\alpha} c_{\alpha}(\mathbf{w})\mathbf{b}^{\alpha} - \hat{\rho} \sum_{\alpha} f_{\alpha}(\mathbf{w})\mathbf{b}^{\alpha} \geq 0$$

and so, $-\hat{\rho} \geq \max_{\mathbf{a}\in A} V(\hat{\mathbf{w}}, \mathbf{a}) \geq V(\hat{\mathbf{w}}, \hat{\mathbf{a}})$. Moreover, constraint (18.b) implies that

$$\varepsilon \geq -\hat{\rho} - V(\hat{\mathbf{w}}, \hat{\mathbf{a}}) \geq \max_{\mathbf{a}\in A} V(\hat{\mathbf{w}}, \mathbf{a}) - V(\hat{\mathbf{w}}, \hat{\boldsymbol{a}}).$$

Thus, property (19) holds.

(2) Under the assumptions of the theorem, Proposition A.5 from the Supplemental Material implies that for each fixed $\bar{\mathbf{w}}$ and a given $\varepsilon > 0$, there exists a d such that $V(\mathbf{w}, \mathbf{a}) - \rho$ has the representation (21) of Putinar's Positivstellensatz with degree d coefficients. For this d, problem (18) has a nonempty feasible region.

(3) Recall the agent's optimal value function $\Psi : W \to \mathbb{R}$ from the proof of Proposition 1. The projection of the set of feasible points of problem (18) to $W \times A$ is a subset of

$$S(\varepsilon) = \{(\mathbf{w}, \mathbf{a}) \in W \times A \mid \Psi(\mathbf{w}) - V(\mathbf{w}, \mathbf{a}) \leq \varepsilon\}$$

and, by point (2), contains $(\bar{\mathbf{w}}, \bar{\mathbf{a}})$. Let $v(\varepsilon) = \max_{(\mathbf{w},\mathbf{a})\in S(\varepsilon)} U(\mathbf{w}, \mathbf{a})$. Then

$$U(\bar{\mathbf{w}}, \bar{\mathbf{a}}) \leq u(\varepsilon) \leq v(\varepsilon).$$

Furthermore, since Ψ and V are continuous (Berge's Maximum Theorem), the set $S(\varepsilon)$ is upper hemicontinuous and uniformly compact near 0.[17] By Hogan (1973, Theorem 5), it follows that v is upper semi-continuous and thus we have

$$\begin{aligned} U(\bar{\mathbf{w}}, \bar{\mathbf{a}}) &\leq \liminf_{\varepsilon\to 0^+} u(\varepsilon) \leq \limsup_{\varepsilon\to 0^+} u(\varepsilon) \\ &\leq \limsup_{\varepsilon\to 0^+} v(\varepsilon) \leq v(0) = U(\bar{\mathbf{w}}, \bar{\mathbf{a}}). \end{aligned}$$

Therefore, $\lim_{\varepsilon\to 0^+} u(\varepsilon) = U(\bar{\mathbf{w}}, \bar{\mathbf{a}})$.

(4) Consider any limit point $(\mathbf{w}_0, \mathbf{a}_0) \in W \times A$ and any sequence $(\mathbf{w}_\varepsilon, \mathbf{a}_\varepsilon)$ converging to it for $\varepsilon \to 0$. Property 2 implies that $U(\mathbf{w}_\varepsilon, \mathbf{a}_\varepsilon) \to U(\bar{\mathbf{w}}, \bar{\mathbf{a}})$. By continuity of Ψ and V, we also have

$$\lim_{\varepsilon\to 0^+} \big(\Psi(\mathbf{w}_\varepsilon) - V(\mathbf{w}_\varepsilon, \mathbf{a}_\varepsilon)\big) = \Psi(\mathbf{w}_0) - V(\mathbf{w}_0, \mathbf{a}_0) = 0.$$

Thus, $(\mathbf{w}_0, \mathbf{a}_0)$ is feasible for (1) and attains the optimal value.
This completes the proof of Theorem 2. *Q.E.D.*

Some comments on the technical convergence results of Theorem 2 are in order. For the one-dimensional effort case, Theorem 1 provides a single well-defined optimization problem that is equivalent to the original principal–agent problem. Ideally, we would like to obtain a similar result for the multidimensional effort case. Unfortunately, in general, that is impossible. A comparison of the sum of squares representation results for univariate and multivariate polynomials reveals the critical difference between the two cases. Proposition 2 in Section 3.2, the "Positivstellensatz" for univariate polynomials, provides a sum of squares representation of nonnegative univariate polynomials with an explicit (small) bound on the degree of the involved sums of squares. Proposition A.2 in the Supplemental Material, Putinar's Positivstellensatz, provides a sum of squares representation of positive multivariate polynomials; however, there is no a priori upper bound on the degree of the involved sums of squares. In fact, from a purely theoretical viewpoint, the necessary degree may be infinite. As a result, any finite-degree representation as in (24) in the Supplemental Material may only constitute a relaxation of the original polynomial optimization problem.

Once we have computed a solution, we can always verify if it is feasible. To accomplish this, we fix $\mathbf{w}$ and solve the polynomial optimization problem for the agent to global optimality using GloptiPoly.

[17] Upper hemicontinuity at 0 means that for any sequence $\varepsilon^k \to 0$, $s^k \in S(\varepsilon^k)$ and $s^k \to s$ imply $s \in S(0)$.

In light of the theoretical difficulties for general multivariate polynomials, it is of great interest to characterize polynomial optimization problems that offer a guaranteed convergence of the relaxation for finite d. If both the objective function and the constraints are s.o.s. convex, then the convergence is finite; see Lasserre (2010, Theorem 5.15).[18] Also, if the objective function is strictly convex and the constraints are convex, then convergence is finite; see Lasserre (2010, Theorem 5.16). The problem of finite convergence continues to be an active research issue in algebraic geometry. For example, Nie (2012) proved finite convergence under a regularity condition on the set of constraints. His approach requires a reformulation of the problem by adding constraints consisting of minors of a Jacobian derived from the KKT conditions. Unfortunately, it appears to be rather difficult to check the regularity condition in applications.

As a final remark, we point out that Schmüdgen's Positivstellensatz (see Schmüdgen (1991)) yields a representation of multivariate positive polynomials that is different than that of Putinar's Positivstellensatz. This representation is slightly more general but requires higher degree sums of squares. Therefore, it appears to be less attractive for economic applications.

6.2. *Multivariate Examples*

In the following, we show the results from two applications of the multivariate polynomial optimization approach. The first example is deliberately simple so that the functions and constraints for a numerical solution of the multivariate problem are straightforward. The second application tackles a production problem from the economics literature.

6.2.1. *Illustrative Example*

Let the set of outcomes be $\{0, 3, 6\}$ with probabilities

$$\left\{\frac{1+a/2+b}{1+a+b}, \frac{b}{1+a+b}, \frac{a/2-b}{1+a+b}\right\},$$

satisfying the constraints

$$b \geq 0, \quad a - 2b \geq 0,$$

which assure that the probability functions are nonnegative. The outcome distribution has mean and variance

$$\frac{3(a-b)}{1+a+b} \quad \text{and} \quad \frac{9\big(2a+a^2-3b+ab-4b^2\big)}{(1+a+b)^2},$$

[18] A polynomial f is called *s.o.s.* convex, iff $\nabla^2 f = WW^T$ for some matrix W.

respectively. Note that the effort a increases both the expected value and the variance of the outcome. On the contrary, the effort b decreases the expectation and the variance.

The principal's and the agent's Bernoulli utility functions are

$$\begin{aligned}
&u(y, w) = -(-6 - w + y)^2 \quad \text{and} \\
&v(a, b, w) = (1 + a + b)\left(-a - \frac{b}{10} + \ln(1 + w)\right),
\end{aligned}$$

respectively. The expected utility of the agent is

$$\begin{aligned}
&\frac{1}{10}\big(-10a - 10a^2 - b - 11ab - b^2 + 10b \ln(1 + w_2) \\
&\quad + 5(a - 2b)\ln(1 + w_3)\big) + \left(1 + \frac{a}{2} + b\right)\ln(1 + w_1)
\end{aligned}$$

and the expected utility of the principal is

$$\begin{aligned}
&-\big(a\big(36 + 12w_1 + w_1^2 + w_3^2\big) \\
&\quad + 2\big((6 + w_1)^2 + b\big(45 + 12w_1 + w_1^2 + 6w_2 + w_2^2 - w_3^2\big)\big)\big) \\
&\quad /\big(2(1 + a + b)\big).
\end{aligned}$$

In sum, in the principal–agent problem, the lower-level problem has an objective function that is polynomial in the agent's decision variables a and b, but nonpolynomial in the principal's decision variables w_1, w_2, w_3. The upper-level problem is rational in a, b and polynomial in w_1, w_2, w_3. And so we can apply Theorem 2 and reformulate the principal–agent problem.

We observe that the largest degree in the variables a and b is 2. So, we can choose the degree of the relaxation to be 2, that is, all the matrices appearing will be of size 3×3, $L_k = (s_{k,i,j})_{i,j=1,2,3}$, where L_k is a lower triangular matrix with nonnegative diagonal. The sum of squares multipliers now appear as follows:

$$\begin{aligned}
\sigma_k = {} & s_{k,1,1}^2 + 2as_{k,1,1}s_{k,2,1} + a^2\big(s_{k,2,1}^2 + s_{k,2,2}^2\big) + 2bs_{k,1,1}s_{k,3,1} \\
& + b^2\big(s_{k,3,1}^2 + s_{k,3,2}^2 + s_{k,3,3}^2\big) + ab(2s_{k,2,1}s_{k,3,1} + 2s_{k,2,2}s_{k,3,2}).
\end{aligned}$$

Thus, the coefficients in the variables a, b of the following polynomial have to be zero:

$$V(a, b, w_1, w_2, w_3) + \rho + \sigma_0 + b\sigma_1 + (a - 2b)\sigma_2 + (1 - a)\sigma_3.$$

This leads to the following equations:

$$
\begin{aligned}
0 &= s_{1,3,1}^2 + s_{1,3,2}^2 + s_{1,3,3}^2 - s_{2,3,1}^2 - s_{2,3,2}^2 - s_{2,3,3}^2, \\
0 &= \frac{1}{2}\left(s_{2,2,1}^2 + s_{2,2,2}^2\right) - s_{3,2,1}^2 - s_{3,2,2}^2, \\
0 &= -1 + s_{0,2,1}^2 + s_{0,2,2}^2 + s_{2,1,1}s_{2,2,1} - 2s_{3,1,1}s_{3,2,1} + s_{3,2,1}^2 + s_{3,2,2}^2, \\
0 &= s_{1,2,1}^2 + s_{1,2,2}^2 - s_{2,2,1}^2 - s_{2,2,2}^2 + s_{2,2,1}s_{2,3,1} \\
&\quad + s_{2,2,2}s_{2,3,2} - 2(s_{3,2,1}s_{3,3,1} + s_{3,2,2}s_{3,3,2}), \\
0 &= -\frac{11}{10} + 2(s_{0,2,1}s_{0,3,1} + s_{0,2,2}s_{0,3,2}) + 2s_{1,1,1}s_{1,2,1} - 2s_{2,1,1}s_{2,2,1} \\
&\quad + s_{2,1,1}s_{2,3,1} - 2s_{3,1,1}s_{3,3,1} + 2(s_{3,2,1}s_{3,3,1} + s_{3,2,2}s_{3,3,2}), \\
0 &= 2(s_{1,2,1}s_{1,3,1} + s_{1,2,2}s_{1,3,2}) - 2(s_{2,2,1}s_{2,3,1} + s_{2,2,2}s_{2,3,2}) \\
&\quad + \frac{1}{2}\left(s_{2,3,1}^2 + s_{2,3,2}^2 + s_{2,3,3}^2\right) - s_{3,3,1}^2 - s_{3,3,2}^2 - s_{3,3,3}^2, \\
0 &= -\frac{1}{10} + s_{0,3,1}^2 + s_{0,3,2}^2 + s_{0,3,3}^2 + 2s_{1,1,1}s_{1,3,1} \\
&\quad - 2s_{2,1,1}s_{2,3,1} + s_{3,3,1}^2 + s_{3,3,2}^2 + s_{3,3,3}^2, \\
0 &= \rho + s_{0,1,1}^2 + s_{3,1,1}^2 + \ln(1 + w_1), \\
0 &= -\frac{1}{10} + 2s_{0,1,1}s_{0,3,1} + s_{1,1,1}^2 - s_{2,1,1}^2 + 2s_{3,1,1}s_{3,3,1} \\
&\quad + \ln(1 + w_1) + \ln(1 + w_2) - \ln(1 + w_3), \\
0 &= -1 + 2s_{0,1,1}s_{0,2,1} + \frac{s_{2,1,1}^2}{2} - s_{3,1,1}^2 + 2s_{3,1,1}s_{3,2,1} \\
&\quad + \frac{1}{2}\ln(1 + w_1) + \frac{1}{2}\ln(1 + w_3).
\end{aligned}
$$

We set the reservation utility to $\frac{3}{2}$ and solve this problem with large-scale NLP solver IPOPT; see Wächter and Biegler (2006). As of the writing of this paper, we cannot apply GloptiPoly to this problem since the number of variables is too large. Therefore, we do not approximate the logarithmic terms but leave them in the optimization problem for IPOPT. We obtain the following solutions:

$$
\begin{aligned}
&a = 0.34156, \quad b = 0.17078, \quad w_1 = 2.7295, \\
&w_2 = 4.0491, \quad 10 \geq w_3 \geq 0.
\end{aligned}
$$

Note that $a-2b=0$ and so the third outcome has probability zero. Therefore, there are a continuum of possible values for w_3. The principal's expected utility is -73.210 and the agent's expected utility is $\frac{3}{2}$. (In this example, it is possible to compute the solution for $\varepsilon=0$ since, for quadratic multivariate polynomials, the sets of sums of squares and nonnegative polynomials are identical; see Appendix A.2.)

6.2.2. *Application: Technology Moral Hazard Program*

The following application is a nonlinear extension of the linear-effort example from Prescott (2004, Section 4.3). The principal is very averse to low output, which is reflected in her discontinuous utility function

$$u(c,q)=\begin{cases} q-c-20, & q\leq 0.21,\\ q-c, & q>0.21.\end{cases}$$

The principal offers a consumption function $c:\mathbb{R}\to\mathbb{R}$ as a contract to the agent with consumption depending on the output q produced by the agent. The agent has a production technology and he controls the mean and the standard deviation of the stochastic output. The agent is risk-averse and receives disutility from increasing the mean or lowering the standard deviation of the output distribution. We replace the linear disutility terms in the Prescott (2004) example by nonlinear expressions and choose as the agent's utility

$$v(c,a_m,a_s)=\sqrt{c}-a_m^2+\sqrt{a_s+1},$$

with a_m and a_s denoting the agent's choices for the mean and the standard deviation of output, respectively.

Following Prescott (2004), we discretize the output levels on the interval $[0,2]$ and approximate a normal distribution of the output. Let $(w_i,q_i)_{i=1,2,\dots,N}$ denote the Gauss–Legendre weights and nodes on the interval $[0,2]$. The output distribution over $\{q_1,\dots,q_N\}$ is as follows:

$$\mathbb{P}(q=q_i|a_m,a_s)=\frac{w_i}{\sqrt{2\pi a_s^2}}\exp\left(-\frac{(w_i-a_m)^2}{2a_s^2}\right)\Big/\left(\sum_{j=1}^{N}\frac{w_j}{\sqrt{2\pi a_s^2}}\exp\left(-\frac{(w_j-a_m)^2}{2a_s^2}\right)\right).$$

(Prescott (2004) implicitly used a similar distribution, however, with equidistant points and no weights.) We obtain the following expected utility for the principal:

$$\sum_j \mathbb{P}(q=q_j|a_m,a_s)\big(q_j-c_j-20\cdot\mathbf{1}_{\leq 0.21}(q_j)\big),$$

where $\mathbf{1}_{\leq 0.21}(q_j)$ denotes the indicator function. The agent's expected utility is

$$\sum_j \mathbb{P}(q = q_j | a_m, a_s)\sqrt{c_j} - a_m^2 + \sqrt{a_s + 1}.$$

Again following Prescott (2004), we compactify the agent's effort set and impose the following constraints:

$$0 \leq a_m \leq 1.5,$$
$$0.3 \leq a_s \leq 0.6.$$

As a reservation utility, we choose $\max_{a_m \geq 0, a_v \leq 0.6} v(0, a_m, a_s) = \sqrt{0.6 + 1} \approx 1.26491$.

We use the tensor product of the Chebyshev polynomials to approximate the agent's utility function with degree pair $(5, 5)$. More precisely, let x_i denote the Chebyshev nodes for degree 5 approximation. Then, after a suitable change of coordinates, we approximate the agent's function on the grid $(x_i, x_j)_{i,j=0,\dots,5}$ with basis functions $\{T_i \otimes T_j\}_{i,j=0,\dots,5}$. In particular, the maximum degree of the basis polynomials is 10. We use 12 nodes for the definition of the probability distribution. And lastly, we choose 12 as the degree for the relaxation. These choices result in the optimal solution displayed in Table III.

Since the principal is very averse to low output, she pays nothing to the agent, $c = 0$, both for the four lowest and the four highest output levels. Such a compensation scheme encourages the agent to choose the output standard deviation, a_s, as low as possible in order to minimize the probability of an extreme outcome. In addition, the agent chooses the output mean, a_m, so that the four

TABLE III

NUMERICAL SOLUTION TO THE PRODUCTION PROBLEM

a_m	a_s	Obj. Agent	Obj. Principal
0.750	0.3	2.000	−2.394

i	1	2	3	4
c_i	0	0	0	0
$\mathbb{P}(q_i \| a_m, a_s)$	0.0216	0.0516	0.0836	0.1147
i	5	6	7	8
c_i	1.5759	5.1741	5.0364	0.3421
$\mathbb{P}(q_i \| a_m, a_s)$	0.1385	0.1477	0.1387	0.1153
i	9	10	11	12
c_i	0	0	0	0
$\mathbb{P}(q_i \| a_m, a_s)$	0.0855	0.0569	0.0328	0.0132

outcomes for which he receives a positive compensation have the highest probabilities.

7. CONCLUSION

In this paper, we have presented a polynomial optimization approach to moral hazard principal–agent problems. Under the assumption that the agent's expected utility function is a rational function of his effort, we have reformulated the agent's maximization problem as an equivalent system of equations and inequalities. This reformulation allowed us to transform the principal–agent problem from a bilevel optimization problem to a nonlinear program. Furthermore, under the assumptions that the principal's expected utility is polynomial and the agent's expected utility is rational in wages (as well as mild assumptions on the effort set and the set of wage choices), we have shown that the resulting NLP is a polynomial optimization problem. Therefore, techniques from global polynomial optimization enable us to solve the NLP to global optimality. We have also shown how to apply the polynomial approach to nonpolynomial problems using Chebyshev approximations of nonpolynomial Bernoulli utility and probability functions.

After the analysis of principal–agent problems with a one-dimensional effort choice for the agent, we have also presented a polynomial optimization approach for problems with multidimensional effort sets. The solution approach for solving such multidimensional problems rests on the same ideas as the approach for the one-dimensional effort model; however, it is technically more difficult. For multidimensional problems, we cannot provide an exact reformulation of the agent's problem but only a relaxation of that problem. Despite this theoretical limitation, the relaxation appears to be often exact in applications.

Our polynomial optimization approach has a number of attractive features. First, we need neither the Mirrlees–Rogerson (or Jewitt) conditions of the classical first-order approach nor the assumption that the agent's utility function is separable. We also do not need to assume that the principal is risk-neutral. Second, under the additional aforementioned assumptions on the utility functions, the final NLP is a polynomial problem that can be solved to global optimality without concerns about constraint qualifications. Third, unlike the first-order approach, the polynomial approach extends quite generally to models with multidimensional effort sets.

The technical assumptions underlying the polynomial approach, while limiting, are not detrimental. The most serious limitation of our polynomial optimization approach is that it is not suited for a subsequent traditional theoretical analysis of the principal–agent model. Despite this shortcoming, the polynomial approach can serve as a useful tool to examine the generality of the insights derived from the very restrictive first-order approach. The ability of the approach to find global solutions to principal–agent problems is one of its hallmarks.

REFERENCES

ABRAHAM, A., S. KOEHNE, AND N. PAVONI (2011): "On the First-Order Approach in Principal–Agent Models With Hidden Borrowing and Lending," *Journal of Economic Theory*, 146, 1331–1361. [757]

ALVI, E. (1997): "First-Order Approach to Principal–Agent Problems: A Generalization," *The GENEVA Papers on Risk and Insurance*, 22, 59–65. [735]

ARAUJO, A., AND H. MOREIRA (2001): "A General Lagrangian Approach for Non-Concave Moral Hazard Problems," *Journal of Mathematical Economics*, 35, 17–39. [731,756]

ARMSTRONG, C. S., D. F. LARCKER, AND C.-L. SU (2010): "Endogenous Selection and Moral Hazard in Compensation Contracts," *Operations Research*, 58, 1090–1106. [731,750,754-756]

BERGE, C. (1963): *Topological Spaces*. Edinburgh: Oliver & Boyd. [737]

BOYD, S., AND L. VANDENBERGHE (2004): *Convex Optimization*. Cambridge: Cambridge University Press. [730]

CONLON, J. R. (2009): "Two New Conditions Supporting the First-Order Approach to Multisignal Principal–Agent Problems," *Econometrica*, 77, 249–278. [731,735,736]

COUZOUDIS, E., AND P. RENNER (2013): "Computing Generalized Nash Equilibria by Polynomial Programming," *Mathematical Methods of Operations Research*, 77, 459–472. [758]

DUBÉ, J.-P., J. T. FOX, AND C.-L. SU (2012): "Improving the Numerical Performance of Static and Dynamic Aggregate Discrete Choice Random Coefficients Demand Estimation," *Econometrica*, 80, 2231–2267. [732]

GREUEL, G.-M., AND G. PFISTER (2002): *A Singular Introduction to Commutative Algebra*. Berlin: Springer-Verlag. [737]

GROSSMAN, S. J., AND O. D. HART (1983): "An Analysis of the Principal–Agent Problem," *Econometrica*, 51, 7–45. [731,757]

HAUBRICH, J. G. (1994): "Risk Aversion, Performance Pay, and the Principal–Agent Problem," *Journal of Political Economy*, 102, 258–276. [754]

HENRION, D., J. LASSERRE, AND J. LÖFBERG (2009): "GloptiPoly 3: Moments, Optimization and Semidefinite Programming," *Optimization Methods and Software*, 24, 761–779. [730,747, 748]

HOGAN, W. W. (1973): "Point-to-Set Maps in Mathematical Programming," *SIAM Review*, 15, 591–603. [761]

HOLMSTRÖM, B. (1979): "Moral Hazard and Observability," *The Bell Journal of Economics*, 10, 74–91. [750]

HOLMSTRÖM, B., AND P. MILGROM (1991): "Multitask Principal–Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law, Economics, & Organization*, 7, 24–52. [731,756]

JEWITT, I. (1988): "Justifying the First-Order Approach to Principal–Agent Problems," *Econometrica*, 56, 1177–1190. [731,735,736,750]

JEWITT, I., O. KADAN, AND J. M. SWINKELS (2008): "Moral Hazard With Bounded Payments," *Journal of Economic Theory*, 143, 59–82. [735,750]

JIBETEAN, D., AND E. DE KLERK (2006): "Global Optimization of Rational Functions: A Semidefinite Programming Approach," *Mathematical Programming*, 106, 93–109. [730,732, 742]

JUDD, K., AND C.-L. SU (2005): "Computation of Moral-Hazard Problems," Working Paper. [731, 732,756]

KADAN, O., AND J. M. SWINKELS (2012): "On the Moral Hazard Problem Without the First-Order Approach," Working Paper. [731,757]

KADAN, O., P. RENY, AND J. M. SWINKELS (2011): "Existence of Optimal Mechanisms in Principal–Agent Problems," Working Paper 2011-002, MFI. [731,736,757]

LASSERRE, J. (2010): *Moments, Positive Polynomials and Their Applications*. Imperial College Press Optimization Series, Vol. 1. London: Imperial College Press. [737-739,742,762]

LAURENT, M. (2009): "Sums of Squares, Moment Matrices and Optimization Over Polynomials," in *Emerging Applications of Algebraic Geometry*. IMA Volumes in Mathematics and Its Applications, Vol. 149. New York: Springer, 157–270. [737,739]

LICALZI, M., AND S. SPAETER (2003): "Distributions for the First-Order Approach to Principal–Agent Problems," *Economic Theory*, 21, 167–173. [731]

MIRRLEES, J. A. (1979): "The Implications of Moral Hazard for Optimal Insurance," Report. [731,735,736]

——— (1999): "The Theory of Moral Hazard and Unobservable Behaviour: Part I," *Review of Economic Studies*, 66, 3–21. [731,735,750,751,754]

NIE, J. (2012): "An Exact Jacobian SDP Relaxation for Polynomial Optimization," *Mathematical Programming*, 137, 1–31. [762]

PRESCOTT, E. S. (1999): "A Primer on Moral-Hazard Models," *Economic Quarterly*, 85, 47–78. [732]

——— (2004): "Computing Solutions to Moral-Hazard Programs Using the Dantzig–Wolfe Decomposition Algorithm," *Journal of Economic Dynamics and Control*, 28, 777–800. [732,765, 766]

RENNER, P., AND K. SCHMEDDERS (2015): "Supplement to 'A Polynomial Optimization Approach to Principal–Agent Problems'," *Econometrica Supplemental Material*, 83, http://dx.doi.org/10.3982/ECTA11351. [732]

ROGERSON, W. P. (1985): "The First-Order Approach to Principal–Agent Problems," *Econometrica*, 53, 1357–1367. [731,735,736,749,750]

SCHMÜDGEN, K. (1991): "The K-Moment Problem for Compact Semi-Algebraic Sets," *Mathematische Annalen*, 289, 203–206. [762]

SINCLAIR-DESGAGNÉ, B. (1994): "The First-Order Approach to Multi-Signal Principal–Agent Problems," *Econometrica*, 62, 459–465. [731,735,736]

SU, C.-L., AND K. L. JUDD (2012): "Constrained Optimization Approaches to Estimation of Structural Models," *Econometrica*, 80, 2213–2230. [732]

TREFETHEN, L. N. (2013): *Approximation Theory and Approximation Practice*. Philadelphia, PA: SIAM. [754]

VANDENBERGHE, L., AND S. BOYD (1996): "Semidefinite Programming," *SIAM Review*, 38, 49–95. [730]

WÄCHTER, A., AND L. T. BIEGLER (2006): "On the Implementation of an Interior-Point Filter Line-Search Algorithm for Large-Scale Nonlinear Programming," *Mathematical Programming*, 106, 25–57. [764]

Hoover Institution, Stanford University, 434 Galvez Mall, Stanford, CA 94305, U.S.A.; phrenner@gmail.com

and

University of Zurich and Swiss Finance Institute, Moussonstrasse 15, 8044 Zurich, Switzerland; karl.schmedders@business.uzh.ch.

Manuscript received January, 2013; final revision received September, 2014.