Journal of Mathematical Economics 64 (2016) 77–98

Contents lists available at ScienceDirect

Journal of Mathematical Economics

journal homepage: www.elsevier.com/locate/jmateco

Solving dynamic public insurance games with endogenous agent distributions: Theory and computational approximation[☆]

Timothy Kam [*], Ronald Stauber

Research School of Economics, Australian National University, ACT 2601, Australia

ARTICLE INFO

Article history:
Received 26 July 2015
Received in revised form
14 March 2016
Accepted 18 March 2016
Available online 7 April 2016

Keywords:
Heterogeneity
Dynamic games
Bilinear programs
Approximation

ABSTRACT

We make two contributions in this paper. First, we extend the characterization of equilibrium payoff correspondences in history-dependent dynamic policy games to a class with endogenously heterogeneous private agents. In contrast to policy games involving representative agents, this extension has interesting consequences as it implies additional nonlinearity (i.e., bilinearity) between the game states (distributions) and continuation/promised values in the policymaker's objective and incentive constraints. The second contribution of our paper is in addressing the computational challenges arising from this payoff-relevant nonlinearity. Exploiting the game's structure, we propose implementable approximate bilinear programming formulations to construct estimates of the equilibrium value correspondence. Our approximation method respects the property of upper hemicontinuity in the target correspondence. We provide small-scale computational examples as proofs of concept.

1. Introduction

The contribution of this paper is twofold. First, we extend the characterization of symmetric sequential equilibrium (SSE) value correspondences in history-dependent dynamic policy games (Phelan and Stacchetti, 2001) to a setting with endogenously heterogeneous private agents.

Second, the extended SSE value correspondence description has an interesting consequence: Agent heterogeneity implies that the domain of the equilibrium payoff correspondence – the relevant game state space – is now a set of distributions over agent types, and these distributions interact nonlinearly with agents' continuation payoffs in the (utilitarian) policymaker's objective and constraints. This payoff-relevant nonlinearity in our class of games poses a fresh computational challenge. We show that a practical way to tackle this is by exploiting the game's structure, which gives rise to an approximate bilinear programming (BLP) formulation for the optimization problems characterizing SSE.[1] This provides a means to computing an approximate SSE payoff correspondence that preserves upper hemicontinuity of the true correspondence. We provide proof-of-concept evidence, in the form of small-scale numerical examples that can be solved on a desktop multi-core computer, and which should be scalable to massively distributed computing facilities.

In addition, we also characterize equilibrium under a class of fixed policies, showing that each of these arguably simple and operational fixed policies induces a unique corresponding long-run equilibrium, among which a socially optimal equilibrium is shown to exist. We show in a numerical example, that it is possible that some of these optimal long-run simple policies can attain a long run (welfare) outcome that can be sustainable as a particular SSE—i.e., one beginning from a game state consistent with the desired long run outcome. In contrast, an assumed variable commitment policy (also known as Ramsey policy) plan is shown not to be an SSE at all.

☆ We are grateful to Árpád Ábrahám, Gaurab Aryal, Richard Brent, Begoña Domínguez, Zhigang Feng, Kenneth Judd, Tom Krebs, Felix Kübler, Paul Leopardi, Jonathan Ligon, Andy McLennan, Nicola Pavoni, Christopher Sleet and Tim Worrall for suggestions and comments. We also thank participants at the 2011 SAET Conference in Faro, Australasian Workshop on Macroeconomic Dynamics at the University of Queensland, and seminar participants at the Universities of Melbourne and Sydney, Monash University, and at the ANU Mathematical Sciences Institute (Computational Group). This paper previously circulated as "Worker Heterogeneity and Limited-commitment Public Insurance".

* Corresponding author.
E-mail addresses: tcy.kam@gmail.com (T. Kam), ronald.stauber@anu.edu.au (R. Stauber).

1 Standard linear programming problems involve the maximization of a linear objective function subject to linear constraints. A canonical bilinear programming problem has the form $\max_{x,y}\{c^T x + x^T Qy + d^T y\}$ subject to $Ax \leq a$ and $By \leq b$, where x, y, c, d, a and b are Euclidean vectors, and Q, A and B are matrices. Hence, such bilinear problems are linear in x for fixed y, and linear in y for fixed x.

http://dx.doi.org/10.1016/j.jmateco.2016.03.004
0304-4068/

Theoretical extension and new computational methods. From a technical point of view, our results are based on the characterization of equilibrium payoff sets that was derived for repeated games by Abreu et al. (1990), and was extended to dynamic games with public state variables by Atkeson (1991), and with private state variables by Phelan and Stacchetti (2001). The general setup we use draws mostly on Phelan and Stacchetti (2001), although in our model all the relevant state variables are public. However, we extend the framework of Phelan and Stacchetti (2001) by considering a game where the government as a large player faces a non-degenerate distribution over heterogeneous agents, who are small in the sense that they have no individual strategic impact on the game. Our main contribution is in constructing a corresponding heterogeneous-agent dynamic game model that can be analyzed in a tractable way using standard theoretical tools and techniques, and for which a computational approximation can be implemented. Furthermore, for the case where the government is able to commit to a fixed transfer policy, we show that this model allows a simple and novel characterization of existence and optimality of resulting steady state outcomes. The key to our approach is in using a finite-dimensional variable that captures the endogenously evolving distribution of private agent heterogeneity (drawn from a finite set) as a state variable that enters the government's objective. This is in contrast to the government in Phelan and Stacchetti (2001), who is faced with a continuum of agents with identical characteristics, so their game essentially reduces to a game between the government and a representative agent. To make our model tractable – also with a view towards the computational implementation – we simplify one feature relative to the setting of Phelan and Stacchetti (2001), by not allowing private agents to individually accumulate capital. Since private agents are not homogeneous, the level of capital accumulation would be determined by their individual histories of actions and personal states, and not just by their current personal state, which would result in a need to represent combinations of the personal states and capital states using very complex, infinite-dimensional state variables.[2]

Applying the methods arising from Abreu et al. (1990) and the related literature to our class of dynamic games, we show that the set of symmetric sequential equilibria (SSE) can be described by an equilibrium payoff correspondence. This correspondence maps the set of distributions over the private agents' personal states – i.e., the set of game states – to payoff vectors summarizing the equilibrium payoffs derived by the government and the private agents. The SSE payoff correspondence can then be characterized as the largest fixed point of a correspondence-valued operator, and can be computed by recursively applying this operator. Computationally, we implement this operator based on the methods suggested by Sleet and Yeltekin (2000) and Judd et al. (2003). However, we show that in our setting this operator can be constructed using approximate bilinear programs (BLP). Our novel algorithm tractably extends standard linear programming (LP) based computational methods for dynamic games (e.g., Feng, 2015).[3]

The method of Feng (2015) discretizes the domain of an upper hemicontinuous SSE value correspondence of interest and utilizes LP formulations (of SSE optimization problems) to construct approximating outer and inner step correspondences for discrete slices of the true payoff correspondence.[4] However, Sleet and Yeltekin (2000) had pointed out that it is no longer clear what one means by an outer- or an inner-approximating step correspondence in this case when the domain is discretized. This is because at each discretized state, the maximal and minimal levels of the "steps" only apply to that state, and, not to a continuous neighborhood of that state, and therefore, may not be outer- nor inner-bounding levels for nearby payoff sets. What is different in our limited-commitment policy problem, is that the policymaker's objective (in terms of calculating worst punishment values) and incentive constraints will involve bilinear interaction terms between a game state (distribution) and a vector of continuation values for the small players.[5] We exploit the game's structure, which gives rise to an approximate bilinear programming formulation (BLP) for the optimization subproblems characterizing SSE.[6] This provides a means to computing an approximate SSE payoff correspondence that preserves upper hemicontinuity of the true correspondence. Thus, our proposed BLP approach does not run into the problem pointed out by Sleet and Yeltekin (2000).

An illustrative heterogenous-agent model. For concreteness, our main analysis is framed in the context of an unemployment-insurance game between a government and a continuum of private agents. However, our results can easily be applied to a wide range of dynamic public insurance games. A corresponding generalization of our model is discussed in Section 2.4. In the unemployment-insurance interpretation of our game, the private agents ("workers") are distinguished by their heterogenous states of individual employment or unemployment durations. Each worker is subject to repeated stochastic transitions between unemployment and employment, with transition probabilities determined jointly by costly and unobservable job search (or job retention) effort levels, and by the duration of the latest unemployment or employment spell. Since agents only receive wages when employed, the objective of the public unemployment insurance scheme is to provide a combination of wage taxes and unemployment benefits that maximizes the total social welfare by insuring the agents against income fluctuations, while at the same time providing incentives for the agents to exert a socially optimal level of job search or job retention effort. We define an agent's "personal state" as the length of his latest unemployment or employment spell. Agents' personal states then evolve according to

[2] A previous version of this paper included government savings as an additional state variable. However, the resulting model does not significantly alter the analysis, while requiring additional notation.

[3] In turn, Feng (2015) extends the seminal work of Judd et al. (2003), who compute equilibrium payoff sets for repeated games. Feng et al. (2014) also use a similar method to fully describe and solve for incomplete-markets recursive equilibrium value correspondences.

[4] Feng (2015) and Feng et al. (2014) define "step correspondences" in a different way from us and Sleet and Yeltekin (2000). For every slice of the true correspondence at each particular discrete state, they find a union of hypercubes that form the largest possible set contained within the true set (their inner approximation), or, a smallest union of hypercubes that cover the true set (outer approximation). The finite collection of these hypercube sets over all the discretized states are defined to be their respective inner- and outer-approximating "step correspondences".

[5] The game state space is a probability simplex. This is in contrast to, for example, an interval, a square, or finite collections of intervals in known dynamic policy games (e.g., Phelan and Stacchetti, 2001; Feng, 2015). Moreover, in these well-known neoclassical growth model settings, the game state does not additionally interact in a nonlinear way with the representative private agent's payoffs which are contingent on those states.

[6] BLP as a special class of nonlinear programs is well studied and there are known results for the existence of an "ϵ-global" optimum, and for algorithms to find it (see, e.g., Horst and Tuy, 1996; Bennett and Mangasarian, 1993; McCormick, 1976). BLP is also widely used in the fields of operations research, chemical engineering and computational image processing (see, e.g., Nahapetyan and Pardalos, 2008; Chandraker and Kriegman, 2008). More recently, BLP as a formulation to approximate large-scale Markov decision problems (MDP) have also been proposed (see Petrik and Zilberstein, 2011). Our contribution here is related to Petrik and Zilberstein (2011) where the operator defining our equilibrium payoff correspondence has a similar spirit to the single decision maker's Bellman operator in the MDP.