Journal of Mathematical Economics 63 (2016) 139–146

Contents lists available at ScienceDirect

Journal of Mathematical Economics

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

Bayesian Nash equilibrium and variational inequalities[☆]

Takashi Ui

Hitotsubashi University, Japan

HIGHLIGHTS

- A Bayesian Nash equilibrium is a solution of a variational inequality.
- A class of Bayesian games with unique equilibria is studied.
- A special case is a potential game with a strictly concave potential function.
- Applications to aggregative games and network games are discussed.

ARTICLE INFO

Article history:
Received 9 October 2015
Received in revised form
29 December 2015
Accepted 10 February 2016
Available online 18 February 2016

Keywords:
Bayesian game
Linear quadratic Gaussian game
Network game
Potential game
Variational inequality
Strict monotonicity

ABSTRACT

This paper provides a sufficient condition for the existence and uniqueness of a Bayesian Nash equilibrium by regarding it as a solution of a variational inequality. The payoff gradient of a game is defined as a vector whose component is a partial derivative of each player's payoff function with respect to the player's own action. If the Jacobian matrix of the payoff gradient is negative definite for each state, then a Bayesian Nash equilibrium is unique. This result unifies and generalizes the uniqueness of an equilibrium in a complete information game by Rosen (1965) and that in a team by Radner (1962). In a Bayesian game played on a network, the Jacobian matrix of the payoff gradient coincides with the weighted adjacency matrix of the underlying graph.

1. Introduction

This paper explores a sufficient condition for the existence and uniqueness of a Bayesian Nash equilibrium in a class of Bayesian games where action sets are closed intervals and each player's payoff function is concave and continuously differentiable with respect to the player's own action. This class of Bayesian games has many applications such as Cournot and Bertrand competition, private provision of public goods, rent seeking, and strategic interaction on networks. A special case of our sufficient condition includes strict concavity of a potential function in Bayesian potential games (Radner, 1962; Ui, 2009).

We formulate a Bayesian Nash equilibrium as a solution of a variational inequality in an infinite-dimensional space (Kinderlehrer and Stampacchia, 1980), which is one representation of the first-order condition for an equilibrium. This representation not only gives us an elementary proof for the uniqueness but also allows us to use the existence theorem for solutions of variational inequalities (Browder, 1965; Hartman and Stampacchia, 1966). It is well known that a Nash equilibrium of a complete information game is a solution of a variational inequality in a finite-dimensional space (Lions and Stampacchia, 1967; Bensoussan, 1974). Thus, it is hardly surprising that a Bayesian Nash equilibrium is a solution of a variational inequality in an infinite-dimensional space. To the best of the author's knowledge, however, the resulting implications are not necessarily well-documented. This paper fills this gap in the literature and shed new light on the variational inequality approach to game theory.

In the main results, we construct a vector whose component is a partial derivative of each player's payoff function with respect to the player's own action. This vector is referred to as the payoff gradient of the game. The payoff gradient is said to be strictly monotone if its Jacobian matrix is negative definite for each state.[1]

☆ An earlier draft of this paper was circulated under the same title in 2004. I am very grateful for valuable input from Yasunori Kimura. I thank the editor and an anonymous referee for detailed comments and suggestions. I also benefited from comments of Kota Murayama and seminar participants at Kyoto University. I acknowledge financial support by Grant-in-Aid for Scientific Research (grant numbers 16530114, 15K03348, 26245024).

E-mail address: oui@econ.hit-u.ac.jp.

[1] To be more precise, negative definiteness of the Jacobian matrix is a sufficient condition for strict monotonicity of the payoff gradient.

http://dx.doi.org/10.1016/j.jmateco.2016.02.004
0304-4068/

It is said to be strongly monotone if it is strictly monotone and the maximum eigenvalue of the Jacobian matrix has a strictly negative supremum over the actions and the states. We show that if the payoff gradient is strictly monotone, then there exists at most one equilibrium, and if the payoff gradient is strongly monotone or if it is strictly monotone and the payoff functions are quadratic, then there exists a unique equilibrium. In particular, we consider a linear quadratic Gaussian (LQG) game, whose payoff functions are quadratic and private signals are normally distributed, and obtain the unique equilibrium in a closed form, which is linear in private signals.

Our condition is an extension of the sufficient condition for the uniqueness of a Nash equilibrium by Rosen (1965), who shows that a Nash equilibrium is unique if the payoff gradient of a complete information game is strictly monotone. As shown by Ui (2008), the unique Nash equilibrium is also a unique correlated equilibrium. We can show the uniqueness of a correlated equilibrium as a special case of our results because a Bayesian game is reduced to a complete information game with a correlation device when payoff functions are independent of the state.

Our condition is also an extension of the sufficient condition for the uniqueness of a Bayesian Nash equilibrium by Radner (1962). Radner (1962) studies a team, an identical interest Bayesian game with a common payoff function,[2] and shows that if the common payoff function is strictly concave in an action profile, a Bayesian Nash equilibrium is a unique maximizer of the expected value of the common payoff function. As a special case, Radner (1962) considers an LQG team and obtains the unique equilibrium in a closed form. Radner's results are used to study Bayesian potential games (Monderer and Shapley, 1996; van Heumen et al., 1996). A Bayesian potential game has the same best-response correspondence as that of a team,[3] the common payoff function of which is referred to as a potential function. If the potential function is strictly concave, a Bayesian Nash equilibrium is a unique maximizer of the expected value of the potential function, as shown by Ui (2009).

Our results generalize Radner's results and the applications to Bayesian potential games in the following sense. A Bayesian game is a Bayesian potential game if and only if the Jacobian matrix of the payoff gradient is symmetric (Monderer and Shapley, 1996), in which case the Jacobian matrix coincides with the Hessian matrix of a potential function. Moreover, the potential function is strictly concave if and only if the payoff gradient is strictly monotone (Ui, 2008). Thus, we can restate the results of Radner (1962) and Ui (2009) as follows: a Bayesian Nash equilibrium is unique if the Jacobian matrix is both symmetric and negative definite. Our results show that the symmetry requirement is not necessary.

For example, most studies on LQG games assume that the Jacobian matrix is symmetric and negative definite, i.e., an LQG game is a Bayesian potential game with a strictly concave potential function.[4] In order to analyze communication in a network, however, Calvó-Armengol et al. (2015) consider an LQG game in which the Jacobian matrix is asymmetric and show the existence and uniqueness of a linear Bayesian Nash equilibrium, while it has been an open question under what condition the linear equilibrium is a unique equilibrium. Our results show that this linear equilibrium is a unique equilibrium if the Jacobian matrix is negative definite.

As an application, we consider aggregative games (Selten, 1970), in which each player's payoff depends on the player's own action and the aggregate of all players' actions. We give a simple sufficient condition for the uniqueness of a Bayesian Nash equilibrium and apply it to a Cournot game and a rent-seeking game. We also consider games played on networks (Ballester et al., 2006; Bramoullé et al., 2014), or network games for short.[5] A Bayesian game with quadratic payoff functions is mathematically equivalent to a Bayesian network game, where the Jacobian matrix of the payoff gradient equals the negative of a weighted adjacency matrix of the underlying graph. Thus, a Bayesian network game has a unique equilibrium if the weighted adjacency matrix is positive definite. We can use this result to study Bayesian network games with random adjacency matrices, whereas most previous studies on Bayesian network games assume a constant adjacency matrix with a special structure (Blume et al., 2015; de Martì and Zenou, 2015; Calvó-Armengol et al., 2015).

The organization of the paper is as follows. Preliminary definitions and results are summarized in Section 2. Section 3 discusses the concept of strictly monotone payoff gradients. Section 4 reports the main results. Section 5 is devoted to applications.

2. Preliminaries

Consider a Bayesian game with a set of players $N = \{1, \ldots, n\}$. Player $i \in N$ has a set of actions $X_i \subseteq \mathbb{R}$, which is a closed interval. We write $X = \prod_{i\in N} X_i$ and $X_{-i} = \prod_{j\neq i} X_j$. Player i's payoff function is a measurable function $u_i : X \times \Omega \to \mathbb{R}$, where $(\Omega, \mathcal{F}, P)$ is a probability space. Player i's information is given by a measurable mapping $\eta_i : \Omega \to Y_i$, where $(Y_i, \mathcal{Y}_i)$ is a measurable space. Player i's strategy is a measurable mapping $\sigma_i : Y_i \to X_i$ with $E[\sigma_i(\eta_i)^2] < \infty$. We regard two strategies σ_i^1, σ_i^2 as the same strategy if $\sigma_i^1(\eta_i(\omega)) = \sigma_i^2(\eta_i(\omega))$ almost everywhere. Let Σ_i denote player i's set of strategies. We write $\Sigma = \prod_{i\in N} \Sigma_i$ and $\Sigma_{-i} = \prod_{j\neq i} \Sigma_j$. We assume that $E[u_i(\sigma, \omega)]$ exists for all $\sigma \in \Sigma$.

We fix N, X, and $(\Omega, \mathcal{F}, P)$ throughout this paper and simply denote a Bayesian game by $(\mathbf{u}, \eta)$, where $\mathbf{u} = (u_i)_{i\in N}$ and $\eta = (\eta_i)_{i\in N}$. We say that $(\mathbf{u}, \eta)$ is *smooth* if $u_i((\cdot, x_{-i}), \omega) : X_i \to \mathbb{R}$ is continuously differentiable for each $x_{-i} \in X_{-i}$, $i \in N$, and a.e. $\omega \in \Omega$, and $E[(\partial u_i(\sigma, \omega)/\partial x_i)^2] < \infty$ for each $\sigma \in \Sigma$ and $i \in N$. We write $\nabla \mathbf{u}(x, \omega) \equiv (\partial u_i(x, \omega)/\partial x_i)_{i\in N}$ and call it the payoff gradient of $\mathbf{u}$. We say that $(\mathbf{u}, \eta)$ is *concave* if $u_i((\cdot, x_{-i}), \omega) : X_i \to \mathbb{R}$ is concave for each $x_{-i} \in X_{-i}$, $i \in N$, and a.e. $\omega \in \Omega$.

A strategy profile $\sigma \in \Sigma$ is a Bayesian Nash equilibrium if, for a.e. $\omega \in \Omega$,

$$E[u_i(\sigma(\eta), \omega) \mid \eta_i] \geq E[u_i((x_i, \sigma_{-i}(\eta_{-i})), \omega) \mid \eta_i] \quad (1)$$

for each $x_i \in X_i$ and $i \in N$, where $\sigma(\eta) = (\sigma_i(\eta_i))_{i\in N}$, $\sigma_{-i}(\eta_{-i}) = (\sigma_j(\eta_j))_{j\neq i}$, and $E[\cdot \mid \eta_i]$ is the conditional expectation operator given $\eta_i(\omega)$.

In this paper, we study an equilibrium of a smooth concave Bayesian game by converting the first-order condition as follows.[6] First, we exchange the order of integration and differentiation to obtain the following representation.[7]

[2] The theory of teams precedes Harsanyi (1967–1968).

[3] Ui (2009) studies a game satisfying this condition and calls it a best-response Bayesian potential game.

[4] Basar and Ho (1974) was the first to use Radner's results to study LQG games that are not teams, followed by many studies on information sharing (Clark, 1983; Vives, 1984; Gal-Or, 1985), information acquisition (Li et al., 1987; Vives, 1988), and social value of information (Morris and Shin, 2002; Angeletos and Pavan, 2007; Ui and Yoshizawa, 2015), among others. LQG games in these studies are Bayesian potential games.

[5] See Jackson and Zenou (2015) for a survey.

[6] Even if a concave Bayesian game is not smooth, we can obtain a similar first-order condition in terms of subderivatives and extend our main results using multi-valued variational inequalities, which is beyond the scope of this paper.

[7] Of course, the same first-order condition is valid under a suitable condition without concavity.