Journal of Mathematical Economics 64 (2016) 113–123

Contents lists available at ScienceDirect

Journal of Mathematical Economics

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

Quantal response methods for equilibrium selection in normal form games

Boyu Zhang

Laboratory of Mathematics and Complex System, Ministry of Education, School of Mathematical Sciences, Beijing Normal University, Beijing, PR China

ARTICLE INFO

Article history:
Received 11 September 2015
Received in revised form
14 April 2016
Accepted 18 April 2016
Available online 23 April 2016

Keywords:
Quantal response equilibrium
Equilibrium selection
Normal form game
Symmetric game
Dominated strategy

ABSTRACT

This paper describes a general framework for equilibrium selection by tracing the graph of the quantal response equilibrium (QRE) correspondence as a function of the variance of random disturbances. If a quantal response function satisfies C^2 continuity, monotonicity and cumulativity, the graph of QRE correspondence generically includes a unique branch that starts at the centroid of the strategy simplex and converges to a unique Nash equilibrium as noises vanish. This equilibrium is called the limiting QRE of the game. We then investigate the limiting QRE in normal form games, and analyze the effects of payoff transformations and adding/eliminating dominated strategies on equilibrium selection. We find that in two-person symmetric games, any strict Nash equilibrium can be selected as the limiting QRE by appropriately adding a single strictly dominated strategy.

1. Introduction

Classical game theory relies on the assumption of perfect rationality. However, people in the real world may make mistakes and their behaviors sometimes deviate from Nash equilibrium strategies. Probabilistic choice models have been developed to incorporate stochastic errors to the analysis of individual decisions (e.g., Luce, 1959; McFadden, 1976). One of the well known model is the quantal response equilibrium (QRE) introduced by McKelvey and Palfrey (1995). In a QRE, players make decisions based on quantal response function (also called perturbed best response function) and believe other players do so as well. A general interpretation of this model is that they observe randomly perturbed payoffs of strategies and choose optimally according to those noisy observations, and a QRE is defined as a fixed point of this process (McKelvey and Palfrey, 1995, 1998; Goeree et al., 2005; Turocy, 2005; Sandholm, 2010).[1]

Formally, a quantal response function maps the vector of expected payoffs into a vector of choice probabilities. Goeree et al. (2005) proposed a "reduced form" definition of QRE. Rather than restricting payoff disturbances explicitly, they define a regular QRE by restricting quantal response functions to satisfy four axioms: continuity, interiority, responsiveness, and monotonicity. The reduced form approach does not require that quantal response functions are derived from some underlying choice models of stochastic utility maximization, therefore allows for a richer set of models for data estimation.

The most common specification of QRE is the logit equilibrium, where the random disturbances on the payoffs follow the extreme value distribution (Blume, 1993, 1995; McKelvey and Palfrey, 1995, 1998; Anderson et al., 2004; Turocy, 2005; Hofbauer and Sandholm, 2002, 2007; Sandholm, 2010). The logit response function has one free parameter λ, whose inverse $1/\lambda$ has been interpreted as the temperature, or the intensity of noise. The set of logit equilibria can be viewed as a correspondence from λ to the set of mixed strategy profiles. At $\lambda = 0$, each strategy is chosen with equal probability and the correspondence contains only the centroid of the strategy simplex. As λ approaches infinity, players choose the best responses and the correspondence converges to a subset of the Nash equilibria.

McKelvey and Palfrey (1995) pointed out that the graph of logit equilibria correspondence generically includes a unique branch that starts at the centroid of the strategy simplex (the only QRE when $\lambda = 0$) and converges to a unique Nash equilibrium as $\lambda \to +\infty$. They then suggested an equilibrium selection by tracing this branch, and called the selected Nash equilibrium the *limiting logit equilibrium* (LLE) of the game. An economics interpretation is that

E-mail address: zhangby@bnu.edu.cn.

[1] Ui (2002) provided an evolutionary interpretation for QRE. In a population game, if a stochastic best response process satisfies the detailed balance condition, then the support of the stationary distribution converges to the set of QRE as the population size goes to infinity. Later, Ui (2006) indicated that the QRE model is also equivalent to an incomplete information game where the actual payoff is the sum of payoffs of some fixed game and independent random terms, and each player's private signal is her own payoffs. A QRE is a probability distribution of action profiles in a Bayesian Nash equilibrium.

http://dx.doi.org/10.1016/j.jmateco.2016.04.003
0304-4068/

as players gain experience from repeated observations, they can be expected to make more precise estimates of the expected payoffs of different strategies. McKelvey and Palfrey (1995) showed that in several repeated game experiments, estimates of λ indeed increased as the game progresses. This then provides an empirical evidence of the equilibrium selection above.[2]

In this paper, we describe a general framework for equilibrium selection by tracing the graph of the QRE correspondence. Following the logit equilibrium, define a *multiplicative QRE* at noise level λ as a fixed point of quantal response functions where payoffs of strategies are multiplied by the factor λ. In order to define the equilibrium selection, some assumptions on the QRE correspondence are needed. First, there is a unique QRE for small λ, so we can trace the branch of the QRE correspondence starting from this QRE. Second, this branch does not include bifurcation point. Third, this branch reaches a Nash equilibrium as λ goes to infinity. Thus, we impose three restrictions on the quantal response functions: continuity, monotonicity and cumulativity. Continuity is a technical property. C^0 continuity is enough to guarantee the existence of a QRE for any given λ, and C^2 continuity implies the uniqueness of the QRE for small λ. Monotonicity is a weak form of rational choice, meaning that strategies with higher payoffs are used more frequently. As a result of monotonicity and continuity, the only QRE at $\lambda = 0$ is the centroid. Cumulativity ensures that players choose best responses for sufficiently large λ. Together with continuity, the QRE correspondence converges to Nash equilibria as λ goes to infinity.

Intuitively, quantal response functions that satisfy the three axioms are smooth generalizations of best response functions. With the three axioms, the graph of the QRE correspondence includes a path that connects the centroid at $\lambda = 0$ to at least one Nash equilibrium. However, C^0 continuity is too weak that the path is not necessarily nicely behaved. In exceptional cases, the path may not be differentiable and bifurcation may arise. Such exceptional cases can be generically excluded by making differentiability assumption. If the quantal response function is C^2 continuous, except for a nowhere dense set of games, the path is diffeomorphic to a C^1 segment. This implies that for almost all normal form games, there is a unique selection from the set of Nash equilibria by "tracing" the graph of the QRE correspondence beginning at the centroid. We call the selected Nash equilibrium the *limiting QRE* of the game.

We then study the properties of the limiting QRE. In normal form games, we show that quantal response method always selects the *majority dominant* (MD) equilibrium (if it exists), where a Nash equilibrium is called MD if it is the unique best response when everyone believes that other players use their component of the MD equilibrium most frequently. This condition is then applied to calculate the limiting QRE in two symmetric n-person coordination games.

It has been shown that the limiting QRE is highly sensitive to linear payoff transformations (Tumennasan, 2013; Zhang and Hofbauer, 2016). Based on an n-person pure coordination game, we find that the influence of a player on the equilibrium selection is positively correlated with the multiplication factor on her payoff. If the factor is large enough, the equilibrium her preferred most will be selected. On the other hand, Goeree and Holt (2001, 2004) observed that the limiting QRE is subject to framing effects in the sense that duplicating a strategy may change the limiting QRE. In this paper, we offer a much stronger proposition for two-person symmetric games: any strict (symmetric) Nash equilibrium can be selected by appropriately adding a single strictly dominated strategy. Therefore, the limiting QRE is also sensitive to the addition and elimination of strictly dominated strategies.

The rest of this paper is organized as follows. Section 2 defines multiplicative QRE and introduces some properties. Section 3 studies the topological structure of the graph of the QRE correspondence in normal form games and symmetric games. Section 4 investigates the limiting QRE in normal form games, and analyzes the effects of payoff transformations and adding/eliminating dominated strategies on equilibrium selection. Section 5 discusses the main results and suggests for further developments.

2. Properties of the quantal response equilibrium

2.1. Notations

The notations in this paper follow that of McKelvey and Palfrey (1995). Consider an n-person normal form game $\Gamma = (N, S, u)$, where $N = \{1, \dots, n\}$ is the set of *players*. For each player $i \in N$, there is a *strategy set* $S_i = \{s_{i1}, \dots, s_{iJ_i}\}$ consisting of J_i pure strategies and a *payoff function*, $u_i : S \to \mathbb{R}$, where $S = \prod_{i\in N} S_i$ is the set of strategy profiles.

Let Δ_i be the set of probability distributions on S_i. Elements of Δ_i are of the form $p_i : S_i \to \mathbb{R}$, where $\sum_{s_{ij}\in S_i} p_i(s_{ij}) = 1$ and $p_i(s_{ij}) \geq 0$ for all $s_{ij} \in S_i$. For convenience, use the notation $p_{ij} = p_i(s_{ij})$ to denote the probability of player i using strategy j. We write the set of mixed strategy profiles by $\Delta = \prod_{i\in N} \Delta_i$ and denote points in Δ by $p = (p_1, \dots, p_n)$. Therefore, given a mixed strategy profile p, player i's expected payoff is $u_i(p) = \sum_{s\in S} p(s)u_i(s)$, where $p(s) = \prod_{i\in N} p_i(s_i)$, where $s_i \in S_i$ denotes the ith element of s. For convenience, for each $i \in N$ and $j \in \{1, \dots, J_i\}$, denote by $u_{ij}(p)$ the expected payoff to player i adopting pure strategy s_{ij} when the other players adopt their components of p. To get an expression of $u_{ij}(p)$, we use the notation (s_{ij}, s_{-i}) to represent the pure strategy profile that player i adopts the strategy s_{ij} and all other players adopt their component of s_{-i}, where $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n) \in \prod_{j\neq i} S_j = S_{-i}$. $u_{ij}(p)$ could be then written as

$$u_{ij}(p) = \sum_{s_{-i}\in S_{-i}} u_i(s_{ij}, s_{-i}) \prod_{t\neq i} p_t(s_t). \tag{1}$$

The space of payoff vectors of player i's pure strategies is $\mathbb{R}^{J_i}$, and write $\mathbb{R}^{\sum J_i} = \prod_{i\in N} \mathbb{R}^{J_i}$. Define $\bar{u} : \Delta \to \mathbb{R}^{\sum J_i}$ by $\bar{u}(p) = (\bar{u}_1(p), \dots, \bar{u}_n(p))$, where $\bar{u}_i(p) = (u_{i1}(p), \dots, u_{iJ_i}(p))$ shows the expected payoffs of player i's J_i pure strategies.

McKelvey and Palfrey (1995) assumed that for each pure strategy s_{ij}, there is an additional payoff disturbance ε_{ij}, and players make decisions based on perturbed observations. Denote the noisy payoff by

$$\tilde{u}_{ij}(p) = u_{ij}(p) + \varepsilon_{ij}. \tag{2}$$

Player i's noise vector, $\varepsilon_i = (\varepsilon_{i1}, \dots, \varepsilon_{iJ_i})$, is distributed according to a joint distribution with density function $f_i(\varepsilon_i)$. $f = (f_1, \dots, f_n)$ is called *admissible* (McKelvey and Palfrey, 1995; Goeree et al., 2005) if

(a) the marginal distribution of f_i exists for each ε_{ij},
(b) disturbances are independent across players (not necessarily across strategies),
(c) $E(\varepsilon_i) = 0$ for all $i \in N$.

Define $B_{ij}(\bar{u}_i)$ to be the set of ε_i such that strategy s_{ij} has the highest disturbed payoff, i.e.,

$$B_{ij}(\bar{u}_i) = \{\varepsilon_i \in \mathbb{R}^{J_i} | u_{ij} + \varepsilon_{ij} \geq u_{ik} + \varepsilon_{ik}, \forall k = 1, \dots, J_i\}. \tag{3}$$

[2] There are many other probabilistic choice models for equilibrium selection. A class of models is based on dynamic approach, which uses learning or evolutionary dynamics to predict equilibrium. Examples include long-run equilibrium (Kandori et al., 1993; Young, 1993) and stochastic fictitious play (Hofbauer and Sandholm, 2002).