Journal of Theoretical Biology 346 (2014) 23–33

Contents lists available at ScienceDirect

Journal of Theoretical Biology

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

Gains from switching and evolutionary stability in multi-player matrix games

Jorge Peña [a,*], Laurent Lehmann [b], Georg Nöldeke [a]

[a] Faculty of Business and Economics, University of Basel, Peter Merian-Weg 6, CH-4002 Basel, Switzerland
[b] Department of Ecology and Evolution, University of Lausanne, Le Biophore, CH-1015 Lausanne, Switzerland

HIGHLIGHTS

- We study the evolutionary dynamics of two-strategy symmetric multi-player matrix games.
- We make use of the theory of polynomials in Bernstein form.
- We unify, simplify and extend previous work on evolutionary multi-player games.

ARTICLE INFO

Article history:
Received 4 October 2013
Received in revised form
12 December 2013
Accepted 16 December 2013
Available online 28 December 2013

Keywords:
Evolutionary game theory
Replicator dynamics
Polynomials in Bernstein form
Public goods games

ABSTRACT

In this paper we unify, simplify, and extend previous work on the evolutionary dynamics of symmetric N-player matrix games with two pure strategies. In such games, gains from switching strategies depend, in general, on how many other individuals in the group play a given strategy. As a consequence, the gain function determining the gradient of selection can be a polynomial of degree $N-1$. In order to deal with the intricacy of the resulting evolutionary dynamics, we make use of the theory of polynomials in Bernstein form. This theory implies a tight link between the sign pattern of the gains from switching on the one hand and the number and stability of the rest points of the replicator dynamics on the other hand. While this relationship is a general one, it is most informative if gains from switching have at most two sign changes, as is the case for most multi-player matrix games considered in the literature. We demonstrate that previous results for public goods games are easily recovered and extended using this observation. Further examples illustrate how focusing on the sign pattern of the gains from switching obviates the need for a more involved analysis.

1. Introduction

Game theory has been widely applied to evolutionary biology (Maynard Smith and Price, 1973; Maynard Smith, 1982; Eshel, 1996; Hofbauer and Sigmund, 1998; Rousset, 2004; Vincent and Brown, 2005; Dercole and Rinaldi, 2008; Broom and Rychtář, 2013). More specifically, the application of game-theoretic concepts has been instrumental in explaining the evolution of traits as diverse as the sex ratio (Hamilton, 1967; Frank, 1987), dispersal (Hamilton and May, 1977; Comins et al., 1980), reciprocity (Axelrod and Hamilton, 1981), group foraging (Clark and Mangel, 1986), policing (Frank, 1995), and anisogamy (Bulmer and Parker, 2002). Evolutionary models of these traits often assume "playing the field" type of interactions (Maynard Smith, 1982, p. 23), where the payoff to an individual depends on an average property of the population or the group with which it interacts.

There are many situations, however, where the payoff to an individual depends critically on the strategy profile in the population (or its group) and where the actions of different individuals cannot be averaged; that is, mass action does not apply. Typical examples involve collective action problems in moderately sized groups, where the change in behavior by a single individual can result in a large, discontinuous change in payoffs to others (e.g., Boyd and Richerson, 1988). Such collective action problems have been modeled as multi-player (or multi-person) matrix games (Broom et al., 1997; Kurokawa and Ihara, 2009; Gokhale and Traulsen, 2010). Except for the very special cases in which group size is taken to be equal to two (so that the well-developed theory of two-player matrix games can be applied, cf. Weibull, 1995; Hofbauer and Sigmund, 1998; Cressman, 2003) or the payoff structure is linear (as in the standard model of the N-person prisoner's dilemma), such games have proven to be difficult to analyze.

* Correspondence to: Research Group for Evolutionary Theory, Max Planck Institute for Evolutionary Biology, August-Thienemann-Str. 2, 24306 Plön, Germany.
E-mail addresses: pena@evolbio.mpg.de (J. Peña), laurent.lehmann@unil.ch (L. Lehmann), georg.noeldeke@unibas.ch (G. Nöldeke).

0022-5193/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.jtbi.2013.12.016

The intrinsic complexity of multi-player matrix games is already evident for the case of symmetric games with two pure strategies A and B on which we focus in this paper. For these games, the average payoff difference in a large and well-mixed population is given by the so-called *gain function* (Bach et al., 2006)

$$g(x) = \sum_{k=0}^{n} \binom{n}{k} x^k (1-x)^{n-k} d_k.$$

Here, n is the number of co-players of a focal player (so that $N = n+1$ is the group size), x is the population fraction of A-strategists, and d_k is the gain a focal player would obtain if switching from strategy B to strategy A when k other group-members play A. The evolutionary solution of the game (such as the set of evolutionarily stable strategies, ESSs, or the set of stable rest points of the replicator dynamics) involves not only finding the roots of the gain function $g(x)$ (a polynomial of degree n) but also, as discussed in Broom et al. (1997), determining the behavior of $g(x)$ in the vicinity of such roots. While this is straightforward for two-player games (for which $g(x)$ is linear in x) and a full classification for three-player games (for which $g(x)$ is quadratic in x) is available (Bukowski and Miękisz, 2004), payoff structures in groups of size larger than five lead to polynomials of degree greater than four that cannot, in general, be solved analytically (Clark, 1984).

In order to deal with such complexity, the vast majority of previous works on multi-player matrix games has considered particular functional forms for the specification of the payoffs and has resorted to lengthy algebra or numerical methods to study the models (Joshi, 1987; Boyd and Richerson, 1988; Dugatkin, 1990; Weesie and Franzen, 1998; Hauert et al., 2006; Zheng et al., 2007; Cuesta et al., 2008; Pacheco et al., 2009; Archetti, 2009; Souza et al., 2009; Archetti and Scheuring, 2011; van Segbroeck et al., 2012). In this way, some non-linear public goods games, including multi-player extensions of well-known two-person matrix games such as the stag hunt (Skyrms, 2004) and the snowdrift game (Sugden, 1986), have been characterized on a case-by-case basis.

In contrast to these efforts, Motro (1991) and Bach et al. (2006) have taken a more systematic approach to the study of non-linear public goods games. Both of these papers consider situations in which each contributor to a public good pays a constant cost, whereas the benefit from the public good, which is obtained by all players, is a function of the number of contributors. Motro (1991) proves that in this case the replicator dynamics has at most one interior rest point if the benefit is concave or convex in the number of contributors. He also provides necessary and sufficient conditions for the existence of such a rest point and characterizes the stability property of all rest points. In a similar spirit, Bach et al. (2006) find sufficient conditions on the shape of the benefits such that there exists a critical cost level with the property that for costs below such a level the replicator dynamics has two interior rest points, whereas for higher costs there is no interior rest point.

Gokhale and Traulsen (2010) have discussed the relationship between the sign pattern of the gains from switching and the number of interior rest points of the replicator dynamics. Specifically, these authors observe that the replicator dynamics has a single interior rest point if the sequence $(d_0, d_1, \ldots, d_n)$, which we refer to as the *gain sequence*, has exactly one sign change. Gokhale and Traulsen (2010) also note that the direction of selection (as given by the sign of the gain function $g(x)$) cannot have more sign changes than the gain sequence. This implies that the number of sign changes of the gain sequence provides an upper bound on the number of interior rest points of the replicator dynamics. The latter observation is also made in Hauert et al. (2006) and Cuesta et al. (2007). When $g(x)$ has no multiple roots, any upper bound on the number of interior rest points translates directly into an upper bound on the number of stable rest points because, as noted in Broom et al. (1997, p. 939), in this case the rest points alternate between being stable and unstable.

In this paper, we show how sign-change conditions like the ones discussed by Gokhale and Traulsen (2010) can be refined by using the fact that the gain function $g(x)$ is a particular kind of polynomial, known as a polynomial in Bernstein form (or Bernstein polynomial), with coefficients given by the gain sequence $(d_0, d_1, \ldots, d_n)$. Our analysis rests on the variation-diminishing property of Bernstein polynomials and a property that we refer to as the preservation of initial and final signs. These properties provide a tight link between the sign pattern of the gain sequence and the sign pattern of the gain function.[1] In particular, if the gain sequence has at most two sign changes, a full characterization of the possible dynamic regimes is easily obtained.

For most of the collective action problems that have been modeled as multi-player matrix games it is straightforward to determine the sign pattern of the gain sequence. Moreover, because the gain sequences of these games have at most two sign changes, our characterization results provide all the information necessary to recover the results on the number and stability of rest points obtained in previous studies. We demonstrate these claims for two classes of public goods games, namely threshold games (e.g., Dugatkin, 1990; Weesie and Franzen, 1998; Zheng et al., 2007; Souza et al., 2009) and constant cost games (e.g., Motro, 1991; Bach et al., 2006; Hauert et al., 2006; Pacheco et al., 2009; Archetti and Scheuring, 2011), and two additional examples taken from Hauert et al. (2006) and van Segbroeck et al. (2012), thus supporting the claim that the approach developed here unifies, simplifies, and extends much of the previous work on multi-player matrix games.

2. Model

Interactions occur in groups of size $N = n+1$, in which a focal individual plays a game against n co-players or opponents. Each individual can choose between one of two different pure strategies, A and B. The game is symmetric so that, from the focal's point of view, any two co-players are exchangeable.

Let a_k denote the payoff to an individual choosing A when k opponents choose A (and hence $n-k$ co-players choose B); likewise, let b_k denote the payoff to an individual choosing B when k opponents choose A. Also let

$$d_k \equiv a_k - b_k$$

denote the gain the focal player makes from choosing A over B, taking the choices of other players (k playing A and $n-k$ playing B) as given. The parameters d_k, which describe the gains from switching, are collected in the gain sequence $\mathbf{d} = (d_0, d_1, \ldots, d_n)$. We assume $\mathbf{d} \neq \mathbf{0}$, thus excluding the uninteresting case in which payoffs are independent of the actions chosen.

Evolution occurs in an infinitely large and well-mixed population with groups randomly formed by binomial sampling. Hence, if the frequency of A-strategists in the whole population is x, the average payoffs obtained by an A-strategist and a B-strategist are respectively given by

$$\pi_A(x) = \sum_{k=0}^{n} \binom{n}{k} x^k (1-x)^{n-k} a_k$$

[1] The fact that the gain function $g(x)$ is a Bernstein polynomial has previously been noted by Cuesta et al. (2007). These authors also suggest that the variation diminishing property of these polynomials may make the analysis of many multi-player games straightforward, but do not pursue this idea.