Automatica 71 (2016) 50–56

Contents lists available at ScienceDirect

Automatica

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

Brief paper

Incentive equilibrium strategies in dynamic games played over event trees[☆]

Elnaz Kanani Kuchesfehani [a], Georges Zaccour [b,1]

[a] GERAD, HEC Montréal, Montréal (Québec), Canada, H3T 2A7

[b] Chair in Game Theory and Management, GERAD, HEC Montréal, Montréal (Québec), Canada, H3T 2A7

ARTICLE INFO

Article history:
Received 10 February 2015
Received in revised form 15 March 2016
Accepted 4 April 2016
Available online 31 May 2016

Keywords:
Dynamic games
Incentive equilibria
Event tree
Cooperation
Linear-state dynamic games
Linear–quadratic dynamic games

ABSTRACT

We characterize incentive equilibrium strategies and their credibility conditions for the class of linear–quadratic dynamic games played over event trees. In such games, the transition from one node to another is nature's decision and cannot be influenced by the players' actions. We show that it is possible for two players wanting to optimize their joint payoff over a given planning horizon to achieve this as an incentive equilibrium. This therefore ensures that cooperation will continue from one node to the next. A simple example illustrates these strategies and the credibility conditions.

1. Introduction

A main issue in cooperative dynamic games is how to sustain cooperation over time, that is, how to ensure that each player will indeed implement her part of the agreement as time goes by. The breakdown of long-term agreements before their maturity has been empirically observed. Schematically, a breakdown will occur either if all the parties agree at an intermediate instant of time to replace the initial agreement by a new one for the remaining periods, or if one of the players finds it (individually) rational to deviate, that is, to switch to her noncooperative strategy from that time onward (Haurie, 1976). The literature in (state-space) dynamic games suggests mainly three approaches to sustain cooperation over time.

Time consistency: A cooperative agreement is time consistent at an initial date and state if, at any intermediate instant of time, the cooperative payoff-to-go of each player dominates, at least weakly, her noncooperative payoff-to-go; see, e.g., Yeung and Petrosyan (2005). Note that a time-consistent payment schedule always exists, and that the cooperative and noncooperative payoffs-to-go are compared along the cooperative state trajectory, which implicitly assumes that the players have so far played cooperatively. A stronger concept is agreeability, which requires the cooperative payoff-to-go to dominate the noncooperative payoff-to-go along *any* state trajectory; see, e.g., Kaitala and Pohjola (1990). For a survey of time consistency, see Zaccour (2008).

Cooperative equilibrium: If the cooperative solution is an equilibrium, then durability of the agreement is not an issue anymore as it will be in the best interest of each player not to deviate from the agreement. To endow the cooperative solution with an equilibrium property, one approach is to use trigger strategies that credibly and effectively punish any player deviating from the agreement; see, e.g., Dockner, Jorgensen, Long, and Sorger (2000), Haurie and Pohjola (1987) and Tolwinski, Haurie, and Leitmann (1986).

Incentive equilibrium: Trigger strategies may embody large discontinuities, i.e., a slight deviation from an agreed-upon path triggers harsh retaliation, generating a very different path from the agreed-upon one. An alternative approach, which will be followed here, is to use incentive strategies that are continuous

☆ Research supported by NSERC (grant no. 37525-2011) and SSHRC (grant no. 435-2013-0532), Canada. The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Kok Lay Teo under the direction of Editor Ian R. Petersen.

E-mail addresses: elnaz.kanani@gerad.ca (E.K. Kuchesfehani), georges.zaccour@gerad.ca (G. Zaccour).

1 Tel.: +1 514 340 6053 6050.

http://dx.doi.org/10.1016/j.automatica.2016.04.026
0005-1098/

in the information. An incentive equilibrium has the property that when both players implement their incentive strategies, the cooperative outcome is realized as an equilibrium. Therefore, no player should be tempted to deviate from the agreement during the course of the game, provided that the incentive strategies are credible. An incentive strategy is credible if it is better for a player who has been cheated to use her strategy than to stick to the coordinated solution. Ehtamo and Hämäläinen (1989, 1993) used linear incentive strategies in a dynamic resource game and demonstrated that such strategies are credible when deviations are not too large.

The concept of incentive strategies has of course been around for a long time in dynamic games (and economics), but it was often understood and used in a leader–follower (or principal–agent) sense. The idea is that the leader designs an incentive to induce the follower to reply in a certain way, which is often meant to be (only) in the leader's best interest, but may also be in the best collective interest (see the early contributions by Ho (1983) and Basar (1984)). In such a case, the incentive is one sided. Here, we focus on two-sided incentive strategies, with the aim of implementing the joint optimization solution.

The objective of this paper is to characterize incentive equilibrium strategies and outcomes for the class of dynamic games played over event trees (DGET). In these games, the transition from one node to another is nature's decision and cannot be influenced by the players' actions. For a detailed description of DGET, see Haurie, Krawczyk, and Zaccour (2012). We focus on linear–quadratic dynamic games, a popular class in applications because it admits closed-form solutions (see, e.g., the books by Engwerda (2005) and Haurie et al. (2012)). Martín-Herrán and Zaccour (2005, 2009) characterized incentive strategies and their credibility for linear-state and linear–quadratic dynamic games (LQDG), but in a deterministic setting.

The rest of the paper is organized as follows. In Section 2, we briefly recall the ingredients of DGET and derive the coordinated solution. In Section 3, we define the incentive equilibrium strategies; and we provide a numerical illustration in Section 4. We briefly conclude in Section 5.

2. Linear–quadratic DGET

Let $\mathcal{T} = \{0, 1, \ldots, T\}$ be the set periods, and denote by $(\xi(t) : t \in \mathcal{T})$ the exogenous stochastic process represented by an event tree, with a root node n^0 in period 0 and a set of nodes $\mathcal{N}^t$ in period $t = 0, 1, \ldots, T$. Each node $n^t \in \mathcal{N}^t$ represents a possible sample value of the history h^t of the $\xi(.)$ process up to time t. Let $a(n^t) \in \mathcal{N}^{t-1}$ be the unique predecessor of node $n^t \in \mathcal{N}^t$ for $t = 0, 1, \ldots, T$, and denote by $S(n^t) \in \mathcal{N}^{t+1}$ the set of all possible direct successors of node $n^t \in \mathcal{N}^t$ for $t = 0, 1, \ldots, T-1$. We call *scenario* any path from node n^0 to a terminal node n^T. Each scenario has a probability, and the probabilities of all scenarios sum up to 1. We denote by π^{n^t} the probability of passing through node n^t, which corresponds to the sum of the probabilities of all scenarios that contain this node. In particular, $\pi^{n^0} = 1$, and π^{n^T} is equal to the probability of the single scenario that terminates in (leaf) node $n^T \in \mathcal{N}^T$. Also, $\sum_{n^t \in \mathcal{N}^t} \pi^{n^t} = 1, \forall t$.

Denote by $u_i(n^t) \in U_i^{n^t} \subseteq \mathbb{R}^{m_i^{n^t}}$ the decision variable of player i at node n^t, where $U_i^{n^t}$ is the control set, $m_i^{n^t}$ is the dimension of the decision variable for player i, $i = 1, 2$. Let $\mathbf{u}(n^t)$ denote the vector of decision variables for both players at node n^t, i.e., $\mathbf{u}(n^t) = (u_1(n^t), u_2(n^t))$. Let $X \subseteq \mathbb{R}^q$ be a state set. Denote by $x(n^t)$ the state vector at node n^t. An admissible S-adapted strategy (where S stands for sample, as in the terminology of Haurie et al. (2012)), for player i over the event tree is a vector $\mathbf{u}_i = \left(u_i(n^t) : n^t \in \mathcal{N}^t, t = 0, \ldots, T-1\right)$, that is, a plan of actions adapted to the history of the random process represented by the event tree.

Assuming a linear–quadratic game structure, the optimization problem of player i is as follows:

$$\max_{\mathbf{u}_i} V_i(x, \mathbf{u}) = \max_{\mathbf{u}_i} \sum_{t=0}^{T-1} \sum_{n^t \in \mathcal{N}^t} \pi^{n^t} \Big(\frac{1}{2} x'(n^t) Q_i(n^t) x(n^t) + p_i'(n^t) x(n^t) + \frac{1}{2} \sum_{j=1}^{2} u_j'(n^t) R_{ij}(n^t) u_j(n^t) \Big) + \sum_{n^T \in \mathcal{N}^T} \pi^{n^T} \Big(\frac{1}{2} x'(n^T) Q_i(n^T) x(n^T) + p_i'(n^T) x(n^T) \Big), \tag{1}$$

subject to

$$x(n^t) = A(a(n^t)) x(a(n^t)) + \sum_{j=1}^{2} B_j(a(n^t)) u_j(a(n^t)),$$

$$x(n^0) = x_0, \tag{2}$$

where $Q_i(n^t) \in \mathbb{R}^{q \times q}, R_{ij}(n^t) \in \mathbb{R}^{m_j^{n^t} \times m_j^{n^t}}, p_i(n^t) \in \mathbb{R}^q, A(n^t) \in \mathbb{R}^{q \times q}$ and $B_j(n^t) \in \mathbb{R}^{q \times m_j^{n^t}}$ for all $n^t \in \mathcal{N}^t, t \in T$.

Assumption 1. The matrices $Q_i(n^t)$ and $Q_i(n^T)$ are symmetric and negative semi definite and $R_{ii}(n^t)$ is negative definite. Additionally, the matrices $R_{ij}(n^t)$, $i \neq j$ are such that $R_{ii}(n^t) + R_{ji}(n^t)$ are negative definite as well.

By Assumption 1, the objective function in (1) will be strictly concave in the control variables. Note that if $\mathcal{N}^t$ consists of one element for all t, then this optimization problem reduces to the standard linear–quadratic optimal-control problem.

2.1. Cooperative solution

Suppose that the two players agree to cooperate and maximize their joint payoff, that is, $\max_{\mathbf{u}_i} \sum_{i=1}^{2} V_i(x, \mathbf{u})$ subject to (2). The Lagrangian associated with the joint optimization problem is given by:

$$\begin{aligned} \mathcal{L}^C = {} & \sum_{t=0}^{T-1} \sum_{n^t \in \mathcal{N}^t} \pi^{n^t} \Big(\frac{1}{2} x'(n^t) \big(Q_1(n^t) + Q_2(n^t)\big) x(n^t) \\ & + \big(p_1(n^t) + p_2(n^t)\big)' x(n^t) + \frac{1}{2} \begin{bmatrix} u_1(n^t) \\ u_2(n^t) \end{bmatrix}' \\ & \times \begin{bmatrix} R_{11}(n^t) + R_{21}(n^t) & 0 \\ 0 & R_{12}(n^t) + R_{22}(n^t) \end{bmatrix} \begin{bmatrix} u_1(n^t) \\ u_2(n^t) \end{bmatrix} \\ & + \sum_{n^T \in \mathcal{N}^T} \pi^{n^T} \Big(\frac{1}{2} x'(n^T) \big(Q_1(n^T) + Q_2(n^T)\big) x(n^T) \\ & + (p_1(n^T) + p_2(n^T))' x(n^T) \Big) + \big(\lambda^C(n^0)\big)' (x(n^0) - x_0) \\ & + \sum_{t=1}^{T} \sum_{n^t \in \mathcal{N}^t} \pi^{n^t} \big(\lambda^C(n^t)\big)' \Big(x(n^t) - A(a(n^t)) x(a(n^t)) \\ & - \sum_{j=1}^{2} B_j(a(n^t)) u_j(a(n^t)) \Big). \end{aligned} \tag{3}$$

This is a standard dynamic optimization problem with the following optimality conditions:

$$\frac{\partial \mathcal{L}^C}{\partial u_i^{n^t}} = \pi^{n^t} u_i'(n^t) \sum_{j=1}^{2} R_{ji}(n^t) + \lambda^C(S(n^t)) B_i^{n^t} = 0,$$

$$\Rightarrow u_i^C(n^t) = -\frac{1}{\pi^{n^t}} \Big(\sum_{j=1}^{2} R_{ji}(n^t) \Big)^{-1} B_i^{n^t} \lambda^C(S(n^t)). \tag{4}$$