Subgame perfect (ϵ-)equilibrium in perfect information games

János Flesch*

April 15, 2016

Abstract

We discuss recent results on the existence and characterization of subgame perfect (ϵ–)equilibrium in perfect information games.

The game. We consider games with perfect information and deterministic transitions. Such games can be given by a directed tree.[1] In this tree, each node is associated with a player, who controls this node. The outgoing arcs at this node represent the actions available to this player at this node. We assume that each node has at least one successor, rather than having terminal nodes.[2] Play of the game starts at the root. At any node z that play visits, the player who controls z has to choose one of the actions at z, which brings play to a next node. This induces an infinite path in the tree from the root, which we call a play. Depending on this play, each player receives a payoff.

Note that these payoffs are fairly general. This setup encompasses the case when the actions induce instantaneous rewards which are then aggregated into a payoff, possibly by taking the total discounted sum or the long-term average. It also includes payoff functions considered in the literature of computer science (reachability games, etc.).

Subgame-perfect (ϵ-)equilibrium. We focus on pure strategies for the players. A central solution concept in such games is subgame-perfect equilibrium, which is a strategy profile that induces a Nash equilibrium in every subgame, i.e. when starting at any node in the tree, no player has an incentive to deviate individually from his continuation strategy. For a tree with finite depth and finite number of actions, a subgame-perfect equilibrium is known to exist and can be found by backward induction. However, if the tree has infinite depth, backward induction is no longer applicable and a subgame-perfect equilibrium may

*Department of Quantitative Economics, Maastricht University, P.O.Box 616, 6200 MD, The Netherlands. E-mail: j.flesch@maastrichtuniversity.nl.

[1] A node encodes a history of play.

[2] Every tree with finite depth can be transformed into a strategically equivalent tree with infinite depth. Indeed, we can extend a finite tree by simply adding one infinite sequence of arcs to every terminal node. So, instead of termination, play will continue along a unique path in which the players have no further strategic choices.

fail to exist even in simple cases. For example, consider the one-player game in which the player can choose between Left and Right every day. If he chooses Right for the first time at day n, then his payoff is $1-\frac{1}{n}$, whereas if he never chooses Right then his payoff is 0. In this game, the player can obtain a payoff arbitrarily close to 1, but no strategy gives him exactly 1. In other words, he has no optimal strategy, and there is no Nash equilibrium or subgame-perfect equilibrium either. Nevertheless, this game admits approximate optimal strategies with arbitrarily small positive error-terms, i.e. for every error-term $\epsilon > 0$, if the player chooses Left at days $1, \ldots, n-1$ and Right at day n, where $n \geq \frac{1}{\varepsilon}$, then his payoff is at least $1-\epsilon$. Such approximate solutions often offer a remedy to the non-existence of exact solutions.

For an error-term $\epsilon > 0$, a strategy profile is called a subgame-perfect ϵ-equilibrium, if it induces a Nash ϵ-equilibrium in every subgame, i.e. no player can gain more than ϵ by a unilateral deviation in any subgame.

Topological setup. We endow the set of plays (i.e. infinite paths from the root) with the following metric, just as in Martin [1975, Borel determinacy] and in the literature on descriptive set theory: Consider two different plays $z = (z_0, z_1, \ldots)$ and $z' = (z'_0, z'_1, \ldots)$, where z_t and z'_t respectively denote the nodes visited at time period t. Then, their distance is defined as $d(z, z') = 2^{-m(z,z')}$, where $m(z, z') \in \mathbb{N}$ is the first coordinate on which z and z' differ. Intuitively, according to this metric, two plays are close if they have a long common prefix. The induced topology is the same as the one where the action sets are endowed with the discrete topology and the set of plays with the product topology. Consequently, the set of plays is compact if and only if there are only finitely many actions at the nodes, and the set of plays is separable if and only if there are countably many actions at the nodes.

Outline of the talk: We discuss recent results on the existence and characterization of subgame perfect (ϵ–)equilibrium in perfect information games, under the above topological structure. We distinguish conditions mainly on the payoff functions of the players (continuity/semi-continuity, finite range/bounded range). Sometimes we need to assume that the number of players is finite, but the set of actions is almost always unrestricted. The outline is roughly as follows:

1. Under the assumption that the payoff function of each player has a finite range: Here we look at subgame-perfect equilibrium in pure strategies under the following further assumptions:

1.1 continuous payoff functions

1.2 weaker assumptions:

- when the payoff functions have only sigma-discrete discontinuities
- lower-semicontinuous payoff functions, finitely many players
- upper-semicontinuous payoff functions

– payoff functions with common preferences at the limit, finitely many players

1.3 characterization of plays that are induced by a pure subgame-perfect equilibrium

2. Under the assumption that the payoff function of each player is bounded: Here we look at subgame-perfect ϵ-equilibrium in pure strategies, for all $\epsilon > 0$. We mainly discuss discretization results (which often lead back to case 1).

3. Under the assumption that the payoff function has a recursive structure, and the tree is induced by a finite graph: Here we look at the existence of subgame-perfect ϵ-equilibrium in randomized strategies, for all $\epsilon > 0$.