Refinements of subgame-perfect ε-equilibrium in games with perfect information

János Flesch
Maastricht University, The Netherlands

— based on joint papers [1] and [2] in the references —

Abstract

We consider games with perfect information and deterministic transitions. A common solution concept is the concept of subgame-perfect ε-equilibrium, where $\varepsilon \geq 0$, which is a strategy profile such that no player can improve his payoff in any subgame by more than ε. We propose and examine a number of refinements of this concept. A major emphasis lies on existence results. Roughly speaking, the most important refinements require the following respective properties: (1) no player makes a big mistake with positive probability, (2) the mistakes vanish as the horizon approaches infinity, i.e. ε depends on the subgame and goes to 0 as play proceeds, and (3) for pure strategy profiles, the induced play paths are continuity points of the payoff functions.

The game: We consider games with perfect information and deterministic transitions, i.e. without chance moves. Such a game can be given by a triple $\mathcal{G} = (N, T, u = (u_i)_{i \in N})$, where

• N is a finite and nonempty set of players.

• T is a directed tree with a root, in which each node is associated with a player, who controls this node. We assume that, at each node in the tree, the number of outgoing arcs is finite and is at least one. This implies that there are no terminal nodes and the tree has an infinite depth.[1] Let $\mathcal{P}$ denote the set of all infinite paths (often called plays) starting at the root. We endow $\mathcal{P}$ with the topology induced by the cylinders.[2]

• $u_i : \mathcal{P} \to \mathbb{R}$ is a payoff function for player i. We assume that u_i is bounded and Borel measurable.[3]

Play of the game starts at the root, and at any node z that play visits, the player who controls z has to choose one of the outgoing arcs at z, which brings play to a next node. This induces an infinite path p in the tree, and each player $i \in N$ receives the corresponding payoff $u_i(p)$.

[1]This way we allow for fairly general games, since every tree with finite depth can be easily 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.

[2]If z is a node in the tree, then the cylinder set corresponding to z is the set of all infinite paths from the root that go through z.

[3]Note that these payoffs are very general. Indeed, only a very special case would be a common situation when these payoffs arise by aggregating certain daily payoffs in the game, possibly by taking the total discounted sum.

Strategy: A strategy σ_i for a player i is a map that, to each node z controlled by player i, assigns a probability distribution $\sigma_i(z)$ on the outgoing arcs. The interpretation is that if node z is visited during the play, then player i should choose the outgoing arc according to $\sigma_i(z)$. A vector of strategies $\sigma = (\sigma_i)_{i \in N}$ is called a strategy profile, and $u_i(\sigma)$ denotes the corresponding expected payoff for player i.

ε-Equilibrium: A strategy profile $\sigma = (\sigma_i)_{i \in N}$ is called an ε-equilibrium, where $\varepsilon \geq 0$, if no player can gain more than ε by deviating alone: for every player $i \in N$ and every strategy σ_i' for player i we have $u_i(\sigma_i', \sigma_{-i}) \leq u_i(\sigma) + \varepsilon$. Here, (σ_i', σ_{-i}) denotes the strategy profile in which player i uses the strategy σ_i' whereas all other players j follow σ_j. One can interpret ε as an error-term. It is shown by Mertens and Neyman (cf. Mertens [4]) that every game $\mathcal{G}$ admits an ε-equilibrium, for every $\varepsilon > 0$.

Subgame-perfect ε-equilibrium: Let $\mathcal{G}$ be a game and let z denote a node in the tree. The subgame $\mathcal{G}_z$ is the game that starts at node z. The payoff function of player i in this subgame is denoted by u_{iz}, and it is simply derived from the original payoff function u_i as follows: if p is an infinite path starting at z, then $u_{iz}(p) = u_i(h_z \otimes p)$ where $h_z \otimes p$ is the path that starts at the root, reaches z and then follows p.

A strategy profile $\sigma = (\sigma_i)_{i \in N}$ is called a subgame-perfect ε-equilibrium, where $\varepsilon \geq 0$, if σ induces an ε-equilibrium in every subgame, i.e. $\sigma_z = (\sigma_{iz})_{i \in N}$ is an ε-equilibrium in $\mathcal{G}_z$ for every node z. Flesch et al. [1] presented a game in which a subgame-perfect ε-equilibrium fails to exist for small $\varepsilon > 0$. Nevertheless, if the payoff functions are all lower-semicontinuous or they are all upper-semicontinuous, then a subgame-perfect ε-equilibrium does exist for every $\varepsilon > 0$ (cf. Flesch et al. [3] and Purves and Sudderth [5]). However, mainly from a conceptual point of view, there are good reasons to look for refinements of the concept of subgame-perfect ε-equilibrium.

Strong subgame-perfect ε-equilibrium: The concept of subgame-perfect ε-equilibrium has the drawback that it does not rule out that a player chooses, with small probability, an outgoing arc that leads to a low payoff. So, if one interprets the probability distribution, which a player's strategy prescribes on the outgoing arcs, as a lottery that the player uses for the choice of an action, and the lottery picks an action with low payoff, then the player may become reluctant to execute this action. We (cf. Flesch et al. [1]) propose and examine a refinement, called strong subgame-perfect ε-equilibrium, which avoids this shortcoming.

Subgame-perfect ε-equilibrium with vanishing errors: The concept of subgame-perfect ε-equilibrium uses the same upper-bound on the error-term, i.e. ε, in every subgame. We (cf. Flesch and Predtetchinski [2]) propose an additional requirement: roughly speaking, the error-terms should go to 0 as the time horizon approaches infinity. More formally, for a strategy profile σ and a node z, let $e(\sigma, z)$ denote the maximal improvement that a player can achieve in the subgame $\mathcal{G}_z$ by deviating alone:

$$e(\sigma, z) = \max_{i \in N} \sup_{\sigma_{iz}'} \left[u_{iz}(\sigma_{iz}', \sigma_{-iz}) - u_{iz}(\sigma_z) \right].$$

A subgame-perfect ε-equilibrium σ is said to have vanishing errors, if in every subgame $\mathcal{G}_z$, the following holds: if the players use σ, the errors $e(\sigma, z_n)$ converge to 0 with probability 1, where z_n denotes the node visited at period n.

The motivation to consider this requirement is two-fold. First, if the payoff functions are continuous (for example, for discounted payoffs), then the following property holds for every infinite path p from the root and every $\delta > 0$: there is a node z on p such that the payoffs are essentially fixed at z, i.e. they can vary at most δ depending on the continuation after z. This means that, if $\delta < \varepsilon$, then requiring that the errors are at most ε is automatically satisfied and the concept of subgame-perfect ε-equilibrium does not have a bite for such deep nodes. (For discounted payoffs, one may wish to require that not only the payoffs, but also the error ε is discounted.) Another motivation might be that once play arrives at deeper nodes, the players may have a better overview of future possibilities. So, it may be reasonable to assume that their mistakes also become smaller.

We (cf. [2]) prove that if all payoff functions are upper-semicontinuous, then a pure subgame-perfect ε-equilibrium with vanishing errors exists, for every $\varepsilon > 0$.[4] However, given errors $r(z) > 0$ for all nodes z, there may be no strategy profile such that $e(\sigma, z) \leq r(z)$ for every node z. Not even when the payoffs are upper-semicontinuous nor when they are lower-semicontinuous, which we (cf. [2]) demonstrate by counter-examples.

Subgame-perfect ε-equilibrium that is robust to small perturbations in the strategies: If the payoff functions are not continuous, and a pure subgame-perfect ε-equilibrium induces an infinite path which is a discontinuity point of the payoff functions, then small perturbations in the strategies may lead to drastically different payoffs. Such small perturbations may occur when players are not able to execute their strategies with full precision. We (cf. [2]) examine when a pure subgame-perfect ε-equilibrium exists which induces a path in every subgame such that these paths are all continuity points of the payoff functions. We prove that if all payoff functions are lower-semicontinuous, then this is indeed the case. The proof is based on the transfinite induction used in Flesch et al. [3]. Then, we present a simple counter-example for the case when the payoff functions are upper-semicontinuous.

References:

[1] Flesch J, Kuipers J, Mashiah-Yaakovi A, Schoenmakers G, Shmaya E, Solan E and Vrieze K [2013]: Equilibrium refinements in perfect information games with infinte horizon. Working paper.

[2] Flesch J and Predtetchinski A [2013]: Subgame-perfect ε-equilibrium with vanishing errors. Working paper.

[3] Flesch J, Kuipers J, Mashiah-Yaakovi A, Schoenmakers G, Solan E and Vrieze K. [2010]: Perfect-information games with lower-semi-continuous payoffs. Mathematics of Operations Research 35, 742-755.

[4] Mertens J.-F. [1987]: Repeated games. Proceedings of the International Congress of Mathematicians, American Mathematical Society, Providence, RI, 1528-1577.

[5] Purves R and Sudderth W [2011]: Perfect information games with upper semicontinuous payoffs. Mathematics of Operations Research 36, 468-473.

[4] This results extends to infinitely many actions. In fact, the result is even stronger: the errors even converge to 0 on every infinite path from the root, and not only on the ones that the strategy profile induces in the subgames.