Subgame perfect ϵ–equilibrium in games of perfect information: a review

Arkadi Predtetchinski
Maastricht University

GAMENET Workshop
Theory and Algorithms in Graph and Stochastic Games
Mons, March 14–15, 2019

A meeting point for

- Computer science
- Game theory
- Economics

The plan

1. Preliminaries
2. Weak subgame perfect ϵ–equilibrium
3. Games with semicontinuous payoffs
4. Extensions
5. Mixed strategies
6. Open questions

The plan

Perfect information games

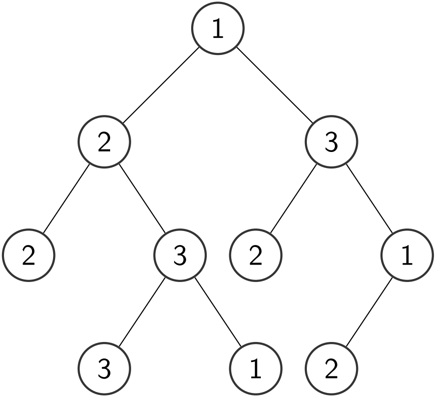

Perfect information games

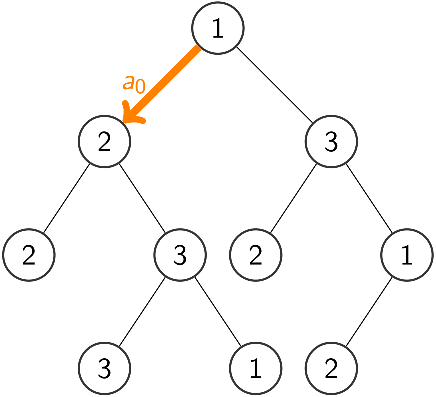

- Player 1 chooses action a_0.

Perfect information games

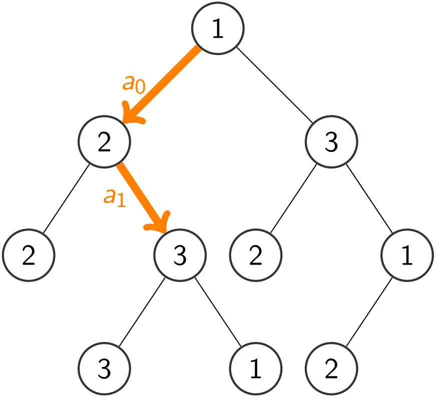

- Player 1 chooses action a_0.
- Player 2 chooses action a_1.

Perfect information games

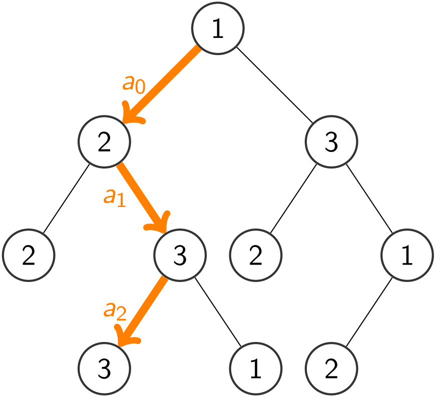

- Player 1 chooses action a_0.
- Player 2 chooses action a_1.
- Player 3 chooses action a_2, and so on.

Perfect information games

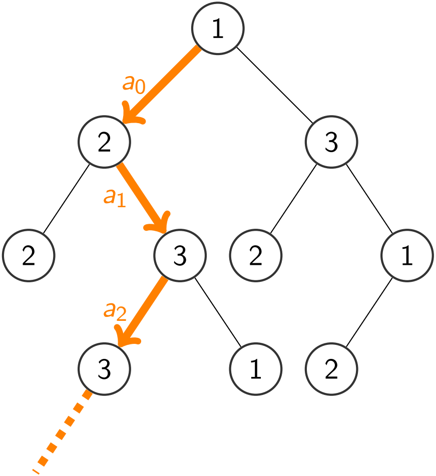

- Player 1 chooses action a_0.
- Player 2 chooses action a_1.
- Player 3 chooses action a_2, and so on.
- This gives an infinite sequence $p = (a_0, a_1, \ldots)$ called a **play**. Player i receives payoff $u_i(p)$.

Perfect information games

- Let T be a pruned tree.
- Nodes (vertices) of T are called (decision) **histories**. Each history is assigned an active player.

Perfect information games

- Let T be a pruned tree.
- Nodes (vertices) of T are called (decision) **histories**. Each history is assigned an active player.
- Let P be the set of infinite branches of T.
- Elements of P are called **plays**.
- The set P is endowed with its usual distance.

Perfect information games

- Let T be a pruned tree.
- Nodes (vertices) of T are called (decision) **histories**. Each history is assigned an active player.
- Let P be the set of infinite branches of T.
- Elements of P are called **plays**.
- The set P is endowed with its usual distance.
- Each player i has a **payoff function** $u_i : P \to \mathbb{R}$.

The distance on P

$d(p,q) = 2^{-m}$ where m is the length of the longest common prefix of p and q.

The distance on P

$d(p, q) = 2^{-m}$ where m is the length of the longest common prefix of p and q.

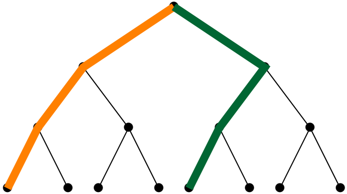

$d(\bullet, \bullet) = 1$

The distance on P

$d(p, q) = 2^{-m}$ where m is the length of the longest common prefix of p and q.

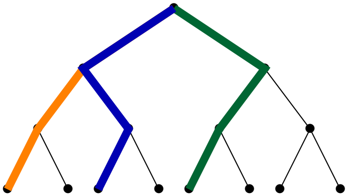

$d(\bullet, \bullet) = 1$

$d(\bullet, \bullet) = 1/2$

The distance on P

$d(p, q) = 2^{-m}$ where m is the length of the longest common prefix of p and q.

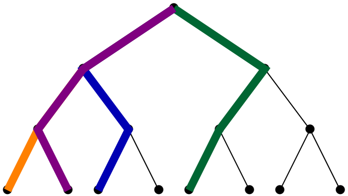

$d(\bullet, \bullet) = 1$

$d(\bullet, \bullet) = 1/2$

$d(\bullet, \bullet) = 1/4$

ϵ–Equilibria

- A joint strategy σ is an **ϵ–equilibrium** a player can gain at most ϵ from a unilateral deviation.

Theorem (Mertens and Neyman [1987], Le Roux and Pauly [2014]): *Consider a game with a finite or a countable set of players where each player's payoff function $u_i : P \to \mathbb{R}$ is bounded and Borel–measurable. Then for each $\epsilon > 0$ the game has an ϵ–equilibrium in pure strategies.*

- This result builds on Borel determinacy (**Martin [1975]**).

Subgame perfect ϵ–equilibria

- **Selten [1965, 1975]**
- A joint strategy σ is a **subgame perfect ϵ–equilibrium** (SPϵE) if it induces an ϵ–equilibrium in each subgame.
- Equivalently, σ is a SPϵE if **in any subgame** a player can gain at most ϵ from a unilateral deviation.

SPE versus equilibrium

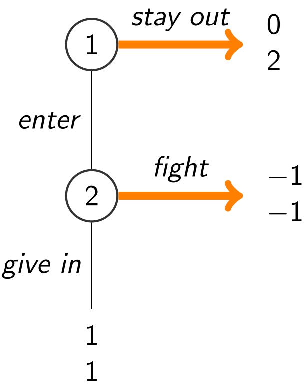

- (*stay out*, *fight*) is an equilibrium, but not a SPE.
- Would player 2 actually fight if given a chance?
- SPE eliminates **non–credible threats**

Finite games

Zermelo's Theorem (Kuhn [1953]): *A finite game has a SPE. All SPE can be found by the process of backward induction.*

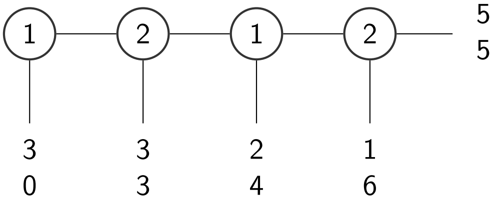

Finite games

Zermelo's Theorem (Kuhn [1953]): *A finite game has a SPE. All SPE can be found by the process of backward induction.*

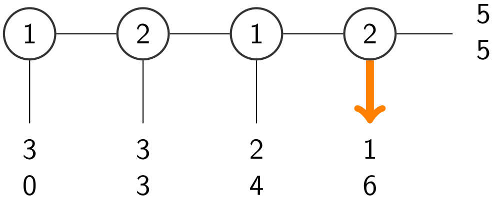

Finite games

Zermelo's Theorem (Kuhn [1953]): *A finite game has a SPE. All SPE can be found by the process of backward induction.*

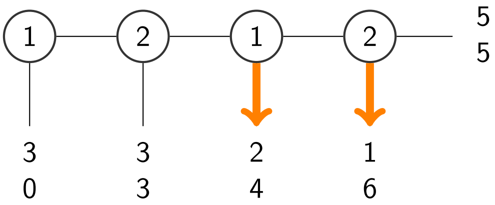

Finite games

Zermelo's Theorem (Kuhn [1953]): *A finite game has a SPE. All SPE can be found by the process of backward induction.*

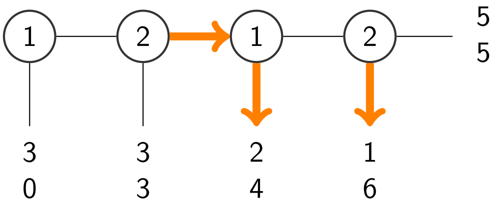

Finite games

Zermelo's Theorem (Kuhn [1953]): *A finite game has a SPE. All SPE can be found by the process of backward induction.*

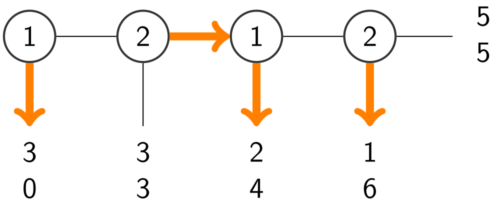

The plan

1. Preliminaries
2. **Weak subgame perfect ϵ–equilibrium**
3. Games with semicontinuous payoffs
4. Extensions
5. Mixed strategies
6. Open questions

An example by Solan and Vieille [2003]

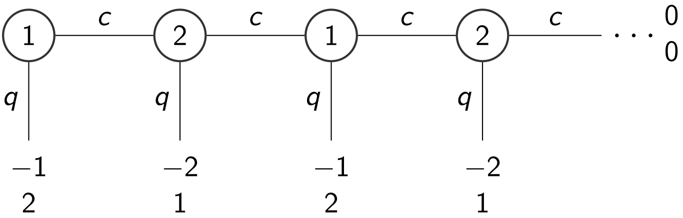

- Always playing q is an equilibrium, but not a SPE.
- The game has no SPϵE in **pure strategies**.

An example by Solan and Vieille [2003]

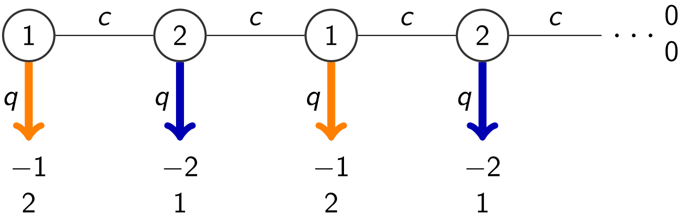

- Always playing q is an equilibrium, but not a SPE.
- The game has no SPϵE in **pure strategies**.

Weak and very weak SPϵE

- **Brihaye, Bruyère, Meunier, Raskin [2015]**
- A joint pure strategy σ is a **weak SPϵE** if in any subgame a player can gain at most ϵ by deviating from σ at **finitely many histories only**.
- A joint pure strategy σ is a **very weak SPϵE** if in any subgame a player can gain at most ϵ by deviating from σ at a **single history only**.

Weak and very weak SPϵE

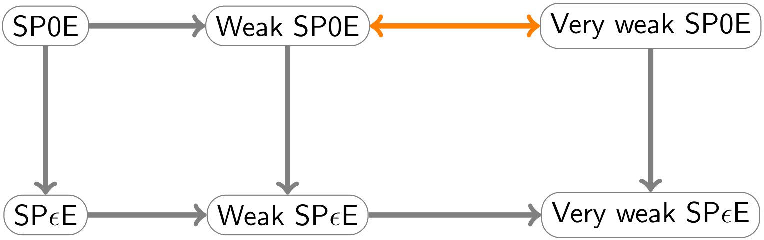

Weak and very weak SPϵE

Remark: If there are **finitely many players**, and each player's payoff function only takes **finitely many values**, then for $\epsilon > 0$ small enough:

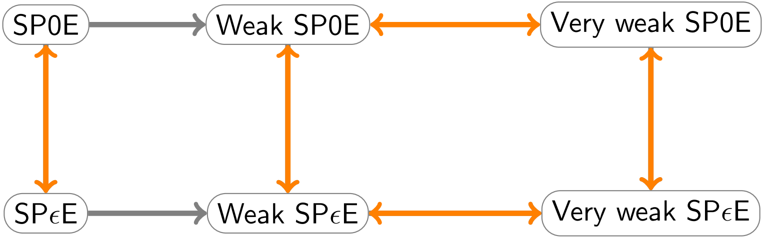

In Solan–Vieille's example

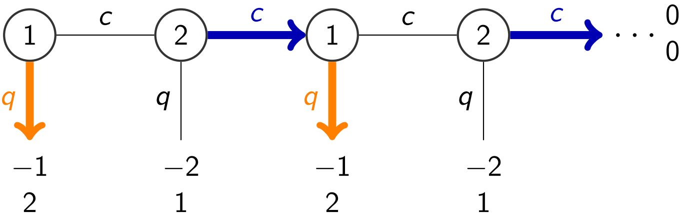

- **Very weak SP0E = weak SP0E**: Player 1 always playing q, player 2 always playing c.

Existence of weak SPϵE

Theorem (Flesch, Kuipers, Mashiah–Yaakovi, Schoenmakers, Solan, Vrieze [2010] and Bruyère, Le Roux, Pauly, Raskin [2017]): *Let $\epsilon > 0$. A game with finitely many players and bounded payoff functions admits a weak SPϵE.*

Existence of weak SPϵE

Theorem (Flesch, Kuipers, Mashiah–Yaakovi, Schoenmakers, Solan, Vrieze [2010] and Bruyère, Le Roux, Pauly, Raskin [2017]): *Let $\epsilon > 0$. A game with finitely many players and bounded payoff functions admits a weak SPϵE.*

- The result leads to several sufficient conditions for the existence of SPϵE.
- One such condition is lower semicontinuity of the payoff functions.

Existence of weak SPϵE: the proof

- A transfinite procedure of elimination of plays/outcomes.
- Start with all plays
- Fix a history and consider the active player's **guarnatee level**: maximizing the utility of the worst possible continuation play after his move.
- Eliminate the plays that do not meet the guarantee levels.
- Repeat...

The plan

1. Preliminaries
2. Weak subgame perfect ϵ–equilibrium
3. **Games with semicontinuous payoffs**
4. Extensions
5. Mixed strategies
6. Open questions

Semicontinuous payoff functions: very informally

- A payoff function is **lower semi–continuous** if the payoff can suddenly jump **downwards** (but not upwards).
- A payoff function is **upper semi–continuous** if the payoff can suddenly jump **upwards**.

Semicontinuous payoff functions

- The function u_i is **lower semi–continuous** if

$$p_n \to p \qquad \Longrightarrow \qquad \liminf_{n\to\infty} u_i(p_n) \geq u_i(p).$$

- The function u_i is **upper semi–continuous** if

$$p_n \to p \qquad \Longrightarrow \qquad \limsup_{n\to\infty} u_i(p_n) \leq u_i(p).$$

Semicontinuous payoff functions

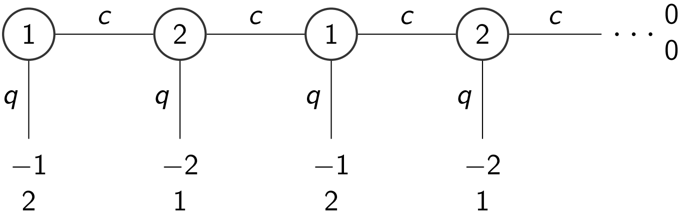

- Player 2's payoff function is **lower** semi–continuous,
- Player 1's payoff function is **upper** semi–continuous.

Semicontinuous payoff functions

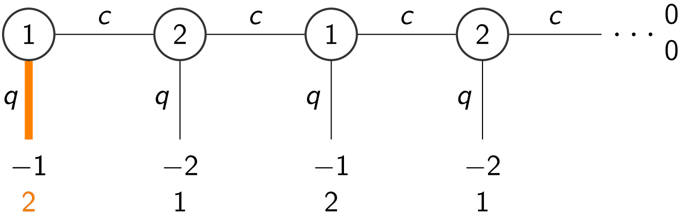

- Player 2's payoff function is **lower** semi–continuous,
- Player 1's payoff function is **upper** semi–continuous.

Semicontinuous payoff functions

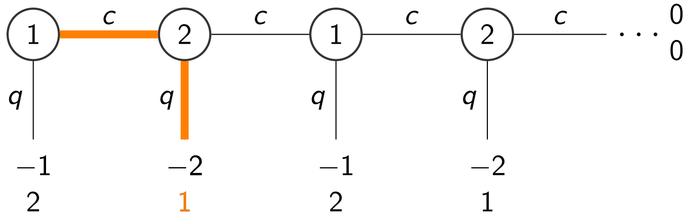

- Player 2's payoff function is **lower** semi–continuous,
- Player 1's payoff function is **upper** semi–continuous.

Semicontinuous payoff functions

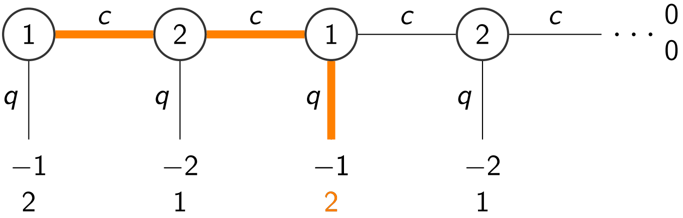

- Player 2's payoff function is **lower** semi–continuous,
- Player 1's payoff function is **upper** semi–continuous.

Semicontinuous payoff functions

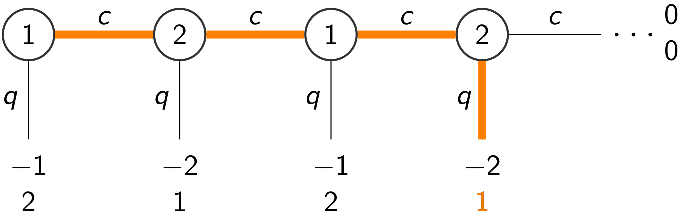

- Player 2's payoff function is **lower** semi–continuous,
- Player 1's payoff function is **upper** semi–continuous.

Semicontinuous payoff functions

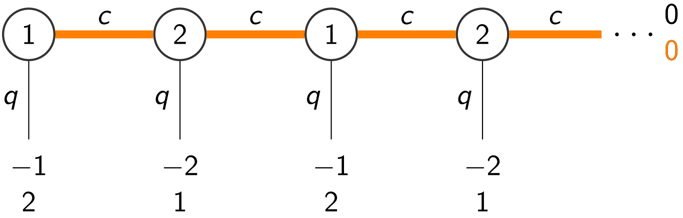

- Player 2's payoff function is **lower** semi–continuous,
- Player 1's payoff function is **upper** semi–continuous.

One–shot deviation principle

Lemma (One–shot deviation principle): *If each player's payoff function is bounded and* **lower semicontinuous***:*

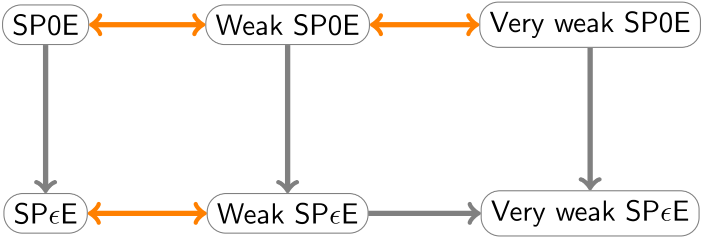

One–shot deviation principle

Corollary (to the existence of weak SPϵE, Flesch et al [2010], Bruyère et al [2017]): *Suppose that there are finitely many players, and that each player's payoff function is bounded and lower semicontinuous. Then for each $\epsilon > 0$ the game admits a SPϵE.*

Upper semicontinuous payoffs

Theorem: *Suppose that there are finitely many players, and that each player's payoff function is bounded and* **lower** *semicontinuous. Then for each $\epsilon > 0$ the game admits a SPϵE.*

Theorem (Purves and Sudderth [2011]): *Suppose that there are finitely many players, and that each player's payoff functions is bounded and* **upper** *semicontinuous. Then for each $\epsilon > 0$ the game admits a SPϵE.*

Upper semicontinuous payoffs: the proof

- Curiously, the proof of the result for upper semicontinuous payoffs is **essentially different** from that for lower semicontinuous payoffs.
- If p is any play, then there is a time t after which p becomes the best play for everyone.
- Backward induction.

The plan

1. Preliminaries
2. Weak subgame perfect ϵ–equilibrium
3. Games with semicontinuous payoffs
4. **Extensions**
   - Common preferences at the limit
   - A criterion for existence
   - Infinitely many players
5. Mixed strategies
6. Open questions

Common preferences at the limit

- A generalization/unification of these two results is possible using the concept of common preferences at the limit (**Flesch and P [2016]**).
- The game exhibits **common preferences at the limit** if for any convergent sequence of plays, either
  - each player's payoff jumps downwards,
  - each player's payoff jumps upwards.

A criterion for preferences that admit SPE

Theorem (Le Roux [2016]): *Consider some preferences over a set of outcomes. The following are equivalent:*

1. *Any game with the given preferences admits a SPE, provided that the outcomes are assigned to plays in a* $\mathbf{\Delta^0_2}$*–measurable fashion.*
2. *There are no two players* 1 *and* 2 *and no outcomes* x, y, *and* z *that form a* **Solan–Vieille cycle***:*

$$z <_1 y <_1 x \text{ and } x <_2 z <_2 y.$$

Infinitely many players

- Non–existence of **pure** SPϵE in a game with lower semicontinuous functions (**Flesch et al [2010]**):

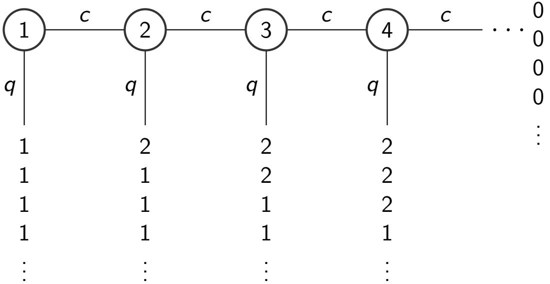

Infinitely many players

Theorem (Flesch and P [2017]): *Suppose that in a game with finitely or* **infinitely** *many players each player's payoff function is bounded and* **upper** *semicontinuous. Then for each $\epsilon > 0$ the game admits a SPϵE.*

- The proof in Purves and Sudderth [2011] is difficult to extend.
- A transfinite procedure of elimination of plays.
- As a by product: a characterization of subgame perfect plays.

The plan

1. Preliminaries
2. Weak subgame perfect ϵ–equilibrium
3. Games with semicontinuous payoffs
4. Extensions
5. **Mixed strategies**
6. Open questions

Solan Vieille's example: a SPϵE in mixed strategies

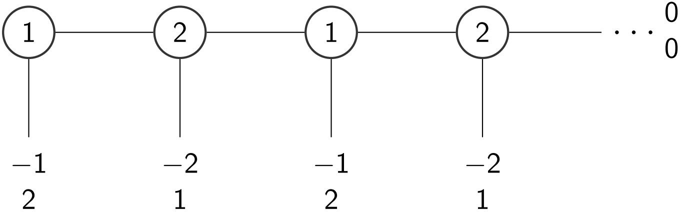

Solan Vieille's example: a SPϵE in mixed strategies

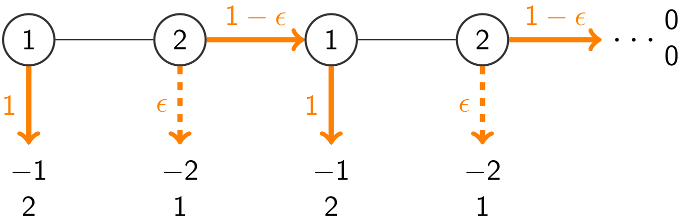

- Player 1 quits with probability 1, player 2 quits with probability ϵ. This works because:
  - Player 1: deviation to "*always continue*" is not profitable, since then player 2 will stop eventually with probability 1, resulting in the payoff of -2.
  - Player 2: is within ϵ of his best possible payoff of 2.

Solan Vieille's example: a SPϵE in mixed strategies

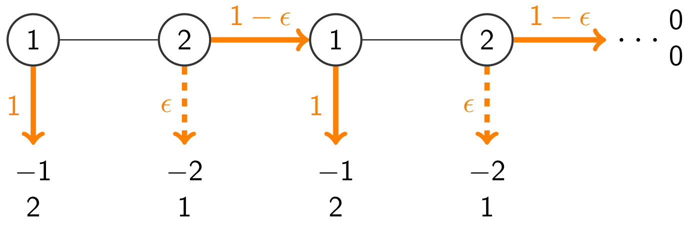

- Player 1 quits with probability 1, player 2 quits with probability ϵ. This works because:
  - Player 1: deviation to "*always continue*" is not profitable, since then player 2 will stop eventually with probability 1, resulting in the payoff of -2.
  - Player 2: is within ϵ of his best possible payoff of 2.
- This does not work for $\epsilon = 0$.

Non–existence of SPϵE

Theorem (Flesch, Kuipers, Mashiah–Yaakovi, Schoenmakers, Shmaya, Solan, Vrieze [2014]) *The following game has no SPϵE, whether in* **pure or mixed** *strategies:*

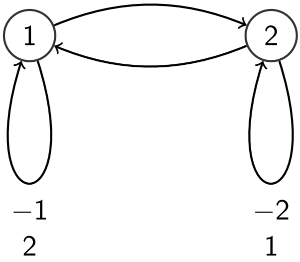

The payoffs are

- $(-1, 2)$ if player 2 is active only finitely many times,
- $(-2, 1)$ if player 1 is active only finitely many times,
- $(0, 0)$ if both players are active infinitely many times.

The main idea behind the non–existence

- The proof of the non–existence is non–trivial. Ultimately it relies on the application of Lèvy's zero–one law.

Sufficient conditions for SPϵE

Positive results are proven for some special classes:

- **Brihaye,Bruyère, De Pril, Gimbert [2012]:** quantitative reachability games.
- **Mashiah-Yaakovi [2014]**: stopping games.
- **Kuipers, Flesch, Schoenmakers, Vrieze [2018]:** recursive perfect information games.
- **Flesch and P [2016]:** games where the set of discontinuity points is sigma–discrete.

Open questions

1. In perfect information games with **chance moves**, is there a weak SPϵE?
2. In a game with infinitely many players and lower semicontinuous functions, is there a (mixed) SPϵE?
3. Suppose that:
   - for each player i, the payoff function u_i only takes finitely many values,
   - for each payoff vector v, the set $\{p \in A^{\mathbb{N}} : u(p) = v\}$ is an open set or a closed set.

   Is there a (mixed) SPϵE?