Limit value in stochastic games

Jérôme Renault, TSE Université Toulouse 1

ReLaX Workshop on Games, Chennai online Feb 1, 2021.

2-player zero-sum stochastic games: $\lim_{\lambda \to 0} v_\lambda$?

Introduction

- Given a A in $I\!R^{I\times J}$, the value of the matrix A is:

$$\mathrm{val}(A) = \mathrm{val}_{\Delta(I)\times\Delta(J)}(A) = \max_{x\in\Delta(I)} \min_{y\in\Delta(J)} \sum_{i\in I}\sum_{j\in J} x_i y_j a_{i,j} = \min_{y\in\Delta(J)} \max_{x\in\Delta(I)} \sum_{i\in I}\sum_{j\in J} x_i y_j a_{i,j}$$

$$\mathrm{val}\begin{pmatrix} 2 & -1 \\ -1 & 1 \end{pmatrix} = \frac{1}{5}$$

- Standard stochastic game: at each stage one matrix will be played, and the matrix of today together with the actions played today determine the matrix played tomorrow.

$$\begin{array}{cc} & \begin{array}{cc} L & R \end{array} \\ \begin{array}{c} T \\ B \end{array} & \begin{pmatrix} 0 & 1^* \\ 1^* & 0^* \end{pmatrix} \end{array}$$

• Zero-sum stochastic game: a dynamic zero-sum game with a Markovian structure, played in discrete time.

Outline:

1. The standard model: finitely many states and actions
1.1 The n-stage game and the λ-discounted game
1.2 Limit value

2. Beyond the standard model
2.1 The 1-player case
2.2 A compact continuous game with no limit value
2.3 A hidden stochastic game with no limit value
2.4 A non zero-sum stochastic game with no limit equilibrium payoff set

3. Concluding Remarks

1. Standard model (Shapley, 1953): a set of states K with an initial state k_1, actions sets I for player 1 and J for player 2, a payoff function $g : K \times I \times J \longrightarrow I\!R$, and a transition $q : K \times I \times J \longrightarrow \Delta(K)$.
K, I and J are assumed to be non empty *finite* sets.

Progress of the game:
- stage 1: players simultaneously choose $i_1 \in I$ and $j_1 \in J$. i_1 and j_1 are publicly announced, and P1's stage payoff is $g(k_1, i_1, j_1)$.
- stage $t \geq 2$: k_t is selected according to $q(k_{t-1}, i_{t-1}, j_{t-1})$ and publicly announced. Players then simultaneously choose $i_t \in I$ et $j_t \in J$. i_t et j_t are announced, and P1's payoff is $g(k_t, i_t, j_t)$.

Example 1:

$$\begin{array}{cc} & \begin{array}{cc} L & R \end{array} \\ \begin{array}{c} T \\ B \end{array} & \left(\begin{array}{cc} 0 & 1^* \\ 1^* & 0^* \end{array} \right) \end{array}$$

3 states. 1^* and 0^* are absorbing states.

A strategy of a player associates to every finite history $(k_1, i_1, j_1, ..., k_{t-1}, i_{t-1}, j_{t-1}, k_t)$, with $t \geq 1$, a mixed action in $\Delta(I)$ or $\Delta(J)$.

A couple of strategies (σ, τ) induces a probability distribution $I\!P_{k_1,\sigma,\tau}$ over the set of plays (endowed with the product σ-algebra).

The n-stage game and the λ-discounted game

• For $n \geq 1$, the n-stage game with initial state k_1 is the zero-sum game with payoff:

$$\gamma_n^{k_1}(\sigma,\tau) = I\!E_{k_1,\sigma,\tau}\left(\frac{1}{n}\sum_{t=1}^{n} g(k_t,i_t,j_t)\right).$$

It has a value: $v_n(k_1) = \max_\sigma \min_\tau \gamma_n^{k_1}(\sigma,\tau) = \min_\tau \max_\sigma \gamma_n^{k_1}(\sigma,\tau)$.

• Given a discount rate λ in $(0,1]$, the λ-discounted game with initial state k_1 is the zero-sum game with payoff :

$$\gamma_\lambda^{k_1}(\sigma,\tau) = I\!E_{k_1,\sigma,\tau}\left(\lambda\sum_{t=1}^{\infty}(1-\lambda)^{t-1} g(k_t,i_t,j_t)\right).$$

It has a value denoted by $v_\lambda(k_1)$.
$\delta = 1-\lambda = \frac{1}{1+r}$ is called the discount factor, and r is the "interest rate".

Proposition v_n and v_λ are characterized by the Shapley equations:

- For $n \geq 0$ and k dans K:

$$(n+1)\, v_{n+1}(k) = \mathrm{Val}_{\Delta(I)\times\Delta(J)}\left(g(k,i,j) + \sum_{k'\in K} q(k'|k,i,j)\, n\, v_n(k')\right).$$

And in any n-stage game, players have Markov optimal strategies.

- For λ in $(0,1]$ and k in K:

$$v_\lambda(k) = \mathrm{Val}_{\Delta(I)\times\Delta(J)}\left(\lambda\, g(k,i,j) + (1-\lambda)\sum_{k'\in K} q(k'|k,i,j)\ v_\lambda(k')\right).$$

And in any λ-discounted game, players have stationary optimal strategies.

Example 1: $v_n = \frac{1}{2-\frac{n-1}{n}v_{n-1}}$ for $n \geq 1$, and $v_\lambda = \frac{1}{1+\sqrt{\lambda}}$ for each λ.

Example 2: A one-player game, with deterministic transitions and actions Black and Blue for Player 1. Payoffs are 1 or 0 in each case.

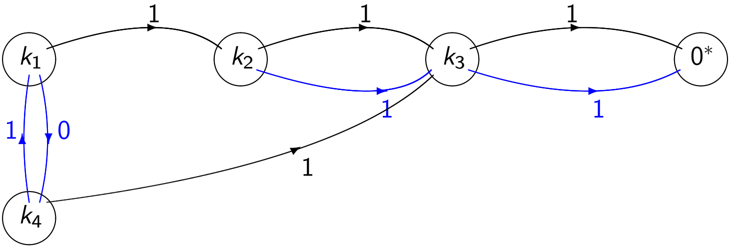

For λ small enough, $v_\lambda(k_1) = \frac{1}{2-\lambda}$ and it is optimal in the λ-discounted game to alternate between states k_1 and k_4.

For $n \geq 0$, $(2n+3)v_{2n+3} = (2n+4)v_{2n+4} = n+3$ (first alternate between k_1 and k_4, then go to k_2 3 or 4 stages before the end).

Limit values

1) The 1-player case: Markov Decision Processes.

For $\lambda > 0$, player 1 has a *pure* stationary optimal strategy in the λ- game. The λ-discounted payoff of a pure stationary strategy $\mathbf{i} : K \to I$ satisfies:

$$\gamma_\lambda^k(\mathbf{i}) = \lambda g(k, \mathbf{i}(k)) + (1-\lambda) \sum_{k' \in K} q(k'|k, \mathbf{i}(k)) \gamma_\lambda^{k'}(\mathbf{i}).$$

Can be written in matrix form: $(I - (1-\lambda)A)v = \lambda \alpha$, where $v = (\gamma_\lambda^k(\mathbf{i}))_k$ and $(I - (1-\lambda)A)$ is invertible. So for each $\mathbf{i}$ and k, the payoff $\gamma_\lambda^k(\mathbf{i})$ is a rational function of λ.

Theorem (Blackwell, 1962): In the 1-player case, there exists $\lambda_0 > 0$ and a pure stationary strategy which is optimal in any game with discount $\lambda \leq \lambda_0$. For $\lambda \leq \lambda_0$ and k in K, the value $v_\lambda(k)$ is a bounded rational fraction of λ, hence converges when λ goes to 0.

Example 2 again: **i** is the strategy which alternates forever between k_1 and k_4. There exists no strategy which is exactly optimal in all n-stage games with n sufficiently large.

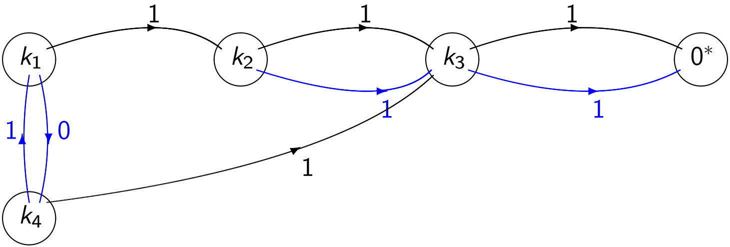

2) Stochastic games: The algebraic approach

For each λ, players have stationary optimal strategies x_λ and y_λ:
Consider the following set:

$$A = \{(\lambda, x_\lambda, y_\lambda, w_\lambda) \in (0,1] \times (I\!R^I)^K \times (I\!R^J)^K \times I\!R^K, \forall k \in K, \\ x_\lambda(k), y_\lambda(k) \text{ stationary optimal in } \Gamma_\lambda(k),\ w_\lambda(k) = v_\lambda(k)\}.$$

A can be written with finitely many polynomial inequalities:

$$\forall i,j,k,\ \sum_i x_\lambda^i(k) = 1, x_\lambda^i(k) \geq 0, \sum_j y_\lambda^j(k) = 1, y_\lambda^j(k) \geq 0,$$

$$\forall j,k,\ \sum_{i \in I} x_\lambda^i(k)(\lambda g(k,i,j) + (1-\lambda)\sum_{k'} q(k'|k,i,j) w_\lambda(k')) \geq w_\lambda(k),$$

$$\forall i,k,\ \sum_{j \in J} y_\lambda^j(k)(\lambda g(k,i,j) + (1-\lambda)\sum_{k'} q(k'|k,i,j) w_\lambda(k')) \leq w_\lambda(k).$$

A is semi-algebraic (can be written a finite union of sets, each of these sets being defined as the conjunction of finitely many weak or strict polynomial inequalities).

The projection of a semi-algebraic set is still semi-algebraic (Tarski-Seidenberg elimination theorem).
So $A^* = \{(\lambda, v_\lambda), \lambda \in (0,1]\}$ is also a semi-algebraic subset of $IR \times IR^K$.
Implies the existence of a bounded Puiseux series development of v_λ in a neighborhood of $\lambda = 0$.
Theorem (Bewley Kohlberg 1976):
There exists $\lambda_0 > 0$, a positive integer M, coefficients $r_m \in IR^K$ for each $m \geq 0$ such that for all $\lambda \in (0, \lambda_0]$, and all k in K:

$$v_\lambda(k) = \sum_{m=0}^{\infty} r_m(k)\, \lambda^{m/M}.$$

Example 1 : $v_\lambda = \frac{1}{1+\sqrt{\lambda}} = (1-\sqrt{\lambda})(1+\lambda+...+\lambda^n+....)$

Corollary: v_λ converges when λ goes to 0.

v_n also converges, and $\lim_{n\to\infty} v_n = \lim_{\lambda\to 0} v_\lambda$. (Uniform convergence of (v_n) and $(v_\lambda)_\lambda$ are equivalent in general stochastic games, Ziliotto 2016)

3) Stochastic games: a recent simple approach

Given stationary strategies $x : K \to \Delta(I)$ and $y : K \to \Delta(J)$, the λ-discounted payoff $\gamma_\lambda^k(x,y)$ is **not** bilinear in x and y, so in general $v_\lambda^k \neq \mathrm{val}(\gamma_\lambda^k(\mathbf{i},\mathbf{j}))_{\mathbf{i},\mathbf{j}}$, where $\mathbf{i}$ and $\mathbf{j}$ are pure stationary strategies.

However $\gamma_\lambda^k(x,y)$ can be computed via a Cramer linear system, with denominator:

$$\Delta_\lambda^0(x,y) = \det(Id - (1-\lambda)Q(x,y)) \geq \lambda^{|K|} > 0$$

with $Q(x,y)$ the stochastic matrix on K induced by x and y. And both $\Delta_\lambda^0(x,y)$ and $\Delta_\lambda^0(x,y)\gamma_\lambda^k(x,y)$ are bilinear in x, y!

Optimal x satisfies: $\forall \mathbf{j}, \Delta_\lambda^0(x,\mathbf{j})(\gamma_\lambda^k(x,\mathbf{j}) - v_\lambda^k) \geq 0$.
Optimal y satisfies: $\forall \mathbf{i}, \Delta_\lambda^0(\mathbf{i},y)(\gamma_\lambda^k(\mathbf{i},y) - v_\lambda^k) \leq 0$

Theorem (Atia and Oliu-Barton, 2019) : $\forall k$, v_λ^k is the unique solution of the equation:

$$0 = \mathrm{val}(\Delta_\lambda^0(\mathbf{i},\mathbf{j})(\gamma_\lambda^k(\mathbf{i},\mathbf{j}) - z))_{\mathbf{i},\mathbf{j}}.$$

Corollary: (Atia-Oliu-Barton 2019) v_λ^k converges to the unique $v \in I\!R$ such that for all $z \in I\!R$:

$$\begin{cases} (z > v) & \Longrightarrow \quad \lim_{\lambda \to 0} \frac{1}{\lambda^{|K|}} \mathrm{val}(\Delta_\lambda^0(\mathbf{i},\mathbf{j})(\gamma_\lambda^k(\mathbf{i},\mathbf{j}) - z))_{\mathbf{i},\mathbf{j}} < 0, \quad \textit{and} \\ (z < v) & \Longrightarrow \quad \lim_{\lambda \to 0} \frac{1}{\lambda^{|K|}} \mathrm{val}(\Delta_\lambda^0(\mathbf{i},\mathbf{j})(\gamma_\lambda^k(\mathbf{i},\mathbf{j}) - z))_{\mathbf{i},\mathbf{j}} > 0. \end{cases}$$

Corollary: (Atia-Oliu-Barton 2019) v_λ^k converges to the unique $v \in I\!R$ such that for all $z \in I\!R$:

$$\begin{cases} (z > v) \implies \lim_{\lambda \to 0} \frac{1}{\lambda^{|K|}} \mathrm{val}(\Delta_\lambda^0(\mathbf{i},\mathbf{j})(\gamma_\lambda^k(\mathbf{i},\mathbf{j}) - z))_{\mathbf{i},\mathbf{j}} < 0, \quad \textit{and} \\ (z < v) \implies \lim_{\lambda \to 0} \frac{1}{\lambda^{|K|}} \mathrm{val}(\Delta_\lambda^0(\mathbf{i},\mathbf{j})(\gamma_\lambda^k(\mathbf{i},\mathbf{j}) - z))_{\mathbf{i},\mathbf{j}} > 0. \end{cases}$$

2.1 Beyond the standard model: the 1-player case

Consider a stochastic dynamic programing problem $\Gamma = (X, F, r, x_0)$ given by a non empty set of states X, an initial state x_0, a transition multifunction F from X to $Z := \Delta_f(X)$ with non empty values, and a reward mapping r from X to $[0,1]$.

Assume X compact metric. Then $\Delta(X)$ is also compact metric space for the Kantorovich-Rubinstein metric: for u, u' in $\Delta(X)$,

$$\begin{aligned} d_{KR}(u,u') &= \sup_{f:X \to I\!R, 1-Lip} \left| \int_{x \in X} f(x)du(x) - \int_{x \in X} f(x)du'(x) \right| \\ &= \min_{\pi \in \Pi(u,u')} \int_{(x,x') \in X \times X} d(x,x')d\pi(x,x'). \end{aligned}$$

X is now viewed as a subset of $\Delta(X)$, and we define the set of *invariant measures* as:

$$R = \{u \in \Delta(X), (u,u) \in \overline{\text{conv}}\text{Graph}(F)\}$$

Theorem (R-Venel 2017) Assume the state space is compact metric, payoffs are continuous and transitions are non expansive for d_{KR}. Then (v_n) and (v_λ) uniformly converge to v^*, where for each initial state x,

$$v^*(x) = \inf\{w(x), w : \Delta(X) \to [0,1] \text{ affine } C^0 \; s.t.$$

$$(1)\ \forall x' \in X, w(x') \geq \sup_{u \in F(x')} w(u)$$

$$(2)\ \forall u \in R, w(u) \geq r(u)\}.$$

2.2 A compact continuous game with no limit value.

Finite set of states, compact action sets, continuous transitions.

First counterexample: Vigeral (2013), with non semi-algebraic transitions (Bolte, Gaubert, Vigeral 2015).

Following counterexample : polynomial transitions but non semi-algebraic action sets: variant of a counter-example of Ziliotto (2016), also mentioned in Sorin Vigeral (2015). The elementary proof here follows Renault (2015).

4 states: $K = \{k_0, k_1, 0^*, 1^*\}$.
In state k_1 player 2 chooses β in $J = [0, 1/2]$.
In state k_0, Player 1 chooses $\alpha \in I = \{0\} \cup \{\frac{1}{4^n}, n \geq 1\}$

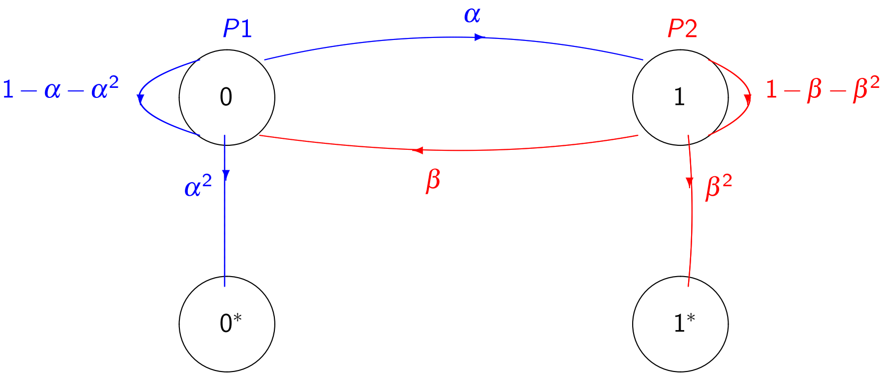

Write $x_\lambda = v_\lambda(k_0)$, $y_\lambda = v_\lambda(k_1)$.
Shapley equations:

$$\begin{aligned} x_\lambda &= \max_{\alpha\in I}(1-\lambda)((1-\alpha-\alpha^2)x_\lambda+\alpha y_\lambda), \\ y_\lambda &= \min_{\beta\in J}\left(\lambda+(1-\lambda)((1-\beta-\beta^2)y_\lambda+\beta x_\lambda+\beta^2)\right). \end{aligned}$$

Can be rewriten:

$$\lambda x_\lambda = (1-\lambda)\max_{\alpha\in I}\left(\alpha(y_\lambda-x_\lambda)-\alpha^2 x_\lambda\right) \quad (1)$$

$$\lambda y_\lambda = \lambda+(1-\lambda)\min_{\beta\in J}(\beta(x_\lambda-y_\lambda)+\beta^2(1-y_\lambda)) \quad (2)$$

Since $x_\lambda > 0$, eq. (**??**) gives that $y_\lambda > x_\lambda$.

Lemma 0: For $\lambda \le 1/5$, $\beta_\lambda = \frac{y_\lambda - x_\lambda}{2(1-y_\lambda)}$ is optimal for player 2 and

$$4\lambda(1-y_\lambda)^2 = (1-\lambda)(y_\lambda - x_\lambda)^2. \quad (3)$$

Consequence: $y_\lambda - x_\lambda \longrightarrow_{\lambda \to 0} 0$.
Let λ_n be a vanishing sequence of discount factors.

Lemma 1: if y_{λ_n} and x_{λ_n} converge to v in $[0,1]$, then $v \leq 1/2$, $y_{\lambda_n} - x_{\lambda_n} \sim 2\sqrt{\lambda_n}(1-v)$ and $\beta_{\lambda_n} \sim \sqrt{\lambda_n}$.

Lemma 2: If for each n, $\sqrt{\lambda_n} \in I$, then y_{λ_n} and x_{λ_n} converge to $1/2$.

Lemma 3: If for each n, the open interval $(\frac{1}{2}\sqrt{\lambda_n}, 2\sqrt{\lambda_n})$ does not intersect I, then $\limsup_n y_{\lambda_n} \leq 4/9$.

Considering the sequences $\lambda_n = \frac{1}{2^{2n}}$ and $\lambda_n = \frac{1}{2^{2n+1}}$ is enough to conclude that there is no limit value.

2.2 A hidden stochastic game with no limit value

Hidden stochastic games: at the beginning of each period, players observe past actions and a public signal (but no longer the current state). Stochastic games with public information.

Model given by: finite sets of states K, of actions I for player 1, of actions J for player 2, and of signals S, a payoff function $g : K \times I \times J \longrightarrow I\!R$, and a transition $q : K \times I \times J \longrightarrow \Delta(K \times S)$.

Ziliotto (2016) constructed a hidden stochastic game with no limit value. ($\liminf v_\delta = 1/2, \limsup v_\delta \geq 5/9$). Disproves 2 conjectures of J.F. Mertens

Here: (R. Ziliotto 2020):

Theorem: For each $\varepsilon > 0$, there exists a zero-sum HSG with payoffs in $[0,1]$ for P1, 6 states, 2 actions for each player, 6 signals, s.t.:

$$\limsup_{\lambda \to 0} v_\lambda \geq 1 - \varepsilon \text{ and } \liminf_{\lambda \to 0} v_\lambda \leq \varepsilon.$$

Construction done in 4 progressive steps: a Markov chain on $[0,1]$, a Markov Decision Process, a stochastic game with infinite state space, and a final HSG.

Step 1: a Markov chain on $[0,1]$, with a parameter $\alpha \in (0,1/4)$.
Initial state $q_0 = 1$.

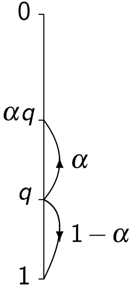

Define for a in $I\!N$, $T_a = \inf\{t \geq 1, q_t \leq \alpha^a\}$.

$$I\!E(T_{a+1}) = \frac{1}{\alpha}(1 + I\!E(T_a)) \text{ (grows exponentially with } a\text{)}$$

Step 2: a Markov Decision Process on $[0,1]$

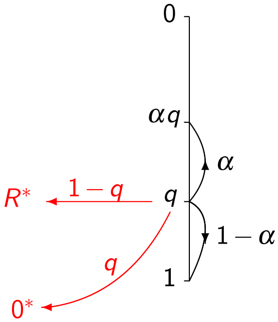

Recall $T_a = \inf\{t \geq 1, q_t \leq \alpha^a\}$.

Payoff of the a-strategy in the MDP with parameter α, reward R and discount δ: $R\, s_{\alpha,\delta}(a)$, with

$$s_{\alpha,\delta}(a) = \frac{(1-\alpha^a)(1-\alpha\delta)}{1-\alpha+(1-\delta)\alpha^{-a}\delta^{-a-1}}.$$

(optimal strategies do not depend on R)

For $R = 1$, the value is:

$$v_{\alpha,\delta} = \mathrm{Max}_{a \in I\!N} s_{\alpha,\delta}(a) = \mathrm{Max}_{a \in I\!N} \frac{(1-\alpha^a)(1-\alpha\delta)}{1-\alpha+(1-\delta)\alpha^{-a}\delta^{-a-1}} \xrightarrow[\delta \to 1]{} 1.$$

Optimal choice for $a \in I\!R_+$ would be $a^* = a^*(\alpha,\delta)$ s.t. $\alpha^{a^*} \simeq \sqrt{\frac{1-\delta}{1-\alpha}}$.
Define $\Delta_1(\alpha) = \{\delta, a^* \in I\!N\} = \{1-(1-\alpha)\alpha^{2a}, a \in I\!N\}$,
and $\Delta_2(\alpha) = \{\delta, a^* \in I\!N + [1/4, 3/4]\}$.

Proposition:
For $\delta \in \Delta_1(\alpha)$, $v_{\alpha,\delta} = 1 - \frac{2}{\sqrt{1-\alpha}}\sqrt{1-\delta} + o(\sqrt{1-\delta})$.
For $\delta \in \Delta_2(\alpha)$, $v_{\alpha,\delta} \leq 1 - \frac{1}{\sqrt{\alpha^{1/2}(1-\alpha)}}\sqrt{1-\delta} + o(\sqrt{1-\delta})$.

Step 3: a stochastic game $\Gamma_{\alpha,\beta}$ with perfect information

States: $X = \{(1,q), q \in [0,1]\} \cup \{(2,l), l \in [0,1]\} \cup 0^* \cup 1^*$, start at $(2,1)$. Sum of payoffs is 1, P1 has payoff 0 in the left part, payoff 1 in the right part.

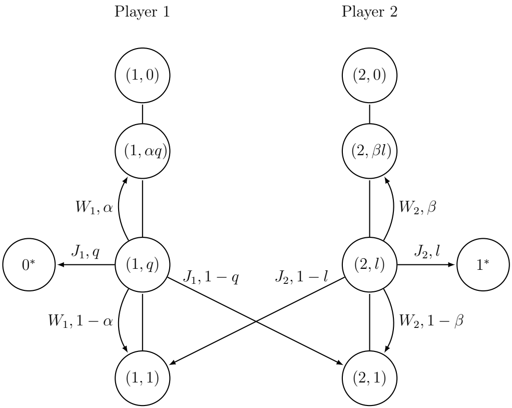

Proposition: The stochastic game restricted to pure stationary strategies has an equilibrium in *dominant strategies*, and value:

$$v_{\alpha,\beta,\delta} = \frac{1 - v_{\beta,\delta}}{1 - v_{\alpha,\delta} v_{\beta,\delta}}.$$

Proposition: Fix $\varepsilon > 0$. For n large enough, fixing $\alpha = 1/n$ and $\beta = 1/(n+1)$ yields:

$$\limsup_{\delta \to 1} v_{\alpha,\beta,\delta} \geq 1 - \varepsilon, \text{ and } \liminf_{\delta \to 1} v_{\alpha,\beta,\delta} \leq \varepsilon.$$

Step 4: a constant-sum hidden stochastic game $\Gamma^*_{\alpha,\beta}$

with 6 states and 6 signals, initial state: (2,1).

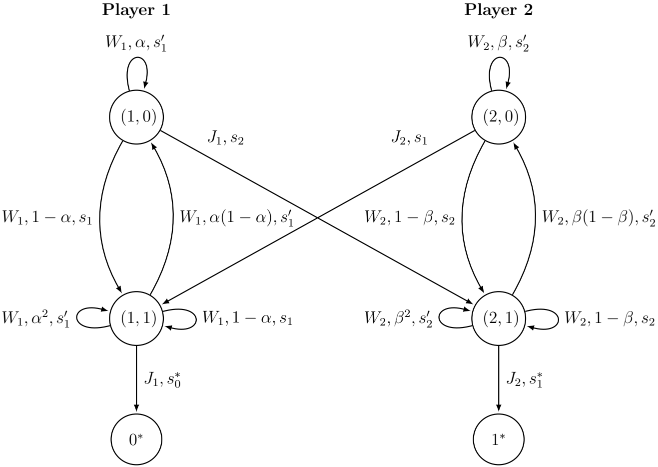

Value $v^*_{\alpha,\beta,\delta} = v_{\alpha,\beta,\delta}$.

2.4 A non zero-sum stochastic game with no limit equilibrium payoff set

Like the standard model, except that each player $i = 1, 2$ now has his own payoff function $g_i : K \times I \times J \to I\!R$.

For each λ, there exists a Nash equilibrium of the λ-discounted game in stationary strategies.

R-Ziliotto 2020: Using the algebraic approach, one can easily prove that the set E''_λ of stationary λ-discounted equilibrium payoffs converges to a non empty compact set of $I\!R^2$.
But there are examples where the set E_λ of λ-discounted equilibrium payoffs does not converge when $\lambda \to 0$.

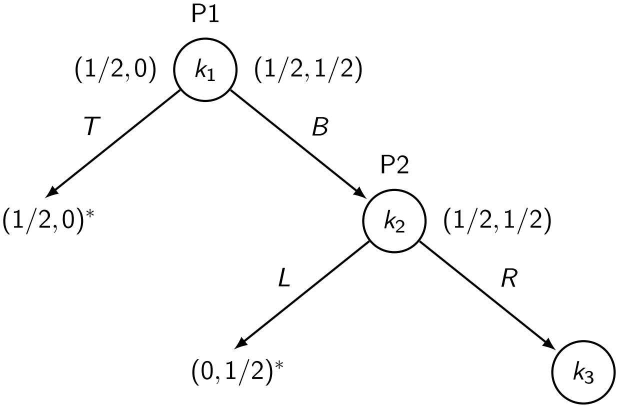

$$\begin{array}{c} \\ T \\ B \end{array} \begin{array}{c} \begin{array}{cc} L & R \end{array} \\ \begin{pmatrix} (1,0) \circlearrowleft & (-1,-1)^* \\ (-1,-1)^* & (0,1)^* \end{pmatrix} \end{array}$$

For all λ, $(1/2, 0) \in E_\lambda$.

If at eq. k_3 is reached, the payoff has to be $(1/2, 1/2)$, and $\lambda \in L = \{1 - (\frac{1}{2})^{1/N}, N \in I\!N_+\}$.

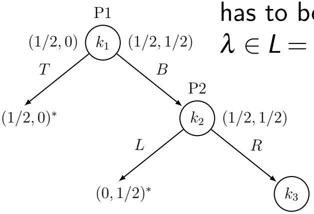

$$\begin{array}{c} \\ T \\ B \end{array} \begin{array}{cc} L & R \\ \end{array}$$

	L	R
T	$(1,0)\circlearrowleft$	$(-1,-1)^*$
B	$(-1,-1)^*$	$(0,1)^*$

For $\lambda \notin L, E_\lambda = \{(1/2, 0)\}$.
For $\lambda \in L, E_\lambda = 1/2 \times [0, 1/2]$.

Concluding Remarks

If we leave the standard model, the limit value may fail to exist.

However, there are positive results in the ergodic case, or in the acyclic case (Laraki, R 2020). Zero-sum counter-examples seems to be due to "nasty cycles".

Several other interesting notions of value not mentioned here (uniform value, limiting average, Borel payoff functions on plays...)

n-player quitting games: at each stage, each player decides to stop or continue. Whenever at least one player stops, the game is absorbed and each player receives a payoff depending on the set of stopping players. Payoff is 0 for everyone if no one ever stops.

Example: $n = 4$

The game is defined by 2^{n-1} vector payoff in $I\!R^n$ Is it true that for each $\varepsilon > 0$, this game has a ε-Nash equilibrium? Open as soon as $n \geq 4$