Convergence of Discrete-Time Games with Imperfect Information to a Continuous-Time Limit

Yuliy Sannikov

Working Paper, February 2006.

Presented by Tomek Piskorski

Motivation and Objective of the Paper

Motivation and Objective of the Paper

- Abreu, Pearce and Stacchetti (1990) showed that any equilibrium payoffs vector of discrete-time repeated game with imperfectly observable actions can be achieved by a recursive equilibrium with players' continuation values as state variables.

Motivation and Objective of the Paper

- Abreu, Pearce and Stacchetti (1990) showed that any equilibrium payoffs vector of discrete-time repeated game with imperfectly observable actions can be achieved by a recursive equilibrium with players' continuation values as state variables.
- In a recursive equilibrium continuation values change location after every observation of the public signal.

Motivation and Objective of the Paper

- Abreu, Pearce and Stacchetti (1990) showed that any equilibrium payoffs vector of discrete-time repeated game with imperfectly observable actions can be achieved by a recursive equilibrium with players' continuation values as state variables.
- In a recursive equilibrium continuation values change location after every observation of the public signal.
- It is not straightforward to see a pattern behind these jumps and the connection between the equilibrium dynamics and the shape of the set of equilibrium payoffs.

Motivation and Objective of the Paper

- Abreu, Pearce and Stacchetti (1990) showed that any equilibrium payoffs vector of discrete-time repeated game with imperfectly observable actions can be achieved by a recursive equilibrium with players' continuation values as state variables.
- In a recursive equilibrium continuation values change location after every observation of the public signal.
- It is not straightforward to see a pattern behind these jumps and the connection between the equilibrium dynamics and the shape of the set of equilibrium payoffs.
- Sannikov (2005) modeling a repeated game in continuous time showed that the set of equilibrium payoffs is characterized by an ODE and that the pattern of interaction that achieve extreme payoffs pairs is **uniquely** determined by the public signals.

Motivation and Objective of the Paper

- Abreu, Pearce and Stacchetti (1990) showed that any equilibrium payoffs vector of discrete-time repeated game with imperfectly observable actions can be achieved by a recursive equilibrium with players' continuation values as state variables.
- In a recursive equilibrium continuation values change location after every observation of the public signal.
- It is not straightforward to see a pattern behind these jumps and the connection between the equilibrium dynamics and the shape of the set of equilibrium payoffs.
- Sannikov (2005) modeling a repeated game in continuous time showed that the set of equilibrium payoffs is characterized by an ODE and that the pattern of interaction that achieve extreme payoffs pairs is **uniquely** determined by the public signals.
- Does the set of payoffs of equilibria of a discrete-time game converge to the set of payoffs of equilibria of a continuous time game as the length between periods goes to zero?

Discrete Time Game: Setup

Discrete Time Game: Setup

Time: Discrete, Infinite, Δ is the period's length.

Discrete Time Game: Setup

Time: Discrete, Infinite, Δ is the period's length.

The two-player stage game is repeated at time points $\{0, \Delta, 2\Delta, \ldots\}$.

Discrete Time Game: Setup

Time: Discrete, Infinite, Δ is the period's length.

The two-player stage game is repeated at time points $\{0, \Delta, 2\Delta, \dots\}$.

Two players $(i = 1,2)$ choose actions from finite action sets $\mathcal{A}_1$ and $\mathcal{A}_2$.

Discrete Time Game: Setup

Time: Discrete, Infinite, Δ is the period's length.

The two-player stage game is repeated at time points $\{0, \Delta, 2\Delta, \ldots\}$.

Two players $(i = 1,2)$ choose actions from finite action sets $\mathcal{A}_1$ and $\mathcal{A}_2$.

The actions the player takes are his private information.

Discrete Time Game: Setup

Time: Discrete, Infinite, Δ is the period's length.

The two-player stage game is repeated at time points $\{0, \Delta, 2\Delta, \dots\}$.

Two players $(i = 1,2)$ choose actions from finite action sets $\mathcal{A}_1$ and $\mathcal{A}_2$.

The actions the player takes are his private information.

A profile of actions $a \in \mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$ determines the probability distribution $\pi^{\Delta}(a)$ over the set of values of public signal $Y = \{1, .., d\}$.

Discrete Time Game: Setup

Time: Discrete, Infinite, Δ is the period's length.

The two-player stage game is repeated at time points $\{0, \Delta, 2\Delta, \dots\}$.

Two players $(i = 1,2)$ choose actions from finite action sets $\mathcal{A}_1$ and $\mathcal{A}_2$.

The actions the player takes are his private information.

A profile of actions $a \in \mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$ determines the probability distribution $\pi^{\Delta}(a)$ over the set of values of public signal $Y = \{1, .., d\}$.

The vector of signal probabilities is determined relative to a fixed strictly positive distribution π as

Discrete Time Game: Setup

Time: Discrete, Infinite, Δ is the period's length.

The two-player stage game is repeated at time points $\{0, \Delta, 2\Delta, \ldots\}$.

Two players $(i = 1,2)$ choose actions from finite action sets $\mathcal{A}_1$ and $\mathcal{A}_2$.

The actions the player takes are his private information.

A profile of actions $a \in \mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$ determines the probability distribution $\pi^{\Delta}(a)$ over the set of values of public signal $Y = \{1, .., d\}$.

The vector of signal probabilities is determined relative to a fixed strictly positive distribution π as

$$\pi^{\Delta}(a) = \pi + \sqrt{\Delta}\,\mu(a)$$

Discrete Time Game: Setup

Time: Discrete, Infinite, Δ is the period's length.

The two-player stage game is repeated at time points $\{0, \Delta, 2\Delta, \ldots\}$.

Two players $(i = 1,2)$ choose actions from finite action sets $\mathcal{A}_1$ and $\mathcal{A}_2$.

The actions the player takes are his private information.

A profile of actions $a \in \mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$ determines the probability distribution $\pi^{\Delta}(a)$ over the set of values of public signal $Y = \{1, .., d\}$.

The vector of signal probabilities is determined relative to a fixed strictly positive distribution π as

$$\pi^{\Delta}(a) = \pi + \sqrt{\Delta}\,\mu(a)$$

where $\mu(a)$ is a d-dimensional vector, whose components add up to 0.

Discrete Time Game: Setup

Time: Discrete, Infinite, Δ is the period's length.

The two-player stage game is repeated at time points $\{0, \Delta, 2\Delta, \dots\}$.

Two players $(i = 1,2)$ choose actions from finite action sets $\mathcal{A}_1$ and $\mathcal{A}_2$.

The actions the player takes are his private information.

A profile of actions $a \in \mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$ determines the probability distribution $\pi^{\Delta}(a)$ over the set of values of public signal $Y = \{1, .., d\}$.

The vector of signal probabilities is determined relative to a fixed strictly positive distribution π as

$$\pi^{\Delta}(a) = \pi + \sqrt{\Delta}\,\mu(a)$$

where $\mu(a)$ is a d-dimensional vector, whose components add up to 0.

Assume that $\pi_k^{\Delta}(a)$ is strictly positive for all $\Delta \in [0.1]$, $a \in \mathcal{A}$, and $k = 1, .., d$.

Discrete Time Game: Setup

Discrete Time Game: Setup

The realized flow of player i's payoffs depends on his action and the public signal.

Discrete Time Game: Setup

The realized flow of player i's payoffs depends on his action and the public signal.

The player i receives this flow of payoff for an interval of time Δ until the next repetition of the stage game.

Discrete Time Game: Setup

The realized flow of player i's payoffs depends on his action and the public signal.

The player i receives this flow of payoff for an interval of time Δ until the next repetition of the stage game.

The payoff flows for action a_i and each signal value are summarized in a vector

Discrete Time Game: Setup

The realized flow of player i's payoffs depends on his action and the public signal.

The player i receives this flow of payoff for an interval of time Δ until the next repetition of the stage game.

The payoff flows for action a_i and each signal value are summarized in a vector

$$f_i^{\Delta}(a_i) = c_i(a_i)\mathbf{1} + \frac{d_i(a_i)}{\sqrt{\Delta}}$$

Discrete Time Game: Setup

The realized flow of player i's payoffs depends on his action and the public signal.

The player i receives this flow of payoff for an interval of time Δ until the next repetition of the stage game.

The payoff flows for action a_i and each signal value are summarized in a vector

$$f_i^{\Delta}(a_i) = c_i(a_i)\mathbf{1} + \frac{d_i(a_i)}{\sqrt{\Delta}}$$

where $c_i(a_i) \in R$, $\mathbf{1} \in R^d$ is a vector of ones, and $d_i(a_i)$ is orthogonal to π.

Discrete Time Game: Setup

The realized flow of player i's payoffs depends on his action and the public signal.

The player i receives this flow of payoff for an interval of time Δ until the next repetition of the stage game.

The payoff flows for action a_i and each signal value are summarized in a vector

$$f_i^{\Delta}(a_i) = c_i(a_i)\mathbf{1} + \frac{d_i(a_i)}{\sqrt{\Delta}}$$

where $c_i(a_i) \in R$, $\mathbf{1} \in R^d$ is a vector of ones, and $d_i(a_i)$ is orthogonal to π.

The expected stage-game flow payoff for player i is given by

Discrete Time Game: Setup

The realized flow of player i's payoffs depends on his action and the public signal.

The player i receives this flow of payoff for an interval of time Δ until the next repetition of the stage game.

The payoff flows for action a_i and each signal value are summarized in a vector

$$f_i^{\Delta}(a_i) = c_i(a_i)\mathbf{1} + \frac{d_i(a_i)}{\sqrt{\Delta}}$$

where $c_i(a_i) \in R$, $\mathbf{1} \in R^d$ is a vector of ones, and $d_i(a_i)$ is orthogonal to π.

The expected stage-game flow payoff for player i is given by

$$g_i(a) = \left(c_i(a_i)\mathbf{1} + \frac{d_i(a_i)}{\sqrt{\Delta}} \right) \cdot \left(\pi + \sqrt{\Delta}\,\mu(a) \right) = c_i(a_i) + d_i(a_i) \cdot \mu(a)$$

Discrete Time Game: Setup

The realized flow of player i's payoffs depends on his action and the public signal.

The player i receives this flow of payoff for an interval of time Δ until the next repetition of the stage game.

The payoff flows for action a_i and each signal value are summarized in a vector

$$f_i^{\Delta}(a_i) = c_i(a_i)\mathbf{1} + \frac{d_i(a_i)}{\sqrt{\Delta}}$$

where $c_i(a_i) \in R$, $\mathbf{1} \in R^d$ is a vector of ones, and $d_i(a_i)$ is orthogonal to π.

The expected stage-game flow payoff for player i is given by

$$g_i(a) = \left(c_i(a_i)\mathbf{1} + \frac{d_i(a_i)}{\sqrt{\Delta}} \right) \cdot \left(\pi + \sqrt{\Delta}\,\mu(a) \right) = c_i(a_i) + d_i(a_i) \cdot \mu(a)$$

Each player cares about his expected lifetime payoff:

Discrete Time Game: Setup

The realized flow of player i's payoffs depends on his action and the public signal.

The player i receives this flow of payoff for an interval of time Δ until the next repetition of the stage game.

The payoff flows for action a_i and each signal value are summarized in a vector

$$f_i^{\Delta}(a_i) = c_i(a_i)\mathbf{1} + \frac{d_i(a_i)}{\sqrt{\Delta}}$$

where $c_i(a_i) \in R$, $\mathbf{1} \in R^d$ is a vector of ones, and $d_i(a_i)$ is orthogonal to π.

The expected stage-game flow payoff for player i is given by

$$g_i(a) = \left(c_i(a_i)\mathbf{1} + \frac{d_i(a_i)}{\sqrt{\Delta}} \right) \cdot \left(\pi + \sqrt{\Delta}\,\mu(a) \right) = c_i(a_i) + d_i(a_i) \cdot \mu(a)$$

Each player cares about his expected lifetime payoff:

$$E\left[(1 - e^{-r\Delta}) \sum_{n=0}^{\infty} e^{-rn\Delta} g_i(a_{n\Delta}) \right]$$

Discrete Time Game: Setup

The realized flow of player i's payoffs depends on his action and the public signal.

The player i receives this flow of payoff for an interval of time Δ until the next repetition of the stage game.

The payoff flows for action a_i and each signal value are summarized in a vector

$$f_i^{\Delta}(a_i) = c_i(a_i)\mathbf{1} + \frac{d_i(a_i)}{\sqrt{\Delta}}$$

where $c_i(a_i) \in R$, $\mathbf{1} \in R^d$ is a vector of ones, and $d_i(a_i)$ is orthogonal to π.

The expected stage-game flow payoff for player i is given by

$$g_i(a) = \left(c_i(a_i)\mathbf{1} + \frac{d_i(a_i)}{\sqrt{\Delta}} \right) \cdot \left(\pi + \sqrt{\Delta}\,\mu(a) \right) = c_i(a_i) + d_i(a_i) \cdot \mu(a)$$

Each player cares about his expected lifetime payoff:

$$E\left[(1 - e^{-r\Delta}) \sum_{n=0}^{\infty} e^{-rn\Delta} g_i(a_{n\Delta}) \right]$$

where r is a common discount factor.

Discrete Time Game:
Public Perfect Equilibrium Payoffs

Discrete Time Game:
Public Perfect Equilibrium Payoffs

Denote the repeated game with frequency of actions Δ by G_Δ.

Discrete Time Game:
Public Perfect Equilibrium Payoffs

Denote the repeated game with frequency of actions Δ by G_Δ.

In game G_Δ, a strategy of a player i is **public** if his action at time Δn is completely determined by the sequence of public signals until that time.

Discrete Time Game:
Public Perfect Equilibrium Payoffs

Denote the repeated game with frequency of actions Δ by G_Δ.

In game G_Δ, a strategy of a player i is **public** if his action at time Δn is completely determined by the sequence of public signals until that time.

A **public perfect equilibrium (PPE)** is a profile of public strategies, s.t. after any history of public signals, continuation strategies form a Nash equilibrium.

Discrete Time Game:
Public Perfect Equilibrium Payoffs

Denote the repeated game with frequency of actions Δ by G_Δ.

In game G_Δ, a strategy of a player i is **public** if his action at time Δn is completely determined by the sequence of public signals until that time.

A **public perfect equilibrium (PPE)** is a profile of public strategies, s.t. after any history of public signals, continuation strategies form a Nash equilibrium.

Denote by $\mathcal{E}_\Delta(r)$ the set of payoff pairs, achievable by PPE with public randomization of game G_Δ.

Discrete Time Game:
Public Perfect Equilibrium Payoffs

Denote the repeated game with frequency of actions Δ by G_Δ.

In game G_Δ, a strategy of a player i is **public** if his action at time Δn is completely determined by the sequence of public signals until that time.

A **public perfect equilibrium (PPE)** is a profile of public strategies, s.t. after any history of public signals, continuation strategies form a Nash equilibrium.

Denote by $\mathcal{E}_\Delta(r)$ the set of payoff pairs, achievable by PPE with public randomization of game G_Δ.

What is $\lim_{\Delta\to 0}\mathcal{E}_\Delta(r)$?

Continuous Time Game: Setup

Continuous Time Game: Setup

Time: Continuous, Infinite.

Continuous Time Game: Setup

Time: Continuous, Infinite.

Two players participate in a continuous-time game.

Continuous Time Game: Setup

Time: Continuous, Infinite.

Two players participate in a continuous-time game.

Consider a driftless d-dimensional Brownian motion $X = \{X_1, \ldots, X_d, \mathcal{F}_t, 0 \leq t < \infty\}$ with covariance matrix Σ^{π} on the probability space $\{\Omega, \mathcal{F}, Q\}$.

Continuous Time Game: Setup

Time: Continuous, Infinite.

Two players participate in a continuous-time game.

Consider a driftless d-dimensional Brownian motion $X = \{X_1, \ldots, X_d, \mathcal{F}_t, 0 \leq t < \infty\}$ with covariance matrix Σ^π on the probability space $\{\Omega, \mathcal{F}, Q\}$.

$$\Sigma^\pi = \begin{bmatrix} \pi_1(1-\pi_1) & -\pi_1\pi_2 & \ldots & -\pi_1\pi_d \\ -\pi_2\pi_1 & \pi_2(1-\pi_2) & \ldots & -\pi_2\pi_d \\ \vdots & \vdots & & \vdots \\ -\pi_d\pi_1 & -\pi_d\pi_2 & \ldots & \pi_d(1-\pi_d) \end{bmatrix}$$

Continuous Time Game: Setup

Time: Continuous, Infinite.

Two players participate in a continuous-time game.

Consider a driftless d-dimensional Brownian motion $X = \{X_1, \ldots, X_d, \mathcal{F}_t, 0 \leq t < \infty\}$ with covariance matrix Σ^π on the probability space $\{\Omega, \mathcal{F}, Q\}$.

$$\Sigma^\pi = \begin{bmatrix} \pi_1(1-\pi_1) & -\pi_1\pi_2 & \ldots & -\pi_1\pi_d \\ -\pi_2\pi_1 & \pi_2(1-\pi_2) & \ldots & -\pi_2\pi_d \\ \vdots & \vdots & & \vdots \\ -\pi_d\pi_1 & -\pi_d\pi_2 & \ldots & \pi_d(1-\pi_d) \end{bmatrix}$$

A strategy of player i is a progressively measurable process A_t^i with values in a finite action set $\mathcal{A}^i$. Let $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$.

Continuous Time Game: Setup

Time: Continuous, Infinite.

Two players participate in a continuous-time game.

Consider a driftless d-dimensional Brownian motion $X = \{X_1, \ldots, X_d, \mathcal{F}_t, 0 \leq t < \infty\}$ with covariance matrix Σ^π on the probability space $\{\Omega, \mathcal{F}, Q\}$.

$$\Sigma^\pi = \begin{bmatrix} \pi_1(1-\pi_1) & -\pi_1\pi_2 & \ldots & -\pi_1\pi_d \\ -\pi_2\pi_1 & \pi_2(1-\pi_2) & \ldots & -\pi_2\pi_d \\ \vdots & \vdots & & \vdots \\ -\pi_d\pi_1 & -\pi_d\pi_2 & \ldots & \pi_d(1-\pi_d) \end{bmatrix}$$

A strategy of player i is a progressively measurable process A_t^i with values in a finite action set $\mathcal{A}^i$. Let $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$.

The actions $a \in \mathcal{A}$ of the players modify the underlying probability measure to endow signal X with drift $\mu(a)$, that is

Continuous Time Game: Setup

Time: Continuous, Infinite.

Two players participate in a continuous-time game.

Consider a driftless d-dimensional Brownian motion $X = \{X_1, \ldots, X_d, \mathcal{F}_t, 0 \leq t < \infty\}$ with covariance matrix Σ^π on the probability space $\{\Omega, \mathcal{F}, Q\}$.

$$\Sigma^\pi = \begin{bmatrix} \pi_1(1-\pi_1) & -\pi_1\pi_2 & \ldots & -\pi_1\pi_d \\ -\pi_2\pi_1 & \pi_2(1-\pi_2) & \ldots & -\pi_2\pi_d \\ \vdots & \vdots & & \vdots \\ -\pi_d\pi_1 & -\pi_d\pi_2 & \ldots & \pi_d(1-\pi_d) \end{bmatrix}$$

A strategy of player i is a progressively measurable process A_t^i with values in a finite action set $\mathcal{A}^i$. Let $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$.

The actions $a \in \mathcal{A}$ of the players modify the underlying probability measure to endow signal X with drift $\mu(a)$, that is

$$dX_t = \mu(A_t)\, dt + dZ_t^A$$

Continuous Time Game: Setup

Time: Continuous, Infinite.

Two players participate in a continuous-time game.

Consider a driftless d-dimensional Brownian motion $X = \{X_1, \ldots, X_d, \mathcal{F}_t, 0 \leq t < \infty\}$ with covariance matrix Σ^π on the probability space $\{\Omega, \mathcal{F}, Q\}$.

$$\Sigma^\pi = \begin{bmatrix} \pi_1(1-\pi_1) & -\pi_1\pi_2 & \ldots & -\pi_1\pi_d \\ -\pi_2\pi_1 & \pi_2(1-\pi_2) & \ldots & -\pi_2\pi_d \\ \vdots & \vdots & & \vdots \\ -\pi_d\pi_1 & -\pi_d\pi_2 & \ldots & \pi_d(1-\pi_d) \end{bmatrix}$$

A strategy of player i is a progressively measurable process A_t^i with values in a finite action set $\mathcal{A}^i$. Let $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$.

The actions $a \in \mathcal{A}$ of the players modify the underlying probability measure to endow signal X with drift $\mu(a)$, that is

$$dX_t = \mu(A_t)\, dt + dZ_t^A$$

where Z^A is a d-dimensional Brownian motion with covariance matrix Σ^π under the modified measure Q^A.

Continuous Time Game: Setup and PPE Payoffs

Continuous Time Game: Setup and PPE Payoffs

The instantaneous payoff of a player depends on his action and public signal X.

Continuous Time Game: Setup and PPE Payoffs

The instantaneous payoff of a player depends on his action and public signal X.

Players discount their payoffs at a common rate r.

Continuous Time Game: Setup and PPE Payoffs

The instantaneous payoff of a player depends on his action and public signal X.

Players discount their payoffs at a common rate r.

The expected discounted average payoff (continuation value) of player i at time t is

Continuous Time Game: Setup and PPE Payoffs

The instantaneous payoff of a player depends on his action and public signal X.

Players discount their payoffs at a common rate r.

The expected discounted average payoff (continuation value) of player i at time t is

$$W_t^i(A) = E^A\left[r\int_t^\infty e^{-r(s-t)}\left(c_i(A_s^i)\,ds + d_i(A_s^i)\cdot dX_s\right) \mid \mathcal{F}_t\right]$$

Continuous Time Game: Setup and PPE Payoffs

The instantaneous payoff of a player depends on his action and public signal X.

Players discount their payoffs at a common rate r.

The expected discounted average payoff (continuation value) of player i at time t is

$$W_t^i(A) = E^A\left[r\int_t^\infty e^{-r(s-t)}\left(c_i(A_s^i)\,ds + d_i(A_s^i)\cdot dX_s\right) \mid \mathcal{F}_t\right]$$
$$= E^A\left[r\int_t^\infty e^{-r(s-t)} g_i(A_s)\,ds \mid \mathcal{F}_t\right]$$

Continuous Time Game: Setup and PPE Payoffs

The instantaneous payoff of a player depends on his action and public signal X.

Players discount their payoffs at a common rate r.

The expected discounted average payoff (continuation value) of player i at time t is

$$W_t^i(A) = E^A\left[r\int_t^\infty e^{-r(s-t)}\left(c_i(A_s^i)\,ds + d_i(A_s^i)\cdot dX_s\right) \mid \mathcal{F}_t\right]$$
$$= E^A\left[r\int_t^\infty e^{-r(s-t)} g_i(A_s)\,ds \mid \mathcal{F}_t\right]$$

E^A denotes the expectation under the modified measure $\mathcal{Q}^A$.

Continuous Time Game: Setup and PPE Payoffs

The instantaneous payoff of a player depends on his action and public signal X.

Players discount their payoffs at a common rate r.

The expected discounted average payoff (continuation value) of player i at time t is

$$W_t^i(A) = E^A\left[r\int_t^\infty e^{-r(s-t)}\left(c_i(A_s^i)\,ds + d_i(A_s^i)\cdot dX_s\right) \mid \mathcal{F}_t\right]$$
$$= E^A\left[r\int_t^\infty e^{-r(s-t)} g_i(A_s)\,ds \mid \mathcal{F}_t\right]$$

E^A denotes the expectation under the modified measure $\mathcal{Q}^A$.

$c_i : \mathcal{A}^i \to \Re$ and $d_i : \mathcal{A}^i \to \Re^d$.

Continuous Time Game: Setup and PPE Payoffs

The instantaneous payoff of a player depends on his action and public signal X.

Players discount their payoffs at a common rate r.

The expected discounted average payoff (continuation value) of player i at time t is

$$W_t^i(A) = E^A\left[r\int_t^\infty e^{-r(s-t)}\left(c_i(A_s^i)\,ds + d_i(A_s^i)\cdot dX_s\right) \mid \mathcal{F}_t\right]$$
$$= E^A\left[r\int_t^\infty e^{-r(s-t)} g_i(A_s)\,ds \mid \mathcal{F}_t\right]$$

E^A denotes the expectation under the modified measure $\mathcal{Q}^A$.

$c_i : \mathcal{A}^i \to \Re$ and $d_i : \mathcal{A}^i \to \Re^d$.

A pair of strategies $A = (A^1, A^2)$ is a **perfect public equilibrium (PPE)** if for $i = 1,2$ A_i maximizes $W_t^i(A^i, A^j)$, A^j given after all public histories.

Continuous Time Game: Setup and PPE Payoffs

The instantaneous payoff of a player depends on his action and public signal X.

Players discount their payoffs at a common rate r.

The expected discounted average payoff (continuation value) of player i at time t is

$$W_t^i(A) = E^A\left[r\int_t^\infty e^{-r(s-t)}\left(c_i(A_s^i)\,ds + d_i(A_s^i)\cdot dX_s\right) \mid \mathcal{F}_t\right]$$
$$= E^A\left[r\int_t^\infty e^{-r(s-t)} g_i(A_s)\,ds \mid \mathcal{F}_t\right]$$

E^A denotes the expectation under the modified measure $\mathcal{Q}^A$.

$c_i : \mathcal{A}^i \to \Re$ and $d_i : \mathcal{A}^i \to \Re^d$.

A pair of strategies $A = (A^1, A^2)$ is a **perfect public equilibrium (PPE)** if for $i = 1,2$ A_i maximizes $W_t^i(A^i, A^j)$, A^j given after all public histories.

Let $\mathcal{E}(r)$ be the set of value pairs that can be achieved by some **PPE**.

Continuous Time Game:
Sannikov's (2005) Characterization of Set of PPE Payoffs

Continuous Time Game: Sannikov's (2005) Characterization of Set of PPE Payoffs

Let $\mathcal{N}$ be the convex hull of pure-strategy Nash equilibrium payoffs.

Continuous Time Game: Sannikov's (2005) Characterization of Set of PPE Payoffs

Let $\mathcal{N}$ be the convex hull of pure-strategy Nash equilibrium payoffs.

Sannikov (2005) showed that in a PPE with value $W_0 \in \partial\mathcal{E}(r)/\mathcal{N}$ the future continuation values of the players, W_t, are on the boundary of $\mathcal{E}(r)$.

Continuous Time Game: Sannikov's (2005) Characterization of Set of PPE Payoffs

Let $\mathcal{N}$ be the convex hull of pure-strategy Nash equilibrium payoffs.

Sannikov (2005) showed that in a PPE with value $W_0 \in \partial\mathcal{E}(r)/\mathcal{N}$ the future continuation values of the players, W_t, are on the boundary of $\mathcal{E}(r)$.

Locally the motion of continuation values is given by

Continuous Time Game: Sannikov's (2005) Characterization of Set of PPE Payoffs

Let $\mathcal{N}$ be the convex hull of pure-strategy Nash equilibrium payoffs.

Sannikov (2005) showed that in a PPE with value $W_0 \in \partial\mathcal{E}(r)/\mathcal{N}$ the future continuation values of the players, W_t, are on the boundary of $\mathcal{E}(r)$.

Locally the motion of continuation values is given by

$$dW_t = r(W_t - g(A_t))\,dt + r\mathbf{T}\beta_t \cdot (dX_t - \mu(A_t)\,dt)$$

Continuous Time Game: Sannikov's (2005) Characterization of Set of PPE Payoffs

Let $\mathcal{N}$ be the convex hull of pure-strategy Nash equilibrium payoffs.

Sannikov (2005) showed that in a PPE with value $W_0 \in \partial\mathcal{E}(r)/\mathcal{N}$ the future continuation values of the players, W_t, are on the boundary of $\mathcal{E}(r)$.

Locally the motion of continuation values is given by

$$dW_t = r(W_t - g(A_t))\,dt + r\mathbf{T}\beta_t \cdot (dX_t - \mu(A_t)\,dt)$$

where $W_t \in \partial\mathcal{E}(r)$ is a current pair of continuation values, A_t is a current action profile, and T is a unit tangent vector to the set $\mathcal{E}(r)$ at point W_t.

Continuous Time Game: Sannikov's (2005) Characterization of Set of PPE Payoffs

Let $\mathcal{N}$ be the convex hull of pure-strategy Nash equilibrium payoffs.

Sannikov (2005) showed that in a PPE with value $W_0 \in \partial\mathcal{E}(r)/\mathcal{N}$ the future continuation values of the players, W_t, are on the boundary of $\mathcal{E}(r)$.

Locally the motion of continuation values is given by

$$dW_t = \underbrace{r(W_t - g(A_t))}_{\text{PROMISE KEEPING}}\, dt + r\mathbf{T}\beta_t \cdot (dX_t - \mu(A_t)\, dt)$$

where $W_t \in \partial\mathcal{E}(r)$ is a current pair of continuation values, A_t is a current action profile, and T is a unit tangent vector to the set $\mathcal{E}(r)$ at point W_t.

Continuous Time Game: Sannikov's (2005) Characterization of Set of PPE Payoffs

Let $\mathcal{N}$ be the convex hull of pure-strategy Nash equilibrium payoffs.

Sannikov (2005) showed that in a PPE with value $W_0 \in \partial\mathcal{E}(r)/\mathcal{N}$ the future continuation values of the players, W_t, are on the boundary of $\mathcal{E}(r)$.

Locally the motion of continuation values is given by

$$dW_t = \underbrace{r(W_t - g(A_t))}_{\text{PROMISE KEEPING}}\, dt + \underbrace{r\mathbf{T}\beta_t \cdot (dX_t - \mu(A_t)\, dt)}_{\text{INCENTIVES}}$$

where $W_t \in \partial\mathcal{E}(r)$ is a current pair of continuation values, A_t is a current action profile, and T is a unit tangent vector to the set $\mathcal{E}(r)$ at point W_t.

Continuous Time Game: Sannikov's (2005) Characterization of Set of PPE Payoffs

Let $\mathcal{N}$ be the convex hull of pure-strategy Nash equilibrium payoffs.

Sannikov (2005) showed that in a PPE with value $W_0 \in \partial\mathcal{E}(r)/\mathcal{N}$ the future continuation values of the players, W_t, are on the boundary of $\mathcal{E}(r)$.

Locally the motion of continuation values is given by

$$dW_t = \underbrace{r(W_t - g(A_t))}_{\text{PROMISE KEEPING}}\, dt + \underbrace{r\mathbf{T}\beta_t \cdot \overbrace{(dX_t - \mu(A_t)\, dt)}^{dZ_t^A}}_{\text{INCENTIVES}}$$

where $W_t \in \partial\mathcal{E}(r)$ is a current pair of continuation values, A_t is a current action profile, and T is a unit tangent vector to the set $\mathcal{E}(r)$ at point W_t.

Continuous Time Game: Sannikov's (2005) Characterization of Set of PPE Payoffs

Let $\mathcal{N}$ be the convex hull of pure-strategy Nash equilibrium payoffs.

Sannikov (2005) showed that in a PPE with value $W_0 \in \partial\mathcal{E}(r)/\mathcal{N}$ the future continuation values of the players, W_t, are on the boundary of $\mathcal{E}(r)$.

Locally the motion of continuation values is given by

$$dW_t = \underbrace{r(W_t - g(A_t))}_{\text{PROMISE KEEPING}}\, dt + \underbrace{r\mathbf{T}\beta_t \cdot \overbrace{(dX_t - \mu(A_t)\, dt)}^{dZ_t^A}}_{\text{INCENTIVES}}$$

where $W_t \in \partial\mathcal{E}(r)$ is a current pair of continuation values, A_t is a current action profile, and T is a unit tangent vector to the set $\mathcal{E}(r)$ at point W_t.

There is a **unique** action profile played at each point of $\partial\mathcal{E}(r)/\mathcal{N}$ and β_t is also **unique**.

Continuous Time Game:
Sannikov's (2005) Characterization of Set of PPE Payoffs

Continuous Time Game:
Sannikov's (2005) Characterization of Set of PPE Payoffs

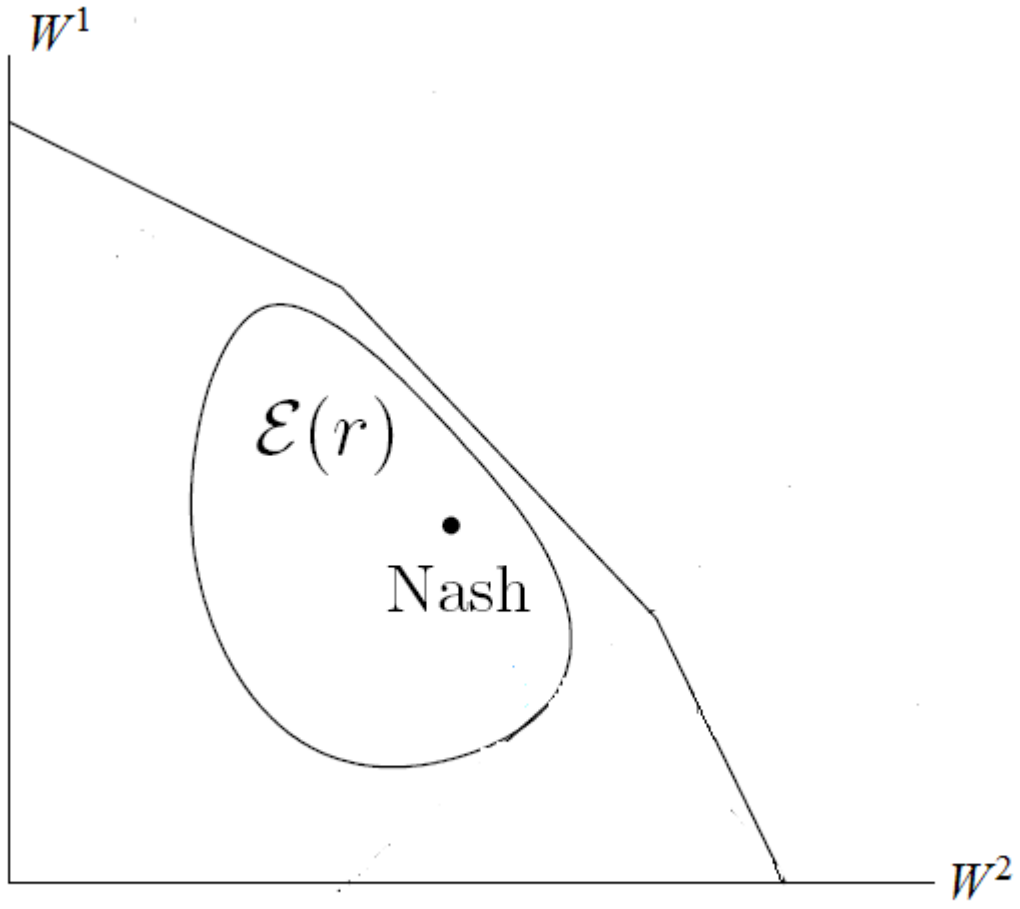

Continuous Time

Continuous Time Game:
Sannikov's (2005) Characterization of Set of PPE Payoffs

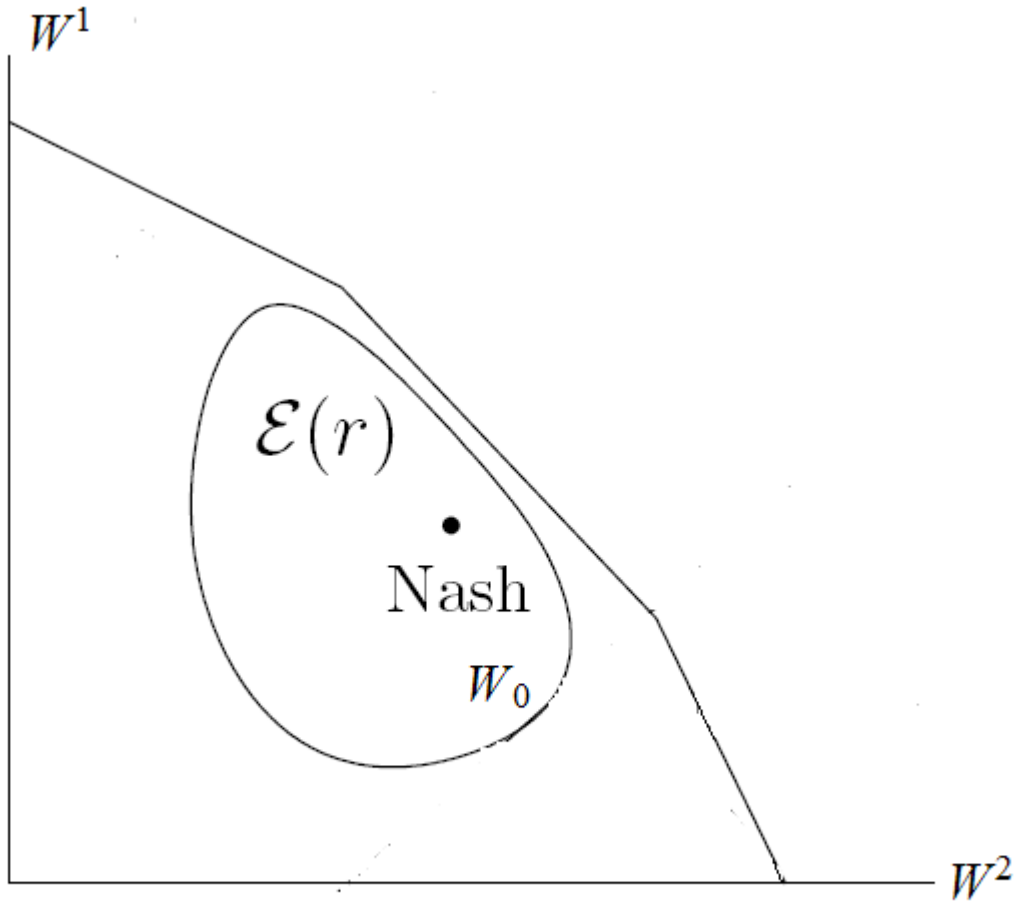

Continuous Time

Continuous Time Game:
Sannikov's (2005) Characterization of Set of PPE Payoffs

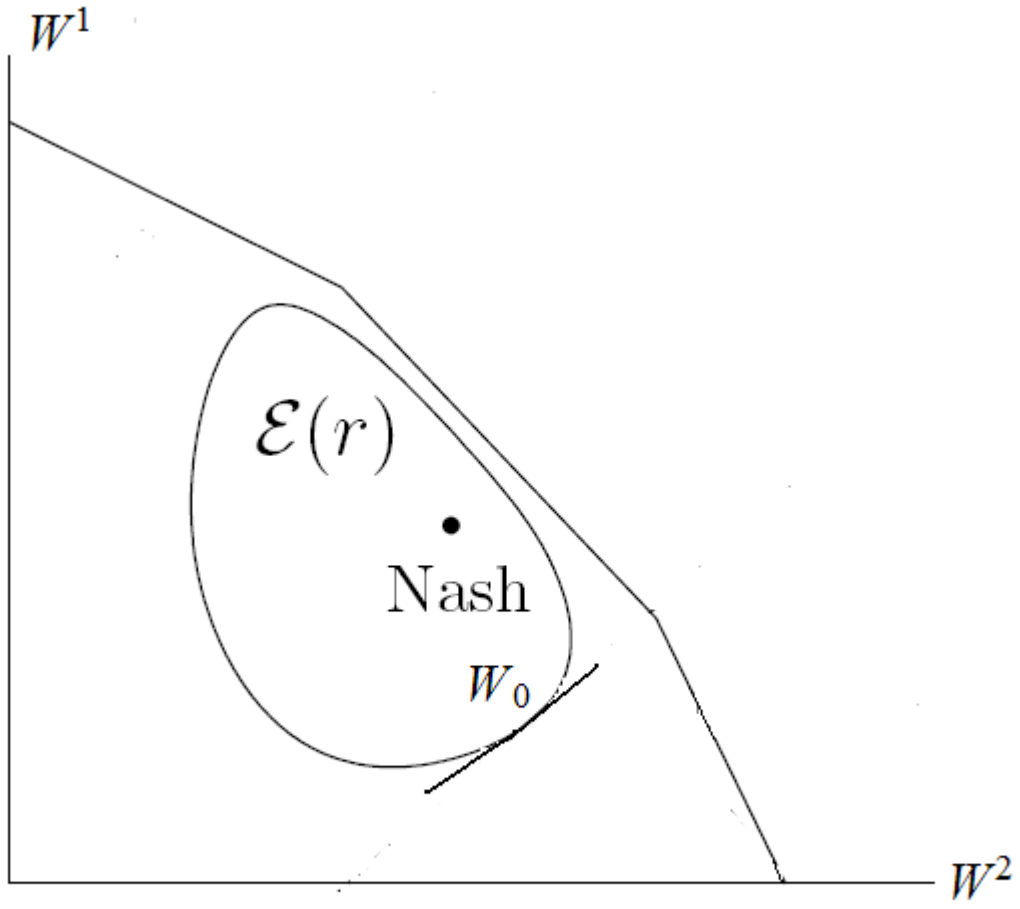

Continuous Time

Continuous Time Game:
Sannikov's (2005) Characterization of Set of PPE Payoffs

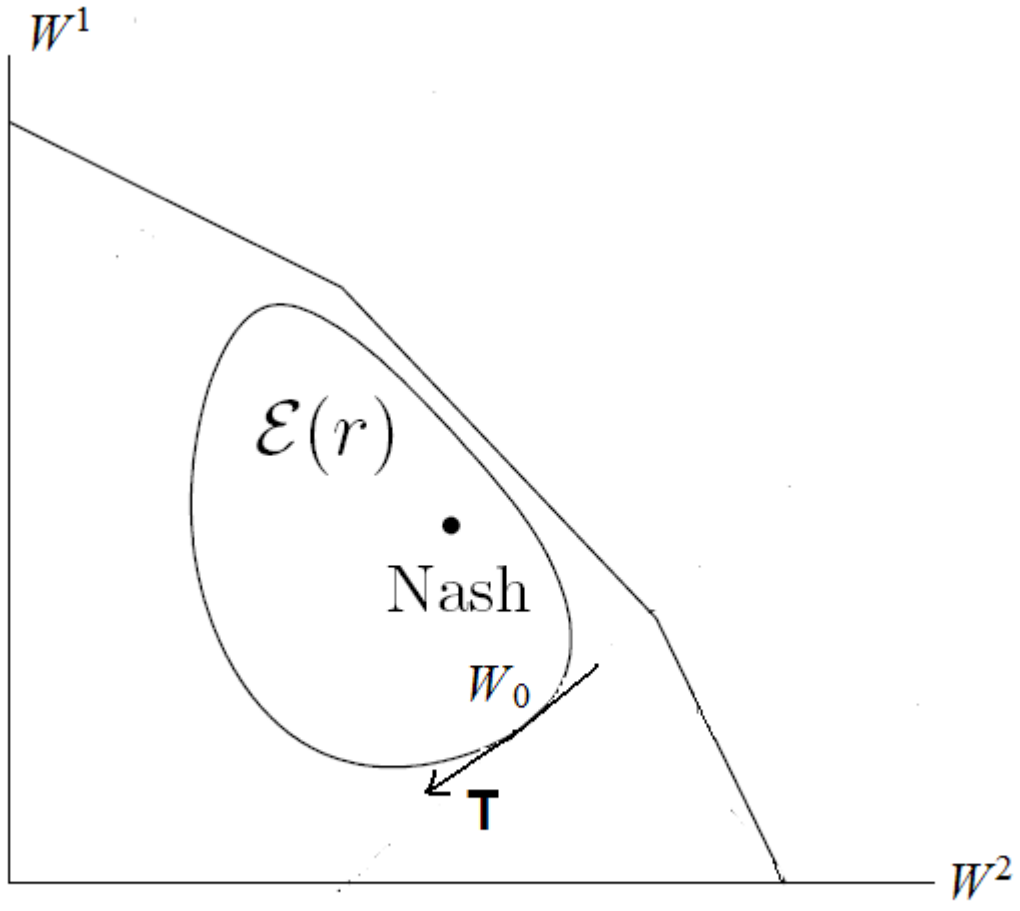

Continuous Time

Main Result

Main Result

$\mathcal{E}^{\Delta}(r) \to \mathcal{E}(r)$ for all discount rates r where $\mathcal{E}(\bullet)$ is continuous in r.

Main Result

$\mathcal{E}^{\Delta}(r) \to \mathcal{E}(r)$ for all discount rates r where $\mathcal{E}(\bullet)$ is continuous in r.

$\mathcal{E}^{\Delta}(r) \to \mathcal{E}(r)$ may fail at values of r where $\mathcal{E}(\bullet)$ is discontinuous in r.

Main Result

$\mathcal{E}^{\Delta}(r) \to \mathcal{E}(r)$ for all discount rates r where $\mathcal{E}(\cdot)$ is continuous in r.

$\mathcal{E}^{\Delta}(r) \to \mathcal{E}(r)$ may fail at values of r where $\mathcal{E}(\cdot)$ is discontinuous in r.

Theorem: Fix $r > 0$ and $\varepsilon > 0$. Then there exists $\underline{\Delta} > 0$ such that for all $\Delta \leq \underline{\Delta}$

Main Result

$\mathcal{E}^{\Delta}(r) \to \mathcal{E}(r)$ for all discount rates r where $\mathcal{E}(\cdot)$ is continuous in r.

$\mathcal{E}^{\Delta}(r) \to \mathcal{E}(r)$ may fail at values of r where $\mathcal{E}(\cdot)$ is discontinuous in r.

Theorem: Fix $r > 0$ and $\varepsilon > 0$. Then there exists $\underline{\Delta} > 0$ such that for all $\Delta \leq \underline{\Delta}$

$$\mathcal{E}(r+\varepsilon) \subseteq \mathcal{E}^{\Delta}(r) \subseteq \mathcal{E}(r-\varepsilon)$$

Discussion

Discussion

- It may seem surprising that as the players act more and more frequently, one gets a nontrivial limit: neither the Folk Theorem nor a collapse of equilibrium payoffs to a Nash equilibrium.

Discussion

- It may seem surprising that as the players act more and more frequently, one gets a nontrivial limit: neither the Folk Theorem nor a collapse of equilibrium payoffs to a Nash equilibrium.

- The key property of the limit is that as players act more and more frequently, the signals they see in each period become less and less informative, but the amount of information they learn per unit of time stays roughly constant.