Coalgebraic Analysis of Equilibria in Infinite Games

Samson Abramsky
Joint work with Viktor Winschel

Overview

Overview

We shall use coalgebraic methods to model and reason about **infinite extensive games**.

Overview

We shall use coalgebraic methods to model and reason about **infinite extensive games**.

Some notable points:

Overview

We shall use coalgebraic methods to model and reason about **infinite extensive games**.

Some notable points:

- We use coalgebraic methods to define **subgame perfect equilibria**, and to prove that certain strategy profiles are such equilibria.

Overview

We shall use coalgebraic methods to model and reason about **infinite extensive games**.

Some notable points:

- We use coalgebraic methods to define **subgame perfect equilibria**, and to prove that certain strategy profiles are such equilibria.
- This leads us to formulate an (apparently) novel notion of **predicate coinduction**.

Overview

We shall use coalgebraic methods to model and reason about **infinite extensive games**.

Some notable points:

- We use coalgebraic methods to define **subgame perfect equilibria**, and to prove that certain strategy profiles are such equilibria.
- This leads us to formulate an (apparently) novel notion of **predicate coinduction**.
- We can also prove a version of a standard game theory result — the 'one deviation principle' — without needing **future discounting** assumptions. This is promising for further applications to economic modelling.

Overview

We shall use coalgebraic methods to model and reason about **infinite extensive games**.

Some notable points:

- We use coalgebraic methods to define **subgame perfect equilibria**, and to prove that certain strategy profiles are such equilibria.
- This leads us to formulate an (apparently) novel notion of **predicate coinduction**.
- We can also prove a version of a standard game theory result — the 'one deviation principle' — without needing **future discounting** assumptions. This is promising for further applications to economic modelling.

These results are quite preliminary, but promising.

Overview

We shall use coalgebraic methods to model and reason about **infinite extensive games**.

Some notable points:

- We use coalgebraic methods to define **subgame perfect equilibria**, and to prove that certain strategy profiles are such equilibria.
- This leads us to formulate an (apparently) novel notion of **predicate coinduction**.
- We can also prove a version of a standard game theory result — the 'one deviation principle' — without needing **future discounting** assumptions. This is promising for further applications to economic modelling.

These results are quite preliminary, but promising.

Our work is inspired by recent work by Pierre Lescanne:
Rationality and escalation in infinite extensive games: arXiv:1112.1185, 2011.

Overview

We shall use coalgebraic methods to model and reason about **infinite extensive games**.

Some notable points:

- We use coalgebraic methods to define **subgame perfect equilibria**, and to prove that certain strategy profiles are such equilibria.
- This leads us to formulate an (apparently) novel notion of **predicate coinduction**.
- We can also prove a version of a standard game theory result — the 'one deviation principle' — without needing **future discounting** assumptions. This is promising for further applications to economic modelling.

These results are quite preliminary, but promising.

Our work is inspired by recent work by Pierre Lescanne:
Rationality and escalation in infinite extensive games: arXiv:1112.1185, 2011.

Our account is explicitly coalgebraic, leading to a mathematically richer and more general approach.

Motivating example: the dollar auction

Motivating example: the dollar auction

M. Shubik. The dollar auction game: A paradox in noncooperative behavior and escalation. Journal of Coflict Resolution, pages 109–111, 1971.

Motivating example: the dollar auction

M. Shubik. The dollar auction game: A paradox in noncooperative behavior and escalation. Journal of Coflict Resolution, pages 109–111, 1971.

The real value of the asset being bid for is r; in the original example, r is one dollar, or 100 cents. The players take it in turns to make bids. The asset goes to the highest bidder, who is left with a profit of $r - b$, where b is the value of his highest bid. The loser must also pay the value of his highest bid, while getting nothing in return. A player either gives up and finishes the game, conceding the auction to the other player and accepting their loss, or continues, hoping that the other player will give up. Both players have an incentive to continue playing well beyond the point where both will make a loss, in order to try to minimize their losses.

Games

Games

We shall consider possibly infinite games of perfect information in extensive form.
We fix the following sets:

Games

We shall consider possibly infinite games of perfect information in extensive form. We fix the following sets:

- A set $\mathcal{A}$ of **agents** or **players**. In our examples, we shall confine ourselves to two-player games, with $\mathcal{A} = \{A, B\}$.

Games

We shall consider possibly infinite games of perfect information in extensive form. We fix the following sets:

- A set $\mathcal{A}$ of **agents** or **players**. In our examples, we shall confine ourselves to two-player games, with $\mathcal{A} = \{A, B\}$.
- A set $\mathcal{C}$ of **choices**. We shall restrict our discussion to games where a player has a choice of two options, *left* or *right*, at every stage in the game where it is their turn to play. Thus $\mathcal{C} = \{\mathsf{l}, \mathsf{r}\}$.

Games

We shall consider possibly infinite games of perfect information in extensive form. We fix the following sets:

- A set $\mathcal{A}$ of **agents** or **players**. In our examples, we shall confine ourselves to two-player games, with $\mathcal{A} = \{A, B\}$.
- A set $\mathcal{C}$ of **choices**. We shall restrict our discussion to games where a player has a choice of two options, *left* or *right*, at every stage in the game where it is their turn to play. Thus $\mathcal{C} = \{\mathsf{l}, \mathsf{r}\}$.
- A set $\mathcal{U}$ of **utility functions**, which assign a utility to each player. We shall take $\mathcal{U} = \mathbb{R}^{\mathcal{A}}$.

Games

We shall consider possibly infinite games of perfect information in extensive form. We fix the following sets:

- A set $\mathcal{A}$ of **agents** or **players**. In our examples, we shall confine ourselves to two-player games, with $\mathcal{A} = \{A, B\}$.
- A set $\mathcal{C}$ of **choices**. We shall restrict our discussion to games where a player has a choice of two options, *left* or *right*, at every stage in the game where it is their turn to play. Thus $\mathcal{C} = \{\mathsf{l}, \mathsf{r}\}$.
- A set $\mathcal{U}$ of **utility functions**, which assign a utility to each player. We shall take $\mathcal{U} = \mathbb{R}^{\mathcal{A}}$.

The set of game trees is defined to be (the carrier of) the final coalgebra $(\mathcal{G}, \gamma)$ of the **Set**-functor

$$F_{\mathsf{G}} : X \mapsto \mathcal{U} + \mathcal{A} \times X \times X.$$

Games

We shall consider possibly infinite games of perfect information in extensive form. We fix the following sets:

- A set $\mathcal{A}$ of **agents** or **players**. In our examples, we shall confine ourselves to two-player games, with $\mathcal{A} = \{A, B\}$.
- A set $\mathcal{C}$ of **choices**. We shall restrict our discussion to games where a player has a choice of two options, *left* or *right*, at every stage in the game where it is their turn to play. Thus $\mathcal{C} = \{\mathsf{l}, \mathsf{r}\}$.
- A set $\mathcal{U}$ of **utility functions**, which assign a utility to each player. We shall take $\mathcal{U} = \mathbb{R}^{\mathcal{A}}$.

The set of game trees is defined to be (the carrier of) the final coalgebra $(\mathcal{G}, \gamma)$ of the **Set**-functor

$$F_{\mathsf{G}} : X \;\mapsto\; \mathcal{U} + \mathcal{A} \times X \times X.$$

Thus a game tree is a possibly infinite binary tree. The binary nodes have the form

$$\langle \alpha, g_{\mathsf{l}}, g_{\mathsf{r}} \rangle,$$

where α is the agent label, and g_{l}, g_{r} are the sub-games corresponding to the left and right choices respectively.

Example: the 0/1 game

Example: the 0/1 game

We firstly define utility functions

$$v := [A \mapsto 0,\ B \mapsto 1], \qquad w := [A \mapsto 1,\ B \mapsto 0].$$

Example: the 0/1 game

We firstly define utility functions

$$v := [A \mapsto 0,\ B \mapsto 1], \qquad w := [A \mapsto 1,\ B \mapsto 0].$$

The game is defined by the following pair of mutually corecursive equations:

$$G = \langle A, v, H \rangle, \qquad H = \langle B, w, G \rangle.$$

Example: the 0/1 game

We firstly define utility functions

$$v := [A \mapsto 0,\ B \mapsto 1], \qquad w := [A \mapsto 1,\ B \mapsto 0].$$

The game is defined by the following pair of mutually corecursive equations:

$$G = \langle A, v, H \rangle, \qquad H = \langle B, w, G \rangle.$$

We can depict the game tree as follows:

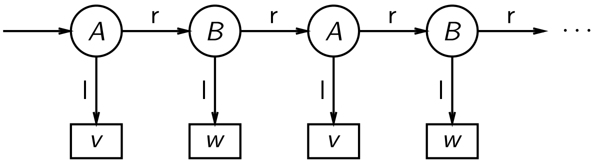

Example: the 0/1 game

We firstly define utility functions

$$v := [A \mapsto 0,\ B \mapsto 1], \qquad w := [A \mapsto 1,\ B \mapsto 0].$$

The game is defined by the following pair of mutually corecursive equations:

$$G = \langle A, v, H \rangle, \qquad H = \langle B, w, G \rangle.$$

We can depict the game tree as follows:

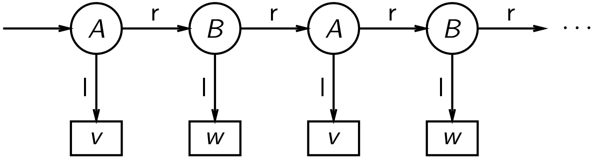

More formally, the equations define an F_{G}-colagebra $\alpha : \{G, H\} \to F_{\mathsf{G}}\{G, H\}$ on the set $\{G, H\}$. The 0/1 game is $[\![G]\!]$, where

$$[\![\cdot]\!] : (\{G, H\}, \alpha) \longrightarrow (\mathcal{G}, \gamma)$$

is the unique coalgebra morphism from $(\{G, H\}, \alpha)$ to the final coalgebra.

The dollar auction

The dollar auction

The $0/1$ game is 'finite-state'; it can be represented by a finite system of equations. The dollar auction game, a well-known example in game theory, is an infinite-state refinement, where the utility functions change as we move down the tree.

The dollar auction

The 0/1 game is ‘finite-state’; it can be represented by a finite system of equations. The dollar auction game, a well-known example in game theory, is an infinite-state refinement, where the utility functions change as we move down the tree.

We fix a real number r, and define utility functions v_n, w_n for each $n \in \mathbb{N}$:

$$v_n := [A \mapsto -n, B \mapsto r - n], \qquad w_n := [A \mapsto r - (n+1), B \mapsto -n].$$

The dollar auction

The 0/1 game is ‘finite-state’; it can be represented by a finite system of equations. The dollar auction game, a well-known example in game theory, is an infinite-state refinement, where the utility functions change as we move down the tree.

We fix a real number r, and define utility functions v_n, w_n for each $n \in \mathbb{N}$:

$$v_n := [A \mapsto -n, B \mapsto r - n], \qquad w_n := [A \mapsto r - (n+1), B \mapsto -n].$$

We can depict the game tree as follows:

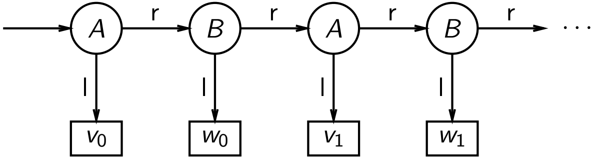

Formal definition of dollar auction game

Formal definition of dollar auction game

We define a game by simultaneous corecursion on the infinite set of variables $\{G_n, H_n \mid n \in \mathbb{N}\}$:

$$\begin{array}{rcl} G_n & = & \langle A, v_n, H_n \rangle \\ H_n & = & \langle B, w_n, G_{n+1} \rangle . \end{array}$$

Formal definition of dollar auction game

We define a game by simultaneous corecursion on the infinite set of variables $\{G_n, H_n \mid n \in \mathbb{N}\}$:

$$\begin{array}{rcl} G_n & = & \langle A, v_n, H_n \rangle \\ H_n & = & \langle B, w_n, G_{n+1} \rangle . \end{array}$$

The dollar auction game is $[\![G_0]\!]$.

Strategy profiles

Strategy profiles

Intuitively, a strategy for a player P of a game G specifies a choice (left or right) for every node of G at which it is P's turn to move. A strategy profile specifies a strategy for every player. Following Lescanne, we shall define the set of strategy profiles directly, as the final coalgebra $(\mathcal{S}, \sigma)$ of the functor

$$F_{\mathsf{SP}} : X \;\mapsto\; \mathcal{U} + \mathcal{A} \times \mathcal{C} \times X \times X.$$

Strategy profiles

Intuitively, a strategy for a player P of a game G specifies a choice (left or right) for every node of G at which it is P's turn to move. A strategy profile specifies a strategy for every player. Following Lescanne, we shall define the set of strategy profiles directly, as the final coalgebra $(\mathcal{S}, \sigma)$ of the functor

$$F_{\mathsf{SP}} : X \mapsto \mathcal{U} + \mathcal{A} \times \mathcal{C} \times X \times X.$$

There is an evident natural transformation $t : F_{\mathsf{SP}} \stackrel{\cdot}{\longrightarrow} F_{\mathsf{G}}$ defined by projection, which induces a functor from the category of F_{SP}-coalgebras to the category of F_{G}-coalgebras. It sends a strategy profile $s \in \mathcal{S}$ to the underlying game tree $\mathsf{game}(s) \in \mathcal{G}$. We say that s is a strategy profile for the game G if $G = \mathsf{game}(s)$.

Example: the 0/1 game

Example: the 0/1 game

We define the strategy profile AsBc, where A always stops (plays left) and B always continues (plays right), by the following simultaneous corecursion:

$$\begin{array}{rcl} \mathsf{AsBc} & = & \langle A, \mathsf{l}, v, \mathsf{BcAs} \rangle \\ \mathsf{BcAs} & = & \langle B, \mathsf{r}, w, \mathsf{AsBc} \rangle \end{array}$$

Example: the 0/1 game

We define the strategy profile AsBc, where A always stops (plays left) and B always continues (plays right), by the following simultaneous corecursion:

$$\begin{aligned} \mathsf{AsBc} &= \langle A, \mathsf{l}, v, \mathsf{BcAs} \rangle \\ \mathsf{BcAs} &= \langle B, \mathsf{r}, w, \mathsf{AsBc} \rangle \end{aligned}$$

We can depict this strategy profile as follows:

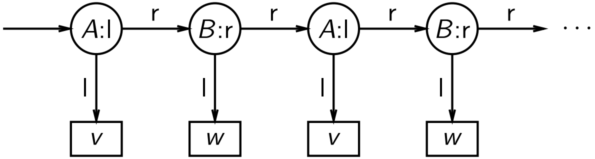

Example: the 0/1 game

We define the strategy profile AsBc, where A always stops (plays left) and B always continues (plays right), by the following simultaneous corecursion:

$$\begin{aligned} \mathsf{AsBc} &= \langle A, \mathsf{l}, v, \mathsf{BcAs} \rangle \\ \mathsf{BcAs} &= \langle B, \mathsf{r}, w, \mathsf{AsBc} \rangle \end{aligned}$$

We can depict this strategy profile as follows:

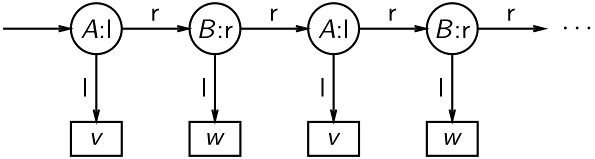

Formally, the strategy profile is $[\![\mathsf{AsBc}]\!]$.

Example: the dollar auction

Example: the dollar auction

In a similar fashion, we can define strategy profiles for the dollar auction.

Example: the dollar auction

In a similar fashion, we can define strategy profiles for the dollar auction.

We define a strategy profile by simultaneous corecursion on an infinite family of variables $\{\mathsf{AsBc}_n, \mathsf{BcAs}_n \mid n \in \mathbb{N}\}$:

$$\begin{aligned} \mathsf{AsBc}_n &= \langle A, \mathsf{l}, v_n, \mathsf{BcAs}_n \rangle \\ \mathsf{BcAs}_n &= \langle B, \mathsf{r}, w_n, \mathsf{AsBc}_{n+1} \rangle. \end{aligned}$$

Example: the dollar auction

In a similar fashion, we can define strategy profiles for the dollar auction.

We define a strategy profile by simultaneous corecursion on an infinite family of variables $\{\mathsf{AsBc}_n, \mathsf{BcAs}_n \mid n \in \mathbb{N}\}$:

$$\begin{array}{rcl} \mathsf{AsBc}_n & = & \langle A, \mathsf{l}, v_n, \mathsf{BcAs}_n \rangle \\ \mathsf{BcAs}_n & = & \langle B, \mathsf{r}, w_n, \mathsf{AsBc}_{n+1} \rangle. \end{array}$$

Formally, $[\![\mathsf{AsBc}_0]\!]$ is a strategy profile for the dollar auction game in which A always stops and B always continues.

Example: the dollar auction

In a similar fashion, we can define strategy profiles for the dollar auction.

We define a strategy profile by simultaneous corecursion on an infinite family of variables $\{\mathsf{AsBc}_n, \mathsf{BcAs}_n \mid n \in \mathbb{N}\}$:

$$\begin{aligned} \mathsf{AsBc}_n &= \langle A, \mathsf{l}, v_n, \mathsf{BcAs}_n \rangle \\ \mathsf{BcAs}_n &= \langle B, \mathsf{r}, w_n, \mathsf{AsBc}_{n+1} \rangle. \end{aligned}$$

Formally, $[\![\mathsf{AsBc}_0]\!]$ is a strategy profile for the dollar auction game in which A always stops and B always continues.

The strategy profile $[\![\mathsf{AcBs}_0]\!]$ where A always continues and B always stops is defined symmetrically.

Properties of strategy profiles I: convergence

Properties of strategy profiles I: convergence

We shall introduce two predicates on strategy profiles, of **weak** and **strong convergence**. A strategy profile is weakly convergent if following the choices it specifies from the root eventually leads to a leaf; it is strongly convergent if this holds in every sub-profile.

Properties of strategy profiles I: convergence

We shall introduce two predicates on strategy profiles, of **weak** and **strong convergence**. A strategy profile is weakly convergent if following the choices it specifies from the root eventually leads to a leaf; it is strongly convergent if this holds in every sub-profile.

We define weak convergence by:

$$s\downarrow \iff (s = U) \vee (s = \langle P, \mathsf{l}, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle \wedge s_{\mathsf{l}}\downarrow) \vee (s = \langle P, \mathsf{r}, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle \wedge s_{\mathsf{r}}\downarrow).$$

Properties of strategy profiles I: convergence

We shall introduce two predicates on strategy profiles, of **weak** and **strong convergence**. A strategy profile is weakly convergent if following the choices it specifies from the root eventually leads to a leaf; it is strongly convergent if this holds in every sub-profile.

We define weak convergence by:

$$s{\downarrow} \iff (s = U) \lor (s = \langle P, \mathsf{l}, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle \land s_{\mathsf{l}}{\downarrow}) \lor (s = \langle P, \mathsf{r}, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle \land s_{\mathsf{r}}{\downarrow}).$$

More formally, weak convergence is an element of the powerset $\mathcal{P}(\mathcal{S})$. It is defined as the least fixpoint of the monotone function

$$f_{\downarrow} : \mathcal{P}(\mathcal{S}) \longrightarrow \mathcal{P}(\mathcal{S}) \; :: \; S \mapsto \mathcal{U} \cup \{\langle P, \mathsf{l}, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle \mid s_{\mathsf{l}} \in S\} \cup \{\langle P, \mathsf{r}, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle \mid s_{\mathsf{r}} \in S\}.$$

Properties of strategy profiles I: convergence

We shall introduce two predicates on strategy profiles, of **weak** and **strong convergence**. A strategy profile is weakly convergent if following the choices it specifies from the root eventually leads to a leaf; it is strongly convergent if this holds in every sub-profile.

We define weak convergence by:

$$s{\downarrow} \iff (s = U) \vee (s = \langle P, \mathsf{l}, s_\mathsf{l}, s_\mathsf{r} \rangle \wedge s_\mathsf{l}{\downarrow}) \vee (s = \langle P, \mathsf{r}, s_\mathsf{l}, s_\mathsf{r} \rangle \wedge s_\mathsf{r}{\downarrow}).$$

More formally, weak convergence is an element of the powerset $\mathcal{P}(\mathcal{S})$. It is defined as the least fixpoint of the monotone function

$$f_{\downarrow} : \mathcal{P}(\mathcal{S}) \longrightarrow \mathcal{P}(\mathcal{S}) \;::\; S \mapsto \mathcal{U} \cup \{\langle P, \mathsf{l}, s_\mathsf{l}, s_\mathsf{r} \rangle \mid s_\mathsf{l} \in S\} \cup \{\langle P, \mathsf{r}, s_\mathsf{l}, s_\mathsf{r} \rangle \mid s_\mathsf{r} \in S\}.$$

Strong convergence is defined as follows:

$$s{\Downarrow} \iff (s = U) \vee (s = \langle P, c, s_\mathsf{l}, s_\mathsf{r} \rangle \wedge s{\downarrow} \wedge s_\mathsf{l}{\Downarrow} \wedge s_\mathsf{r}{\Downarrow}).$$

More formally, it is the greatest fixpoint of the monotone function $f_{\Downarrow}$ defined analogously to $f_{\downarrow}$.

Properties of strategy profiles I: convergence

We shall introduce two predicates on strategy profiles, of **weak** and **strong convergence**. A strategy profile is weakly convergent if following the choices it specifies from the root eventually leads to a leaf; it is strongly convergent if this holds in every sub-profile.

We define weak convergence by:

$$s\downarrow \iff (s = U) \vee (s = \langle P, \mathsf{l}, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle \wedge s_{\mathsf{l}}\downarrow) \vee (s = \langle P, \mathsf{r}, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle \wedge s_{\mathsf{r}}\downarrow).$$

More formally, weak convergence is an element of the powerset $\mathcal{P}(\mathcal{S})$. It is defined as the least fixpoint of the monotone function

$$f_{\downarrow} : \mathcal{P}(\mathcal{S}) \longrightarrow \mathcal{P}(\mathcal{S}) \ :: \ S \ \mapsto \ \mathcal{U} \cup \{\langle P, \mathsf{l}, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle \mid s_{\mathsf{l}} \in S\} \cup \{\langle P, \mathsf{r}, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle \mid s_{\mathsf{r}} \in S\}.$$

Strong convergence is defined as follows:

$$s\Downarrow \iff (s = U) \vee (s = \langle P, c, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle \wedge s\downarrow \wedge s_{\mathsf{l}}\Downarrow \wedge s_{\mathsf{r}}\Downarrow).$$

More formally, it is the greatest fixpoint of the monotone function $f_{\Downarrow}$ defined analogously to $f_{\downarrow}$.

These fixpoints exist by the Tarski (or Knaster-Tarski) fixed point theorem.

Discussion

Discussion

Note that weak convergence is an **inductive** notion; the profile must specify a finite path from the root to a leaf. Thus it is defined as a least fixpoint.

Discussion

Note that weak convergence is an **inductive** notion; the profile must specify a finite path from the root to a leaf. Thus it is defined as a least fixpoint.

Strong convergence, by contrast, expresses a constraint on all subtrees of an infinite tree, and hence is defined **coinductively**, as a greatest fixpoint.

Discussion

Note that weak convergence is an **inductive** notion; the profile must specify a finite path from the root to a leaf. Thus it is defined as a least fixpoint.

Strong convergence, by contrast, expresses a constraint on all subtrees of an infinite tree, and hence is defined **coinductively**, as a greatest fixpoint.

All the strategy profiles described in the previous section are strongly convergent.

Modal coalgebra view

Modal coalgebra view

The relationship between the two notions can be nicely characterized in terms of coalgebraic modal logic.

Modal coalgebra view

The relationship between the two notions can be nicely characterized in terms of coalgebraic modal logic.

We can define an 'always' modality $\Box$ as an operator on $\mathcal{P}(\mathcal{S})$:

$$s \models \Box\phi \;\equiv\; s \models \phi \;\wedge\; s = \langle P, c, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle \;\Rightarrow\; s_{\mathsf{l}} \models \Box\phi \;\wedge\; s_{\mathsf{r}} \models \Box\phi.$$

This operator is defined coinductively as a greatest fixpoint.

Modal coalgebra view

The relationship between the two notions can be nicely characterized in terms of coalgebraic modal logic.

We can define an 'always' modality $\Box$ as an operator on $\mathcal{P}(\mathcal{S})$:

$$s \models \Box\phi \;\equiv\; s \models \phi \;\wedge\; s = \langle P, c, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle \;\Rightarrow\; s_{\mathsf{l}} \models \Box\phi \;\wedge\; s_{\mathsf{r}} \models \Box\phi.$$

This operator is defined coinductively as a greatest fixpoint.

This modality generalizes straighforwardly to any polynomial functor, and can in fact be defined in a much more general way in the context of coalgebraic modal logic.

Modal coalgebra view

The relationship between the two notions can be nicely characterized in terms of coalgebraic modal logic.

We can define an 'always' modality $\Box$ as an operator on $\mathcal{P}(\mathcal{S})$:

$$s \models \Box\phi \;\equiv\; s \models \phi \;\wedge\; s = \langle P, c, s_\mathsf{l}, s_\mathsf{r} \rangle \;\Rightarrow\; s_\mathsf{l} \models \Box\phi \;\wedge\; s_\mathsf{r} \models \Box\phi.$$

This operator is defined coinductively as a greatest fixpoint.

This modality generalizes straighforwardly to any polynomial functor, and can in fact be defined in a much more general way in the context of coalgebraic modal logic.

If we write WC for the weak convergence predicate, and SC for the strong convergence predicate, we have the following:

Proposition

$\mathsf{SC} = \Box\mathsf{WC}$.

Utility functions induced by strategy profiles

Utility functions induced by strategy profiles

Given a strategy profile s, we define the utility function $\hat{s}$ which it induces. In general, this function may be partial; however, if s is weakly convergent, $\hat{s}$ is always a well-defined total function in $\mathcal{U}$.

Utility functions induced by strategy profiles

Given a strategy profile s, we define the utility function $\hat{s}$ which it induces. In general, this function may be partial; however, if s is weakly convergent, $\hat{s}$ is always a well-defined total function in $\mathcal{U}$.

This function is defined as follows:

Utility functions induced by strategy profiles

Given a strategy profile s, we define the utility function $\hat{s}$ which it induces. In general, this function may be partial; however, if s is weakly convergent, $\hat{s}$ is always a well-defined total function in $\mathcal{U}$.

This function is defined as follows:

- If $s = U$, then $\hat{s} = U$.

Utility functions induced by strategy profiles

Given a strategy profile s, we define the utility function $\hat{s}$ which it induces. In general, this function may be partial; however, if s is weakly convergent, $\hat{s}$ is always a well-defined total function in $\mathcal{U}$.

This function is defined as follows:

- If $s = U$, then $\hat{s} = U$.
- If $s = \langle P, \mathsf{l}, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle$, then $\hat{s} = \hat{s}_{\mathsf{l}}$.

Utility functions induced by strategy profiles

Given a strategy profile s, we define the utility function $\hat{s}$ which it induces. In general, this function may be partial; however, if s is weakly convergent, $\hat{s}$ is always a well-defined total function in $\mathcal{U}$.

This function is defined as follows:

- If $s = U$, then $\hat{s} = U$.
- If $s = \langle P, \mathsf{l}, s_\mathsf{l}, s_\mathsf{r} \rangle$, then $\hat{s} = \hat{s}_\mathsf{l}$.
- If $s = \langle P, \mathsf{r}, s_\mathsf{l}, s_\mathsf{r} \rangle$, then $\hat{s} = \hat{s}_\mathsf{r}$.

Subgame perfect equilibria

Subgame perfect equilibria

We shall now show how the game-theoretic notion of **subgame perfect equilibrium** can be defined coalgebraically.

Subgame perfect equilibria

We shall now show how the game-theoretic notion of **subgame perfect equilibrium** can be defined coalgebraically.

We firstly define a predicate PE on strategy profiles:

$$\mathrm{PE}(s) \iff s{\Downarrow} \wedge (s = \langle P, c, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle \Rightarrow \hat{s}_c(P) \geq \hat{s}_{c'}(P))$$

Subgame perfect equilibria

We shall now show how the game-theoretic notion of **subgame perfect equilibrium** can be defined coalgebraically.

We firstly define a predicate PE on strategy profiles:

$$\mathrm{PE}(s) \iff s{\Downarrow} \wedge (s = \langle P, c, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle \Rightarrow \hat{s}_c(P) \geq \hat{s}_{c'}(P))$$

This predicate says that the choice at the root for player P results in a better payoff for P than the other choice would have done.

Subgame perfect equilibria

We shall now show how the game-theoretic notion of **subgame perfect equilibrium** can be defined coalgebraically.

We firstly define a predicate PE on strategy profiles:

$$\mathrm{PE}(s) \iff s{\Downarrow} \wedge (s = \langle P, c, s_\mathsf{l}, s_\mathsf{r} \rangle \Rightarrow \hat{s}_c(P) \geq \hat{s}_{c'}(P))$$

This predicate says that the choice at the root for player P results in a better payoff for P than the other choice would have done.

We now define the subgame perfect equilibrium predicate SPE on $\mathcal{S}$, coinductively as a greatest fixpoint.

$$\mathrm{SPE}(s) \iff \mathrm{PE}(s) \wedge (s = \langle P, c, s_\mathsf{l}, s_\mathsf{r} \rangle \Rightarrow \mathrm{SPE}(s_\mathsf{l}) \wedge \mathrm{SPE}(s_\mathsf{r})).$$

Subgame perfect equilibria

We shall now show how the game-theoretic notion of **subgame perfect equilibrium** can be defined coalgebraically.

We firstly define a predicate PE on strategy profiles:

$$\mathsf{PE}(s) \iff s\Downarrow \wedge \ (s = \langle P, c, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle \ \Rightarrow \ \hat{s}_c(P) \geq \hat{s}_{c'}(P))$$

This predicate says that the choice at the root for player P results in a better payoff for P than the other choice would have done.

We now define the subgame perfect equilibrium predicate SPE on $\mathcal{S}$, coinductively as a greatest fixpoint.

$$\mathsf{SPE}(s) \iff \mathsf{PE}(s) \wedge \ (s = \langle P, c, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle \ \Rightarrow \ \mathsf{SPE}(s_{\mathsf{l}}) \wedge \ \mathsf{SPE}(s_{\mathsf{r}})).$$

This says that the PE predicate holds at every node of the tree. In fact, we have:

Proposition

$\mathsf{SPE} = \Box\mathsf{PE}.$

Predicate coinduction

Predicate coinduction

The next step is to show that all the strategy profiles discussed in the previous section are in fact subgame perfect equilibria. In order to do this, we will need to have an appropriate proof principle in place.

Predicate coinduction

The next step is to show that all the strategy profiles discussed in the previous section are in fact subgame perfect equilibria. In order to do this, we will need to have an appropriate proof principle in place.

The main emphasis in coinductive proofs has been on proving equations; the main tool for this is provided by the notion of bisimulation. However, as emphasized by Kozen, the scope of coinductive methods is broader than this. In our case, we are interested in **predicates** (properties) rather than equations. In particular, we wish to show that various elements of the final coalgebra $\mathcal{S}$ satisfy the SPE predicate.

Predicate coinduction

The next step is to show that all the strategy profiles discussed in the previous section are in fact subgame perfect equilibria. In order to do this, we will need to have an appropriate proof principle in place.

The main emphasis in coinductive proofs has been on proving equations; the main tool for this is provided by the notion of bisimulation. However, as emphasized by Kozen, the scope of coinductive methods is broader than this. In our case, we are interested in **predicates** (properties) rather than equations. In particular, we wish to show that various elements of the final coalgebra $\mathcal{S}$ satisfy the SPE predicate.

We shall formulate a proof principle which is adequate to carry out these proofs, and justify it in terms of Kozen's metric coinduction principle. It should be possible to give a much more general account; this is an interesting challenge, which is left to future work.

The setting

The setting

We shall formulate the principle in terms of an arbitrary polynomial functor T; we will apply it in the case $T = F_{\mathsf{SP}}$. The final coalgebra of T is denoted $(\mathcal{S}, \sigma)$.

The setting

We shall formulate the principle in terms of an arbitrary polynomial functor T; we will apply it in the case $T = F_{\mathsf{SP}}$. The final coalgebra of T is denoted $(\mathcal{S}, \sigma)$.

Some notation. Suppose we have a T-coalgebra (X, α), which we think of as a corecursive system of equations on the set of variables X. We define a map

$$\bar{\alpha} : \mathcal{S}^X \longrightarrow \mathcal{S}^X :: \eta \mapsto [x \mapsto \sigma^{-1} \circ T\eta \circ \alpha(x)].$$

The setting

We shall formulate the principle in terms of an arbitrary polynomial functor T; we will apply it in the case $T = F_{\mathsf{SP}}$. The final coalgebra of T is denoted $(\mathcal{S}, \sigma)$.

Some notation. Suppose we have a T-coalgebra (X, α), which we think of as a corecursive system of equations on the set of variables X. We define a map

$$\bar{\alpha} : \mathcal{S}^X \longrightarrow \mathcal{S}^X :: \eta \mapsto [x \mapsto \sigma^{-1} \circ T\eta \circ \alpha(x)].$$

Here $\sigma^{-1} : T\mathcal{S} \longrightarrow \mathcal{S}$ is the inverse of σ, which is an isomorphism by the Lambek lemma.

The setting

We shall formulate the principle in terms of an arbitrary polynomial functor T; we will apply it in the case $T = F_{\mathsf{SP}}$. The final coalgebra of T is denoted $(\mathcal{S}, \sigma)$.

Some notation. Suppose we have a T-coalgebra (X, α), which we think of as a corecursive system of equations on the set of variables X. We define a map

$$\bar{\alpha} : \mathcal{S}^X \longrightarrow \mathcal{S}^X :: \eta \mapsto [x \mapsto \sigma^{-1} \circ T\eta \circ \alpha(x)].$$

Here $\sigma^{-1} : T\mathcal{S} \longrightarrow \mathcal{S}$ is the inverse of σ, which is an isomorphism by the Lambek lemma.

The following proposition follows directly by unravelling the definitions and applying the final coalgebra property:

Proposition

The map $\bar{\alpha}$ has a unique fixpoint $\bar{\alpha}^* \in \mathcal{S}^X$; moreover, $\bar{\alpha}^* = [\![\cdot]\!]$, the unique coalgebra morphism from (X, α) to the final coalgebra.

The predicate coinduction principle

The predicate coinduction principle

Now let $\phi \subseteq \mathcal{S}$ be a predicate on $\mathcal{S}$. We lift this to $\phi^X \subseteq \mathcal{S}^X$:

$$\eta \in \phi^X \iff \forall x \in X.\, \eta(x) \in \phi.$$

We say that ϕ is an **invariant** if $\phi = \Box\psi$ for some predicate ψ.

The predicate coinduction principle

Now let $\phi \subseteq \mathcal{S}$ be a predicate on $\mathcal{S}$. We lift this to $\phi^X \subseteq \mathcal{S}^X$:

$$\eta \in \phi^X \iff \forall x \in X.\, \eta(x) \in \phi.$$

We say that ϕ is an **invariant** if $\phi = \Box\psi$ for some predicate ψ.

We formulate the predicate coinduction principle as a proof rule:

$$\frac{\phi \text{ invariant}, \ \phi \neq \varnothing, \quad \eta \in \phi^X \ \Rightarrow \ \bar{\alpha}(\eta) \in \phi^X}{\bar{\alpha}^* \in \phi^X}$$

The predicate coinduction principle

Now let $\phi \subseteq \mathcal{S}$ be a predicate on $\mathcal{S}$. We lift this to $\phi^X \subseteq \mathcal{S}^X$:

$$\eta \in \phi^X \iff \forall x \in X.\, \eta(x) \in \phi.$$

We say that ϕ is an **invariant** if $\phi = \Box\psi$ for some predicate ψ.

We formulate the predicate coinduction principle as a proof rule:

$$\frac{\phi \text{ invariant, } \phi \neq \varnothing, \quad \eta \in \phi^X \;\Rightarrow\; \bar{\alpha}(\eta) \in \phi^X}{\bar{\alpha}^* \in \phi^X}$$

Using Proposition 3, we can restate the conclusion as follows:

$$\forall x \in X.\, [\![x]\!] \in \phi.$$

Metric coinduction

Metric coinduction

In order to show the soundness of this principle, we shall reduce it to the following metric coinduction principle due to Kozen:

Metric coinduction

In order to show the soundness of this principle, we shall reduce it to the following metric coinduction principle due to Kozen:

Proposition

Let M be a complete metric space, $u : M \to M$ a contractive map, and $C \subseteq M$ a non-empty closed subset of M. We write u^* for the unique fixpoint of u, which exists by the Banach fixpoint theorem. Then the following proof rule is valid:

$$\frac{x \in C \;\Rightarrow\; u(x) \in C}{u^* \in C}$$

Metric coinduction

In order to show the soundness of this principle, we shall reduce it to the following metric coinduction principle due to Kozen:

Proposition

Let M be a complete metric space, $u : M \to M$ a contractive map, and $C \subseteq M$ a non-empty closed subset of M. We write u^* for the unique fixpoint of u, which exists by the Banach fixpoint theorem. Then the following proof rule is valid:

$$\frac{x \in C \;\Rightarrow\; u(x) \in C}{u^* \in C}$$

Proof We have $u^* = \lim_{n\to\infty} u^n(a)$ for any $a \in M$. Since C is non-empty, we can take $a \in C$. By the premise of the rule, $u^n(a) \in C$ for all n. Since C is closed, $u^* \in C$. □

Reduction of predicate coinduction to metric coinduction

Reduction of predicate coinduction to metric coinduction

The reduction is an immediate consequence of the following result.

Reduction of predicate coinduction to metric coinduction

The reduction is an immediate consequence of the following result.

Proposition (Induced ultrametric)

We can define a distance function $d : \mathcal{S}^2 \to [0,1]$ *such that the following holds:*

1. $(\mathcal{S}, d)$ *is a complete ultrametric space.*
2. *An invariant* $\phi \subseteq \mathcal{S}$ *is closed in this space.*
3. *For any set* X, $(\mathcal{S}^X, d^X)$ *is a complete ultrametric space, where*

$$d^X(\eta, \mu) \;:=\; \sup_{x \in X} d(\eta_x, \mu_x).$$

4. *For any coalgebra* (X, α), *the map* $\bar{\alpha}$ *is contractive.*

Reduction of predicate coinduction to metric coinduction

The reduction is an immediate consequence of the following result.

Proposition (Induced ultrametric)

We can define a distance function $d : \mathcal{S}^2 \to [0,1]$ *such that the following holds:*

1. $(\mathcal{S}, d)$ *is a complete ultrametric space.*
2. *An invariant* $\phi \subseteq \mathcal{S}$ *is closed in this space.*
3. *For any set* X, $(\mathcal{S}^X, d^X)$ *is a complete ultrametric space, where*

$$d^X(\eta, \mu) := \sup_{x \in X} d(\eta_x, \mu_x).$$

4. *For any coalgebra* (X, α), *the map* $\bar{\alpha}$ *is contractive.*

This is essentially a variation of well-known results by Barr et al.

Proof of Induced ultrametric proposition

Proof of Induced ultrametric proposition

The basic point is that T is cocontinuous, and hence the final T-coalgebra is the limit of the ω^{op}-diagram

$$\mathbf{1} \longleftarrow T\mathbf{1} \longleftarrow T^2\mathbf{1} \longleftarrow \cdots \longleftarrow T^k\mathbf{1} \longleftarrow \cdots \tag{1}$$

where $\mathbf{1}$ is the terminal object in **Set**. The connecting maps are $T^k! : T^{k+1}\mathbf{1} \longrightarrow T^k\mathbf{1}$, where $! : T\mathbf{1} \longrightarrow \mathbf{1}$ is the unique map to the terminal object.

Proof of Induced ultrametric proposition

The basic point is that T is cocontinuous, and hence the final T-coalgebra is the limit of the ω^{op}-diagram

$$\mathbf{1} \longleftarrow T\mathbf{1} \longleftarrow T^2\mathbf{1} \longleftarrow \cdots \longleftarrow T^k\mathbf{1} \longleftarrow \cdots \tag{1}$$

where $\mathbf{1}$ is the terminal object in **Set**. The connecting maps are $T^k! : T^{k+1}\mathbf{1} \longrightarrow T^k\mathbf{1}$, where $! : T\mathbf{1} \longrightarrow \mathbf{1}$ is the unique map to the terminal object.

We define maps $\pi_k : \mathcal{S} \to T^k\mathbf{1}$. Firstly π_0 is the unique map to the terminal. We define inductively

$$\pi_{k+1} \; := \; \mathcal{S} \xrightarrow{\sigma} T\mathcal{S} \xrightarrow{T\pi_k} TT^k\mathbf{1}.$$

Then $(\mathcal{S}, \{\pi_k\})$ is the limit cone of the diagram (1).

Proof of Induced ultrametric proposition

The basic point is that T is cocontinuous, and hence the final T-coalgebra is the limit of the ω^{op}-diagram

$$\mathbf{1} \longleftarrow T\mathbf{1} \longleftarrow T^2\mathbf{1} \longleftarrow \cdots \longleftarrow T^k\mathbf{1} \longleftarrow \cdots \qquad (1)$$

where $\mathbf{1}$ is the terminal object in **Set**. The connecting maps are $T^k! : T^{k+1}\mathbf{1} \longrightarrow T^k\mathbf{1}$, where $! : T\mathbf{1} \longrightarrow \mathbf{1}$ is the unique map to the terminal object.

We define maps $\pi_k : \mathcal{S} \to T^k\mathbf{1}$. Firstly π_0 is the unique map to the terminal. We define inductively

$$\pi_{k+1} \;:=\; \mathcal{S} \xrightarrow{\sigma} T\mathcal{S} \xrightarrow{T\pi_k} TT^k\mathbf{1}.$$

Then $(\mathcal{S}, \{\pi_k\})$ is the limit cone of the diagram (1).

This has the following consequences:

- For all $x, y \in \mathcal{S}$:
$$x = y \iff \forall k.\, \pi_k x = \pi_k y.$$
- Given a sequence $\{x_k\}$ with $x_k \in T^k\mathbf{1}$ and $x_k = T^k!(x_{k+1})$, there is a unique $x \in \mathcal{S}$ such that for all k, $x_k = \pi_k x$.

Proof ctd

Proof ctd

We can define the ultrametric on $\mathcal{S}$ by

$$d(x,y) = \begin{cases} 0, & x = y \\ 2^{-k}, & \text{least } k \text{ such that } \pi_k x \neq \pi_k y,\ x \neq y. \end{cases}$$

Proof ctd

We can define the ultrametric on $\mathcal{S}$ by

$$d(x,y) = \begin{cases} 0, & x = y \\ 2^{-k}, & \text{least } k \text{ such that } \pi_k x \neq \pi_k y,\ x \neq y. \end{cases}$$

The properties stated in the proposition follow fairly readily from this description. For example, the completeness of d^X follows from the fact that Cauchy convergence in this metric implies that for each k, for some N, for all $x \in X$, $\pi_k(\eta_j(x))$ is fixed for all $j \geq N$. This implies a uniform mode of convergence, and thus justifies taking limits pointwise in $\mathcal{S}^X$.

SPE proofs

SPE proofs

We shall now apply predicate coinduction to show that the strategy profiles we defined are subgame-perfect equilibria.

SPE proofs

We shall now apply predicate coinduction to show that the strategy profiles we defined are subgame-perfect equilibria.

We begin with the strategy profiles for the 0/1-game.

SPE proofs

We shall now apply predicate coinduction to show that the strategy profiles we defined are subgame-perfect equilibria.

We begin with the strategy profiles for the 0/1-game.

Proposition

The strategy profiles $[\![\mathrm{AcBs}]\!]$ *and* $[\![\mathrm{AsBc}]\!]$ *are subgame-perfect equilibria.*

SPE proofs

We shall now apply predicate coinduction to show that the strategy profiles we defined are subgame-perfect equilibria.

We begin with the strategy profiles for the 0/1-game.

Proposition

The strategy profiles $[\![\mathrm{AcBs}]\!]$ *and* $[\![\mathrm{AsBc}]\!]$ *are subgame-perfect equilibria.*

Proof Applying predicate coinduction, to show that $[\![\mathrm{AsBc}]\!]$ is SPE, we must show that $\langle A, \mathrm{l}, v, \mathrm{BcAs}\rangle$ and $\langle B, \mathrm{r}, w, \mathrm{AsBc}\rangle$ are SPE, under the assumption that AsBc and BcAs are SPE.

SPE proofs

We shall now apply predicate coinduction to show that the strategy profiles we defined are subgame-perfect equilibria.

We begin with the strategy profiles for the 0/1-game.

Proposition

The strategy profiles $[\![\mathrm{AcBs}]\!]$ *and* $[\![\mathrm{AsBc}]\!]$ *are subgame-perfect equilibria.*

Proof Applying predicate coinduction, to show that $[\![\mathrm{AsBc}]\!]$ is SPE, we must show that $\langle A, \mathrm{l}, v, \mathrm{BcAs}\rangle$ and $\langle B, \mathrm{r}, w, \mathrm{AsBc}\rangle$ are SPE, under the assumption that AsBc and BcAs are SPE.

This reduces to verifying the inequalities

$$v(A) = 0 \geq 0 = v(A), \quad v(B) = 1 \geq 0 = w(B).$$

SPE proofs

We shall now apply predicate coinduction to show that the strategy profiles we defined are subgame-perfect equilibria.

We begin with the strategy profiles for the 0/1-game.

Proposition

The strategy profiles $[\![\mathrm{AcBs}]\!]$ *and* $[\![\mathrm{AsBc}]\!]$ *are subgame-perfect equilibria.*

Proof Applying predicate coinduction, to show that $[\![\mathrm{AsBc}]\!]$ is SPE, we must show that $\langle A, \mathrm{l}, v, \mathrm{BcAs}\rangle$ and $\langle B, \mathrm{r}, w, \mathrm{AsBc}\rangle$ are SPE, under the assumption that AsBc and BcAs are SPE.

This reduces to verifying the inequalities

$$v(A) = 0 \geq 0 = v(A), \quad v(B) = 1 \geq 0 = w(B).$$

The verification that $[\![\mathrm{AcBs}]\!]$ is SPE is entirely similar. □

The dollar auction

The dollar auction

Recall the parameter r used to define the utility functions v_n, w_n.

The dollar auction

Recall the parameter r used to define the utility functions v_n, w_n.

Proposition

If $r \geq 1$, the strategy profiles $[\![\mathrm{AcBs}_n]\!]$ and $[\![\mathrm{AsBc}_n]\!]$ are subgame-perfect equilibria for all n.

The dollar auction

Recall the parameter r used to define the utility functions v_n, w_n.

Proposition

If $r \geq 1$, the strategy profiles $[\![\mathrm{AcBs}_n]\!]$ and $[\![\mathrm{AsBc}_n]\!]$ are subgame-perfect equilibria for all n.

Proof The coinduction hypothesis is that AsBc_n and BcAs_n are SPE for all n.

The dollar auction

Recall the parameter r used to define the utility functions v_n, w_n.

Proposition

If $r \geq 1$, the strategy profiles $[\![\mathrm{AcBs}_n]\!]$ and $[\![\mathrm{AsBc}_n]\!]$ are subgame-perfect equilibria for all n.

Proof The coinduction hypothesis is that AsBc_n and BcAs_n are SPE for all n.

We must show that $\langle A, \mathrm{l}, v_n, \mathrm{BcAs}_n \rangle$ and $\langle B, \mathrm{r}, w_n, \mathrm{AsBc}_{n+1} \rangle$ are SPE for all n.

The dollar auction

Recall the parameter r used to define the utility functions v_n, w_n.

Proposition

If $r \geq 1$, the strategy profiles $[\![\mathrm{AcBs}_n]\!]$ and $[\![\mathrm{AsBc}_n]\!]$ are subgame-perfect equilibria for all n.

Proof The coinduction hypothesis is that AsBc_n and BcAs_n are SPE for all n.

We must show that $\langle A, \mathrm{l}, v_n, \mathrm{BcAs}_n \rangle$ and $\langle B, \mathrm{r}, w_n, \mathrm{AsBc}_{n+1} \rangle$ are SPE for all n.

This reduces to verifying the inequalities

$$v_n(A) = -n \geq -(n+1) = v_{n+1}(A), \quad v_{n+1}(B) = r - (n+1) \geq -n = w_n(B).$$

The dollar auction

Recall the parameter r used to define the utility functions v_n, w_n.

Proposition

If $r \geq 1$, the strategy profiles $[\![\mathrm{AcBs}_n]\!]$ and $[\![\mathrm{AsBc}_n]\!]$ are subgame-perfect equilibria for all n.

Proof The coinduction hypothesis is that AsBc_n and BcAs_n are SPE for all n.

We must show that $\langle A, \mathrm{l}, v_n, \mathrm{BcAs}_n \rangle$ and $\langle B, \mathrm{r}, w_n, \mathrm{AsBc}_{n+1} \rangle$ are SPE for all n.

This reduces to verifying the inequalities

$$v_n(A) = -n \geq -(n+1) = v_{n+1}(A), \quad v_{n+1}(B) = r - (n+1) \geq -n = w_n(B).$$

Note that the latter inequalities are satisfied if and only if $r \geq 1$. □

The one-deviation principle

The one-deviation principle

We shall now verify an important property of subgame-perfect equilibria: a strategy profile is SPE if and only if it dominates any profile which differs from it in exactly one choice.

The one-deviation principle

We shall now verify an important property of subgame-perfect equilibria: a strategy profile is SPE if and only if it dominates any profile which differs from it in exactly one choice.

In the standard game-theory literature, this is proved for infinite-horizon games only under strong continuity assumptions on the payoffs, which amount to discounting beyond some finite horizon, thus allowing reduction to standard backwards induction reasoning. We shall need no such assumptions.

The one-deviation principle

We shall now verify an important property of subgame-perfect equilibria: a strategy profile is SPE if and only if it dominates any profile which differs from it in exactly one choice.

In the standard game-theory literature, this is proved for infinite-horizon games only under strong continuity assumptions on the payoffs, which amount to discounting beyond some finite horizon, thus allowing reduction to standard backwards induction reasoning. We shall need no such assumptions.

Given a strategy profile s of the form $\langle P, c, s_\mathsf{l}, s_\mathsf{r} \rangle$, we say that a profile t for the same game is a one-deviation from s if t has one of the following forms:

- $\langle P, c', s_\mathsf{l}, s_\mathsf{r} \rangle$, where $c' \neq c$.
- $\langle P, c, s'_\mathsf{l}, s_\mathsf{r} \rangle$, where s'_l is a one-deviation from s_l.
- $\langle P, c, s_\mathsf{l}, s'_\mathsf{r} \rangle$, where s'_r is a one-deviation from s_r.

This is an inductive definition.

The one-deviation principle

The one-deviation principle

Given a strategy profile s we define a relation $s \succcurlyeq t$, where t is a one-deviation of s, inductively as follows:

- If $s = \langle P, \mathsf{l}, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle$ and $t = \langle P, \mathsf{r}, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle$, then $s \succcurlyeq t$ iff $\hat{s}_{\mathsf{l}}(P) \geq \hat{s}_{\mathsf{r}}(P))$.
- If $s = \langle P, \mathsf{r}, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle$ and $t = \langle P, \mathsf{l}, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle$, then $s \succcurlyeq t$ iff $\hat{s}_{\mathsf{r}}(P) \geq \hat{s}_{\mathsf{l}}(P))$.
- If $s = \langle P, c, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle$ and $t = \langle P, c, s'_{\mathsf{l}}, s_{\mathsf{r}} \rangle$, then $s \succcurlyeq t$ iff $s_{\mathsf{l}} \succcurlyeq s'_{\mathsf{l}}$.
- If $s = \langle P, c, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle$ and $t = \langle P, c, s_{\mathsf{l}}, s'_{\mathsf{r}} \rangle$, then $s \succcurlyeq t$ iff $s_{\mathsf{r}} \succcurlyeq s'_{\mathsf{r}}$.

The one-deviation principle

Given a strategy profile s we define a relation $s \succcurlyeq t$, where t is a one-deviation of s, inductively as follows:

- If $s = \langle P, \mathsf{l}, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle$ and $t = \langle P, \mathsf{r}, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle$, then $s \succcurlyeq t$ iff $\hat{s}_{\mathsf{l}}(P) \geq \hat{s}_{\mathsf{r}}(P))$.
- If $s = \langle P, \mathsf{r}, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle$ and $t = \langle P, \mathsf{l}, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle$, then $s \succcurlyeq t$ iff $\hat{s}_{\mathsf{r}}(P) \geq \hat{s}_{\mathsf{l}}(P))$.
- If $s = \langle P, c, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle$ and $t = \langle P, c, s'_{\mathsf{l}}, s_{\mathsf{r}} \rangle$, then $s \succcurlyeq t$ iff $s_{\mathsf{l}} \succcurlyeq s'_{\mathsf{l}}$.
- If $s = \langle P, c, s_{\mathsf{l}}, s_{\mathsf{r}} \rangle$ and $t = \langle P, c, s_{\mathsf{l}}, s'_{\mathsf{r}} \rangle$, then $s \succcurlyeq t$ iff $s_{\mathsf{r}} \succcurlyeq s'_{\mathsf{r}}$.

Proposition (The one-deviation principle)

A strongly convergent strategy profile s is SPE *if and only if for every one-deviation t, $s \succcurlyeq t$.*

The one-deviation principle

Given a strategy profile s we define a relation $s \succcurlyeq t$, where t is a one-deviation of s, inductively as follows:

- If $s = \langle P, \mathsf{l}, s_\mathsf{l}, s_\mathsf{r} \rangle$ and $t = \langle P, \mathsf{r}, s_\mathsf{l}, s_\mathsf{r} \rangle$, then $s \succcurlyeq t$ iff $\hat{s}_\mathsf{l}(P) \geq \hat{s}_\mathsf{r}(P))$.
- If $s = \langle P, \mathsf{r}, s_\mathsf{l}, s_\mathsf{r} \rangle$ and $t = \langle P, \mathsf{l}, s_\mathsf{l}, s_\mathsf{r} \rangle$, then $s \succcurlyeq t$ iff $\hat{s}_\mathsf{r}(P) \geq \hat{s}_\mathsf{l}(P))$.
- If $s = \langle P, c, s_\mathsf{l}, s_\mathsf{r} \rangle$ and $t = \langle P, c, s'_\mathsf{l}, s_\mathsf{r} \rangle$, then $s \succcurlyeq t$ iff $s_\mathsf{l} \succcurlyeq s'_\mathsf{l}$.
- If $s = \langle P, c, s_\mathsf{l}, s_\mathsf{r} \rangle$ and $t = \langle P, c, s_\mathsf{l}, s'_\mathsf{r} \rangle$, then $s \succcurlyeq t$ iff $s_\mathsf{r} \succcurlyeq s'_\mathsf{r}$.

Proposition (The one-deviation principle)

A strongly convergent strategy profile s is SPE *if and only if for every one-deviation t, $s \succcurlyeq t$.*

Proof Firstly, note that if $\neg(s \succcurlyeq t)$ for some one-deviation t, this means that some subprofile of s does not satisfy PE, and since SPE $= \Box$PE, this implies that s does not satisfy SPE. For the converse, note that if for some one-deviation t, $s \succcurlyeq t$, this implies that the sub-profile of s whose root is at the node where t differs from s satisfies PE. If this holds for all one-deviations, then all sub-profiles of s satisfy PE, and hence s satisfies SPE. $\Box$

Further Results

Further Results

We can give a **complete characterization** of the subgame-perfect equilibria for the dollar auction game.

Further Results

We can give a **complete characterization** of the subgame-perfect equilibria for the dollar auction game.

Theorem

A strategy profile for the dollar auction game is SPE if and only if it is of the form

A always stops and B always continues

or symmetrically

B always stops and A always continues.

Further Results

We can give a **complete characterization** of the subgame-perfect equilibria for the dollar auction game.

Theorem

A strategy profile for the dollar auction game is SPE if and only if it is of the form

A always stops and B always continues

or symmetrically

B always stops and A always continues.

N.B. Details to be checked!

Further Results

We can give a **complete characterization** of the subgame-perfect equilibria for the dollar auction game.

Theorem

A strategy profile for the dollar auction game is SPE if and only if it is of the form

A always stops and B always continues

or symmetrically

B always stops and A always continues.

N.B. Details to be checked!

In fact this applies to any game sharing the following features of the dollar auction game:

1. At any point of the game, it is always better for a given player if the other player stops first.
2. At any point of the game, it is (strictly) better for the player who is the first to stop from that point to stop now rather than later.

Further Directions

Further Directions

- Lescanne proposes to use this form of analysis to explain the **rationality of infinite escalation**. If we don't know which strategy the other player is following, we always have an incentive to continue!

Further Directions

- Lescanne proposes to use this form of analysis to explain the **rationality of infinite escalation**. If we don't know which strategy the other player is following, we always have an incentive to continue!
- Many topics in economics which refer to infinite horizons and reflexivity seem tailor-made for the use of coalgebraic methods. At the same time, they can suggest new challenges and technical directions for coalgebra.