Calculating Games with Higher-Order Functions

Paulo Oliva

(based on joint work with M. Escardó)

Queen Mary, University of London, UK

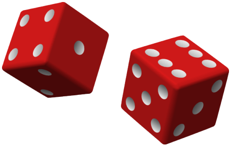

Southampton University

Wednesday, 12 November 2014

Outline

Outline

Game Theory

- Early development in the 19th century
- Formal approach with von Neumann (1930's)

John von Neumann

Game Theory

- Early development in the 19th century
- Formal approach with von Neumann (1930's)

- n players
- n strategy sets $X_1, \ldots, X_n$
- payoff function $q \colon \vec{X} \to \mathbb{R}^n$

John von Neumann

Game Theory

- Early development in the 19th century
- Formal approach with von Neumann (1930's)

- n players
- n strategy sets $X_1, \ldots, X_n$
- payoff function $q \colon \vec{X} \to \mathbb{R}^n$

John von Neumann

How should players choose their strategies

in order to maximise their individual payoffs?

Game Theory

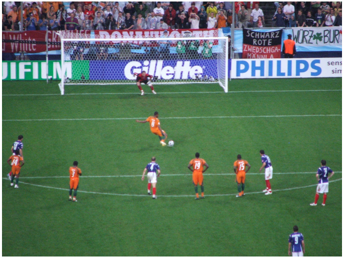

Game Theory

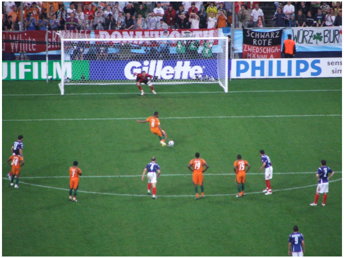

Penalties

Two players

Strategy sets $X_1 = X_2 = \{L, R\}$

Payoff function

f	L	R
L	$(1, 0)$	$(0, 1)$
R	$(0, 1)$	$(1, 0)$

Game Theory

- No **winning** strategy!
- What about strategies in **equilibrium**?

Game Theory

- No **winning** strategy!
- What about strategies in **equilibrium**?

Definition (Nash)

Strategy profile $\vec{x}$ is in equilibrium if no player has an incentive to unilaterally change his strategy

Game Theory

- No **winning** strategy!
- What about strategies in **equilibrium**?

Definition (Nash)

Strategy profile $\vec{x}$ is in equilibrium if no player has an incentive to unilaterally change his strategy

The "penalty" example shows that strategy profiles in equilibrium not necessarily exist either

Game Theory

- What if players choose "mixed" strategies
  i.e. player chooses probability distribution on strategies

Game Theory

- What if players choose "mixed" strategies
  i.e. player chooses probability distribution on strategies

Theorem (Nash)

Mixed strategies in equilibrium always exist

Game Theory

- What if players choose “mixed” strategies
  i.e. player chooses probability distribution on strategies

Theorem (Nash)

Mixed strategies in equilibrium always exist

The “penalty” example is again an illustration of this:
Players randomly choosing left or right is best they can do

Simultaneous versus Sequential Games

- That's all in the case of **simultaneous** games
- With **sequential** games things are simpler and nicer

- Strategies: mappings from previous moves to current move
- Similar definition of Nash equilibrium

Simultaneous versus Sequential Games

- That's all in the case of **simultaneous** games
- With **sequential** games things are simpler and nicer

- Strategies: mappings from previous moves to current move
- Similar definition of Nash equilibrium

But equilibrium always exists and can be computed
by a technique called **backward induction**

Backward Induction

$$q \colon X \times Y \times Z \to \mathbb{R}^3$$

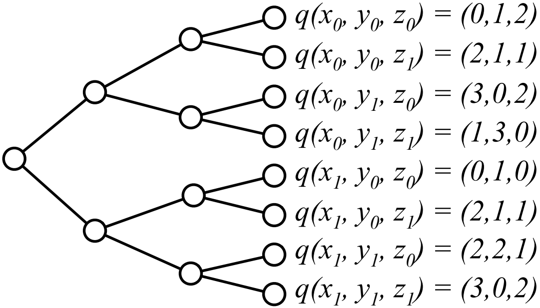

Backward Induction

$$q \colon X \times Y \times Z \to \mathbb{R}^3$$

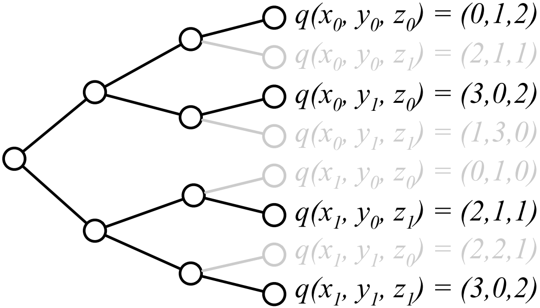

Backward Induction

$$q\colon X \times Y \times Z \to \mathbb{R}^3$$

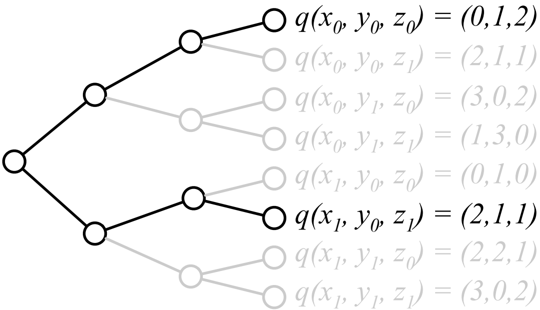

Backward Induction

$$q \colon X \times Y \times Z \to \mathbb{R}^3$$

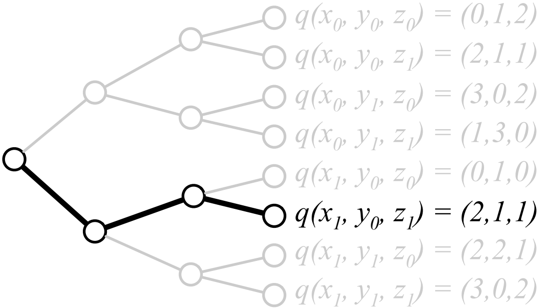

Our Recent Work

1. Generalised notions of sequential game, Nash equilibrium and backward induction

Our Recent Work

1. Generalised notions of sequential game, Nash equilibrium and backward induction
2. Showed how general notions appear in topology, proof theory, and algorithms, amongst others

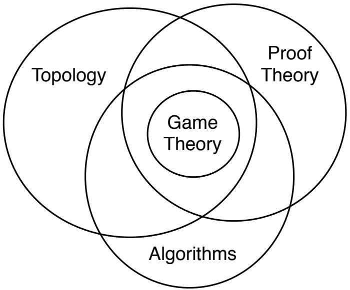

Outline

Single-player Games

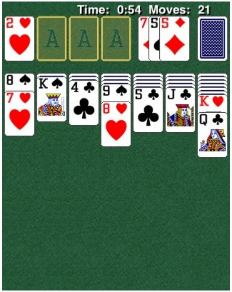

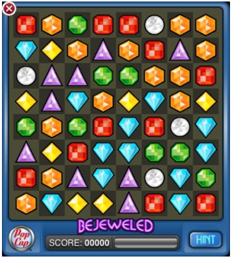

Two-player Games

Two **players**: Black and White

Two-player Games

Two **players**: Black and White

Possible **outcomes**:

- Black wins
- White wins
- Draw

Two-player Games

Two **players**: Black and White

Possible **outcomes**:

- Black wins
- White wins
- Draw

Strategy: Choice of move at round k given previous moves

Another Game

Two **players**: John and Julia

Another Game

Two **players**: John and Julia

John splits a cake. Julia chooses one of the two pieces

Another Game

Two **players**: John and Julia

John splits a cake. Julia chooses one of the two pieces

Possible **outcomes**:

- John gets $N\%$ of the cake (John's payoff)
- Julia gets $(100 - N)\%$ of the cake (Julia's payoff)

Another Game

Two **players**: John and Julia

John splits a cake. Julia chooses one of the two pieces

Possible **outcomes**:

- John gets $N\%$ of the cake (John's payoff)
- Julia gets $(100 - N)\%$ of the cake (Julia's payoff)

Best strategy for John is to split cake into half

It is not a "winning strategy" but it is an **optimal strategy**

It maximises his payoff

Number of Player vs Number of Rounds

Number of players is not essential

It is important what the “goal” at each round is

Rounds with “**same goal**” mean played by “**same player**”

Number of Player vs Number of Rounds

Number of players is not essential

It is important what the "goal" at each round is

Rounds with "**same goal**" mean played by "**same player**"

How to describe the goal at a particular round?

Number of Player vs Number of Rounds

Number of players is not essential

It is important what the "goal" at each round is

Rounds with "**same goal**" mean played by "**same player**"

How to describe the goal at a particular round?

You could say: The goal is to win!

But maybe this is not possible (or might not even make sense)

Instead, the goal should be described as:

a choice of outcome from each set of possible outcomes

As in...

Q: How much would you like to pay for your flight?

As in...

Q: How much would you like to pay for your flight?

A: As little as possible!

Quantifiers

R = set of outcomes

X = set of possible moves

$$\boxed{\phi \in (X \to R) \to R}$$

describes the desired outcome $\phi p \in R$ given $p \in X \to R$

Quantifiers

R = set of outcomes

X = set of possible moves

$$\phi \in (X \to R) \to R$$

describes the desired outcome $\phi p \in R$ given $p \in X \to R$

In the example:

R	=	*prices (real numbers)*
X	=	*possible flights*
$X \to R$	=	*price of each flight*
ϕ	=	*minimal value functional*

Quantifiers

$$\phi : (X \to R) \to R$$

Quantifiers

$$\phi : (X \to R) \to R$$

Other Examples

Operation	ϕ	:	$(X \to R) \to R$
Supremum	$\sup_{[0,1]}$	:	$([0,1] \to \mathbb{R}) \to \mathbb{R}$
Integration	$\int_0^1$	:	$([0,1] \to \mathbb{R}) \to \mathbb{R}$
Limit	$\lim$	:	$(\mathbb{N} \to R) \to R$
Quantifiers	$\forall_X, \exists_X$	:	$(X \to \mathbb{B}) \to \mathbb{B}$
Double negation	$\neg\neg X$	:	$(X \to \bot) \to \bot$
Fixed point operator	fix_X	:	$(X \to X) \to X$

Quantifiers

$$\phi : (X \to R) \to R \qquad (\equiv K_R X)$$

Other Examples

Operation	ϕ	:	$(X \to R) \to R$
Supremum	$\sup_{[0,1]}$	:	$([0,1] \to \mathbb{R}) \to \mathbb{R}$
Integration	$\int_0^1$	:	$([0,1] \to \mathbb{R}) \to \mathbb{R}$
Limit	$\lim$	:	$(\mathbb{N} \to R) \to R$
Quantifiers	$\forall_X, \exists_X$	:	$(X \to \mathbb{B}) \to \mathbb{B}$
Double negation	$\neg\neg X$	:	$(X \to \bot) \to \bot$
Fixed point operator	fix_X	:	$(X \to X) \to X$

Theorem (Maximum Value Theorem)

For any $p \in C[0,1]$ *there is a point* $a \in [0,1]$ *such that*

$$\sup p = p(a)$$

Theorem (Maximum Value Theorem)

For any $p \in C[0,1]$ *there is a point* $a \in [0,1]$ *such that*

$$\sup p = p(a)$$

Theorem (Mean Value Theorem)

For any $p \in C[0,1]$ *there is a point* $a \in [0,1]$ *such that*

$$\int_0^1 p = p(a)$$

Theorem (Witness Theorem)

For any $p \colon X \to \mathbb{B}$ there is a point $a \in X$ such that

$$\exists x^X p(x) \quad \Leftrightarrow \quad p(a)$$

(similar to Hilbert's ε-term).

Theorem (Witness Theorem)

For any $p \colon X \to \mathbb{B}$ there is a point $a \in X$ such that

$$\exists x^X p(x) \quad \Leftrightarrow \quad p(a)$$

(similar to Hilbert's ε-term).

Theorem (Counter-example Theorem)

For any $p \colon X \to \mathbb{B}$ there is a point $a \in X$ such that

$$\forall x^X p(x) \quad \Leftrightarrow \quad p(a)$$

(a is counter-example to p if one exists).

Let $J_R X \equiv (X \to R) \to X$

Let $J_R X \equiv (X \to R) \to X$

Definition (Selection Functions)

$\varepsilon \colon J_R X$ is called a **selection function** for $\phi \colon K_R X$ if

$$\phi(p) = p(\varepsilon p)$$

holds for all $p \colon X \to R$

Let $J_R X \equiv (X \to R) \to X$

Definition (Selection Functions)

$\varepsilon \colon J_R X$ is called a **selection function** for $\phi \colon K_R X$ if

$$\phi(p) = p(\varepsilon p)$$

holds for all $p \colon X \to R$

Definition (Attainable Quantifiers)

A quantifier $\phi \colon K_R X$ is called **attainable** if it has a selection function $\varepsilon \colon J_R X$

For Instance

- $\sup\colon K_{\mathbb{R}}[0,1]$ is an attainable quantifier

$$\sup(p) = p(\operatorname{argsup}(p))$$

where $\operatorname{argsup}\colon J_{\mathbb{R}}[0,1]$

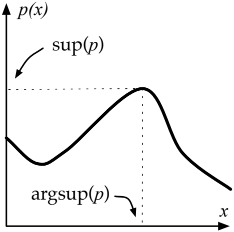

For Instance

- $\sup\colon K_{\mathbb{R}}[0,1]$ is an attainable quantifier

$$\sup(p) = p(\operatorname{argsup}(p))$$

where $\operatorname{argsup}\colon J_{\mathbb{R}}[0,1]$

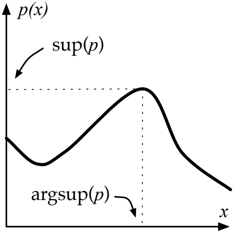

- $\mathsf{fix}\colon K_X X$ is an attainable quantifier

$$\mathsf{fix}(p) = p(\mathsf{fix}(p))$$

where $\mathsf{fix}\colon J_X X \; (= K_X X)$

Selection Functions and Quantifiers

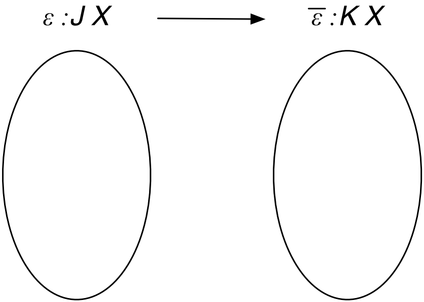

Every selection function $\varepsilon \colon J_R X$ defines a quantifier $\overline{\varepsilon} \colon K_R X$

$$\overline{\varepsilon}(p) \quad = \quad p(\varepsilon(p))$$

Selection Functions and Quantifiers

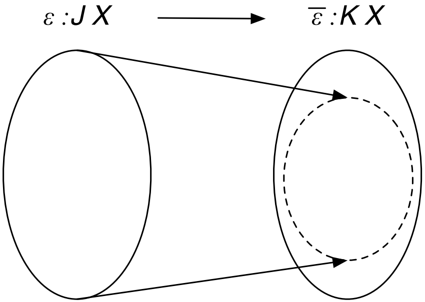

Not all quantifiers are attainable, e.g. $R = \{0, 1\}$

$$\phi(p) \quad = \quad 0$$

Selection Functions and Quantifiers

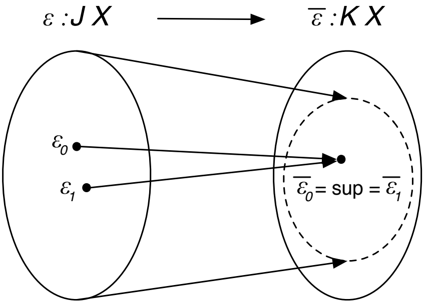

Different ε might define same ϕ, e.g. $X = [0,1]$ and $R = \mathbb{R}$

$$\varepsilon_0(p) = \mu x . \sup p = p(x)$$
$$\varepsilon_1(p) = \nu x . \sup p = p(x)$$

Outline

Finite Sequential Games (n rounds)

Definition (A tuple $(R, (X_i)_{i<n}, (\phi_i)_{i<n}, q)$ where)

- R is the set of **possible outcomes**
- X_i is the set of **available moves** at round i
- $\phi_i \colon K_R X_i$ is the **goal quantifier** for round i
- $q \colon \Pi_{i=0}^{n-1} X_i \to R$ is the **outcome function**

Finite Sequential Games (n rounds)

Definition (A tuple $(R, (X_i)_{i<n}, (\phi_i)_{i<n}, q)$ where)

- R is the set of **possible outcomes**
- X_i is the set of **available moves** at round i
- $\phi_i \colon K_R X_i$ is the **goal quantifier** for round i
- $q \colon \Pi_{i=0}^{n-1} X_i \to R$ is the **outcome function**

Definition (Strategy)

Family of mappings

$$\text{next}_k \colon \Pi_{i=0}^{k-1} X_i \to X_k$$

Definition (Strategic Play)

Given strategy next_k and partial play $\vec{a} = a_0, \ldots, a_{k-1}$, the **strategic extension** of $\vec{a}$ is $\mathbf{b}^{\vec{a}} = b^{\vec{a}}_k, \ldots, b^{\vec{a}}_{n-1}$ where

$$b^{\vec{a}}_i = \mathsf{next}_i(\vec{a}, b^{\vec{a}}_k, \ldots, b^{\vec{a}}_{i-1})$$

Definition (Strategic Play)

Given strategy next_k and partial play $\vec{a} = a_0, \ldots, a_{k-1}$, the **strategic extension** of $\vec{a}$ is $\mathbf{b}^{\vec{a}} = b^{\vec{a}}_k, \ldots, b^{\vec{a}}_{n-1}$ where

$$b^{\vec{a}}_i = \text{next}_i(\vec{a}, b^{\vec{a}}_k, \ldots, b^{\vec{a}}_{i-1})$$

Definition (Optimal Strategy)

Strategy next_k is **optimal** if for any partial play $\vec{a}$

$$q(\vec{a}, \mathbf{b}^{\vec{a}}) = \phi_k(\lambda x_k . q(\vec{a}, x_k, \mathbf{b}^{\vec{a}, x_k}))$$

Definition (Strategic Play)

Given strategy next_k and partial play $\vec{a} = a_0, \ldots, a_{k-1}$, the **strategic extension** of $\vec{a}$ is $\mathbf{b}^{\vec{a}} = b_k^{\vec{a}}, \ldots, b_{n-1}^{\vec{a}}$ where

$$b_i^{\vec{a}} = \mathsf{next}_i(\vec{a}, b_k^{\vec{a}}, \ldots, b_{i-1}^{\vec{a}})$$

Definition (Optimal Strategy)

Strategy next_k is **optimal** if for any partial play $\vec{a}$

$$q(\vec{a}, \mathbf{b}^{\vec{a}}) = \phi_k(\lambda x_k . q(\vec{a}, x_k, \mathbf{b}^{\vec{a}, x_k}))$$

A product of selection functions computes optimal strategies

Standard Game Theory

When quantifiers are $\max$ or $\sup$ over finite or compact set

Then argsup exists (and hence $\sup$ is attainable)

Generalised Game	$\mapsto$	Standard Game
Optimal strategy	$\mapsto$	Strategy in Nash equilibrium
Product of argsup	$\mapsto$	Backward induction!

Fixed Point Theory

Fixed point operators are their own selection function

Generalised Game	$\mapsto$	Operators on product space
Optimal strategy	$\mapsto$	Bekiç's Lemma
Product of fix's	$\mapsto$	The proof!

Proof Theory

Proof interpretation

$$\exists i \leq n \forall x^{X_i} \exists r^R A_i(x, r) \quad \mapsto \quad \forall \varepsilon_{(\cdot)} \exists i \leq n \exists p A_i(\varepsilon_i p, p(\varepsilon_i p))$$

Proof Theory

Proof interpretation

$$\exists i \leq n \forall x^{X_i} \exists r^R A_i(x, r) \quad \mapsto \quad \underline{\forall \varepsilon_{(\cdot)}} \exists i \leq n \exists p A_i(\varepsilon_i p, p(\varepsilon_i p))$$

ε's define quantifiers, which partially define a game

Computational interpretation relies on completing the definition of the game so optimal strategy solves problem

Proof Theory

Proof interpretation

$$\exists i \leq n \forall x^{X_i} \exists r^R A_i(x, r) \quad \mapsto \quad \underline{\forall \varepsilon_{(\cdot)}} \exists i \leq n \exists p A_i(\varepsilon_i p, p(\varepsilon_i p))$$

ε's define quantifiers, which partially define a game

Computational interpretation relies on completing the definition of the game so optimal strategy solves problem

Existence of optimal strategy actually implies the consistency of mathematics!

Outline

Monads

- K_R and J_R are **strong monads**

Monads

- K_R and J_R are **strong monads**
- $J_R \mapsto K_R$ is a **monad morphism**

Monads

- K_R and J_R are **strong monads**
- $J_R \mapsto K_R$ is a **monad morphism**
- **Product** of quantifiers

$$K_R X \times K_R Y \to K_R(X \times Y)$$

calculates optimal outcome

Monads

- K_R and J_R are **strong monads**
- $J_R \mapsto K_R$ is a **monad morphism**
- **Product** of quantifiers

$$K_R X \times K_R Y \to K_R(X \times Y)$$

  calculates optimal outcome
- **Product** of selection functions

$$J_R X \times J_R Y \to J_R(X \times Y)$$

  calculates optimal play

Monads

- K_R and J_R are **strong monads**
- $J_R \mapsto K_R$ is a **monad morphism**
- **Product** of quantifiers

$$K_R X \times K_R Y \to K_R(X \times Y)$$

  calculates optimal outcome
- **Product** of selection functions

$$J_R X \times J_R Y \to J_R(X \times Y)$$

  calculates optimal play
- Infinite product $\Pi_i J_R X_i \to J_R \Pi_i X_i$ exists (*in some models*)

Summary

- Generalised notion of sequential game
- Generalised notion of optimal strategy (equilibrium)
- Product of sel. fct. computes optimal strategies
- Results from fixed point theory, topology, proof theory, etc, can be viewed as optimal strategies in certain games

References

M. Escardó and P. Oliva
Selection functions, bar recursion and backward induction
MSCS, 20(2):127-168, 2010

M. Escardó and P. Oliva
What sequential games, the Tychnoff theorem and the double-negation shift have in common
ACM SIGPLAN MSFP, ACM Press 2010

M. Escardó and P. Oliva
Sequential games and optimal strategies
Proceedings of the Royal Society A, 2011

M. Escardó and P. Oliva
Computing Nash equilibria of unbounded games
The Turing Centenary Conference, 2012