Compositional game theory

Jules Hedges

(University of Oxford)

SYCO 1, Birmingham
21 September 2018

A peek at where we're going

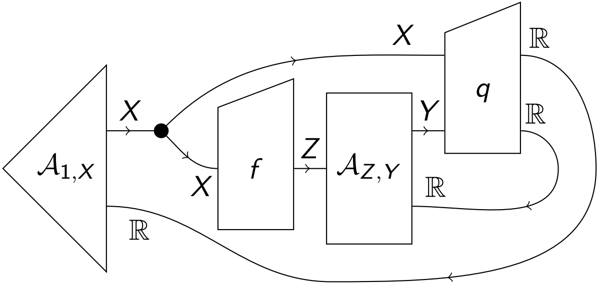

Game theory

- Mathematical theory of interacting "rational" agents
- Players make observations and then make choices
- Group choices determine payoffs
- "Local view" of rationality: players act to maximise payoff
- "Global view": equilibrium strategies

Example: penalty shootout

$$a, b \in \{L, R\}$$

Example: penalty shootout

$$a, b \in \{L, R\}$$

$$\pi(a, b) = \begin{cases} (+1, -1) & \text{if } a \neq b \\ (-1, +1) & \text{if } a = b \end{cases}$$

Example: penalty shootout

$$a, b \in \{L, R\}$$

$$\pi(a,b) = \begin{cases} (+1, -1) & \text{if } a \neq b \\ (-1, +1) & \text{if } a = b \end{cases}$$

Unique (probabilistic) equilibrium: $a = b = \frac{1}{2}|L\rangle + \frac{1}{2}|R\rangle$

Example: penalty shootout

$$a, b \in \{L, R\}$$

$$\pi(a, b) = \begin{cases} (+1, -1) & \text{if } a \neq b \\ (-1, +1) & \text{if } a = b \end{cases}$$

Unique (probabilistic) equilibrium: $a = b = \frac{1}{2} |L\rangle + \frac{1}{2} |R\rangle$

Nash's theorem generalises this situation

Picturing game theory (1945 – 2018)

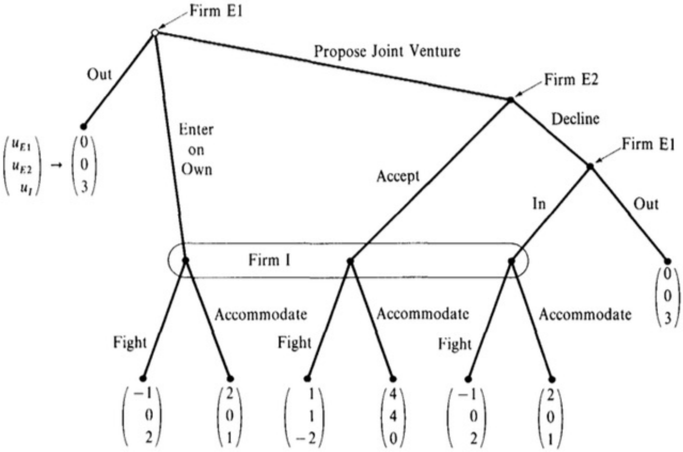

Game theory has some issues

- Well known: equilibrium as behavioural prediction is experimentally falsified (e.g. ultimatum game)

Game theory has some issues

- Well known: equilibrium as behavioural prediction is experimentally falsified (e.g. ultimatum game)
- Harsanyi type spaces are accurate but underfit (and mathematically hard!)

Game theory has some issues

- Well known: equilibrium as behavioural prediction is experimentally falsified (e.g. ultimatum game)
- Harsanyi type spaces are accurate but underfit (and mathematically hard!)
- There is no accepted operational theory (or "equilibriating process") (c.f. evolutionary game theory)

Game theory has some issues

- Well known: equilibrium as behavioural prediction is experimentally falsified (e.g. ultimatum game)
- Harsanyi type spaces are accurate but underfit (and mathematically hard!)
- There is no accepted operational theory (or "equilibriating process") (c.f. evolutionary game theory)
- Serious computability/complexity issues (algorithmic game theory)

Game theory has some issues

- Well known: equilibrium as behavioural prediction is experimentally falsified (e.g. ultimatum game)
- Harsanyi type spaces are accurate but underfit (and mathematically hard!)
- There is no accepted operational theory (or "equilibriating process") (c.f. evolutionary game theory)
- Serious computability/complexity issues (algorithmic game theory)
- Ordinary games do not compose/scale

The fundamental headache of social science

Beliefs have causal effects

Defining **PC**

PC is a category where:

- Objects are pairs of sets $\binom{X}{S}$
- Morphisms $\lambda : \binom{X}{S} \to \binom{Y}{R}$ are pairs of functions:
  - $v_\lambda : X \to Y$
  - $u_\lambda : X \times R \to S$

λ is called a lens

Defining **PC**

PC is a category where:

- Objects are pairs of sets $\binom{X}{S}$
- Morphisms $\lambda : \binom{X}{S} \to \binom{Y}{R}$ are pairs of functions:
  - $v_\lambda : X \to Y$
  - $u_\lambda : X \times R \to S$

λ is called a lens

We draw it like this:

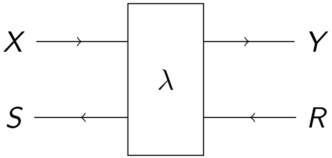

Intuition for **PC**

Approximately . . .

- First part: physical information
  - X and Y are sets of things an agent can observe or choose

Intuition for **PC**

Approximately . . .

- First part: physical information
    - X and Y are sets of things an agent can observe or choose
- Second part: teleological or counterfactual information
    - R and S are sets of things an agent can optimise or have preferences about

Intuition for **PC**

Approximately . . .

- First part: physical information
  - X and Y are sets of things an agent can observe or choose
- Second part: teleological or counterfactual information
  - R and S are sets of things an agent can optimise or have preferences about

A typical example:

- $f : X \to Y$ is a function
- Promote to $\lambda : \binom{X}{\mathbb{R}} \to \binom{Y}{\mathbb{R}}$ with $v_\lambda = f$
- $u_\lambda : X \times \mathbb{R} \to \mathbb{R}$ is backpropagation of value
- If we know x and we know the value of $f(x)$ then u_λ tells us what the value of x was

Example: a decision process

(aka. a Markov decision process without the probability)
Take a state space S, actions A, transition function
$f : S \times A \to S \times \mathbb{R}$

Example: a decision process

(aka. a Markov decision process without the probability)

Take a state space S, actions A, transition function $f : S \times A \to S \times \mathbb{R}$

Every policy function $\sigma : S \to A$ determines a lens $\lambda : \binom{S}{\mathbb{R}} \to \binom{S}{\mathbb{R}}$ by

- $v_\lambda(s) = f(s, \sigma(s))_1$
- $u_\lambda(s, u) = f(s, \sigma(s))_2 + \beta \cdot u$
- $0 < \beta < 1$ is discount factor

Composing lenses

Given

$$\begin{pmatrix} X \\ S \end{pmatrix} \xrightarrow{\lambda} \begin{pmatrix} Y \\ R \end{pmatrix} \xrightarrow{\mu} \begin{pmatrix} Z \\ Q \end{pmatrix}$$

we can compose them to $\mu \circ \lambda : \binom{X}{S} \to \binom{Z}{Q}$

(Important non-obvious fact: this is associative)

Composing lenses

Given

$$\binom{X}{S} \xrightarrow{\lambda} \binom{Y}{R} \xrightarrow{\mu} \binom{Z}{Q}$$

we can compose them to $\mu \circ \lambda : \binom{X}{S} \to \binom{Z}{Q}$

(Important non-obvious fact: this is associative)

Given $\binom{X_1}{S_1} \xrightarrow{\lambda_1} \binom{Y_1}{R_1}$ and $\binom{X_2}{S_2} \xrightarrow{\lambda_2} \binom{Y_2}{R_2}$ we can compose them to

$$\begin{pmatrix} X_1 \times X_2 \\ S_2 \times S_1 \end{pmatrix} \xrightarrow{\lambda_1 \otimes \lambda_2} \begin{pmatrix} Y_1 \times Y_2 \\ R_2 \times R_1 \end{pmatrix}$$

PC is a symmetric monoidal category

Special lenses

$f : X \to Y$ lifts to $f : \binom{X}{1} \to \binom{Y}{1}$ or $f^* : \binom{1}{Y} \to \binom{1}{X}$

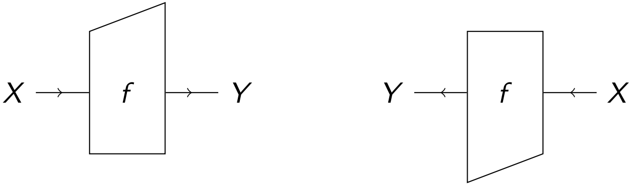

Special lenses

$f : X \to Y$ lifts to $f : \binom{X}{1} \to \binom{Y}{1}$ or $f^* : \binom{1}{Y} \to \binom{1}{X}$

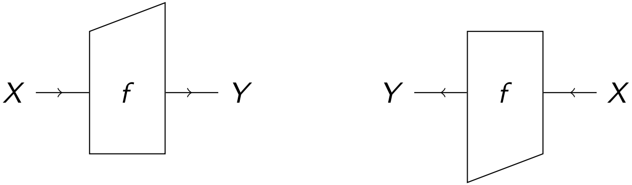

Special case: Every $\binom{X}{1}$ is a comonoid, every $\binom{1}{X}$ is a monoid

Special lenses

$f : X \to Y$ lifts to $f : \binom{X}{1} \to \binom{Y}{1}$ or $f^* : \binom{1}{Y} \to \binom{1}{X}$

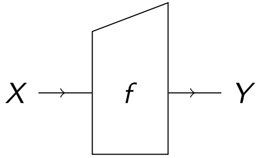

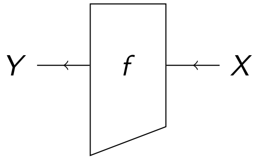

Special case: Every $\binom{X}{1}$ is a comonoid, every $\binom{1}{X}$ is a monoid

There is canonical $\varepsilon_X : \binom{X}{X} \to \binom{1}{1}$ (but no η!)

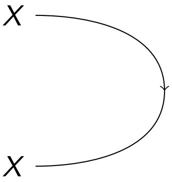

The counit law

Theorem:

$$\varepsilon_Y \circ ((f, 1) \otimes (1, \mathrm{id}_Y)) = \varepsilon_X \circ ((\mathrm{id}_X, 1) \otimes (1, f))$$

aka:

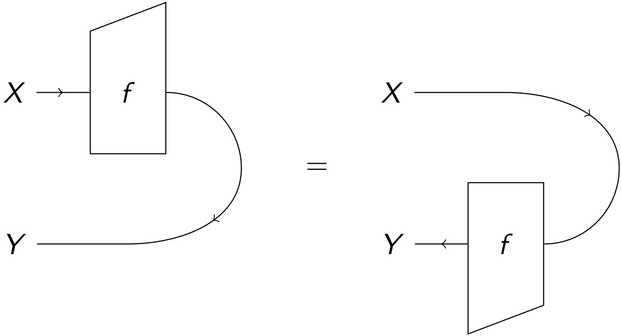

Interesting facts about **PC**

- **PC** is a dialectica category over a 1-valued logic
  - hence, a sound model of linear logic

Interesting facts about **PC**

- **PC** is a dialectica category over a 1-valued logic
  - hence, a sound model of linear logic
- $\binom{X}{S} \mapsto X$, $\lambda \mapsto v_\lambda$ is a fibration
  - It's fibrewise opposite of Jacobs' simple fibration

Interesting facts about **PC**

- **PC** is a dialectica category over a 1-valued logic
  - hence, a sound model of linear logic
- $\binom{X}{S} \mapsto X$, $\lambda \mapsto v_\lambda$ is a fibration
  - It's fibrewise opposite of Jacobs' simple fibration
- Hot off the press: **PC** is complete (if underlying cat is complete, cocomplete, cartesian closed, . . .)
  - Work in progress: game theory using **Span**(**PC**)

Interesting facts about **PC**

- **PC** is a dialectica category over a 1-valued logic
  - hence, a sound model of linear logic
- $\binom{X}{S} \mapsto X$, $\lambda \mapsto v_\lambda$ is a fibration
  - It's fibrewise opposite of Jacobs' simple fibration
- Hot off the press: **PC** is complete (if underlying cat is complete, cocomplete, cartesian closed, . . .)
  - Work in progress: game theory using **Span**(**PC**)
- Really hot off the press: **PC** can be defined over a monoidal category:

$$\text{hom}_{\mathbf{PC}(\mathcal{C})}\left(\binom{X}{S}, \binom{Y}{R}\right) = \int^{A \in \mathcal{C}} \text{hom}_{\mathcal{C}}(X, A \otimes Y) \times \text{hom}_{\mathcal{C}}(A \otimes R, S)$$

  - Needed for probabilistic open games etc
  - Universal property: "freely adding counits"
  - Mitchell Riley, Categories of Optics, arXiv

The context functors

- $\mathbb{V} : \mathbf{PC} \to \mathbf{Set}$, $(X, S) \mapsto X$, $\ell \mapsto v_\ell$
  - It's the view fibration of a lens
  - $\mathbb{V} \cong \mathrm{hom}_{\mathbf{PC}}(I, -)$

The context functors

- $\mathbb{V} : \mathbf{PC} \to \mathbf{Set}$, $(X, S) \mapsto X$, $\ell \mapsto v_\ell$
  - It's the view fibration of a lens
  - $\mathbb{V} \cong \hom_{\mathbf{PC}}(I, -)$
- $\mathbb{K} : \mathbf{PC}^{\mathrm{op}} \to \mathbf{Set}$, $(X, S) \mapsto X \to S$
  - The continuation functor
  - $\mathbb{K} \cong \hom_{\mathbf{PC}}(-, I)$

The context functors

- $\mathbb{V} : \mathbf{PC} \to \mathbf{Set}$, $(X, S) \mapsto X$, $\ell \mapsto v_\ell$
  - It's the view fibration of a lens
  - $\mathbb{V} \cong \mathrm{hom}_{\mathbf{PC}}(I, -)$
- $\mathbb{K} : \mathbf{PC}^{\mathrm{op}} \to \mathbf{Set}$, $(X, S) \mapsto X \to S$
  - The continuation functor
  - $\mathbb{K} \cong \mathrm{hom}_{\mathbf{PC}}(-, I)$

Slogan: points are states, continuations are effects

Defining open games

An open game $\mathcal{G} : \binom{X}{S} \to \binom{Y}{R}$ consists of:

- A set $\Sigma_{\mathcal{G}}$ of strategy profiles

Defining open games

An open game $\mathcal{G} : \binom{X}{S} \to \binom{Y}{R}$ consists of:

- A set $\Sigma_\mathcal{G}$ of strategy profiles
- For every $\sigma : \Sigma_\mathcal{G}$, a lens $\mathcal{G}(\sigma) : \binom{X}{S} \to \binom{Y}{R}$

Defining open games

An open game $\mathcal{G} : \binom{X}{S} \to \binom{Y}{R}$ consists of:

- A set $\Sigma_\mathcal{G}$ of strategy profiles
- For every $\sigma : \Sigma_\mathcal{G}$, a lens $\mathcal{G}(\sigma) : \binom{X}{S} \to \binom{Y}{R}$
- For every context $(h, k) : \mathbb{V}\binom{X}{S} \times \mathbb{K}\binom{Y}{R}$, a set $\mathbf{E}_\mathcal{G}(h, k) \subseteq \Sigma_\mathcal{G}$ of Nash equilibria

Defining open games

An open game $\mathcal{G} : \binom{X}{S} \to \binom{Y}{R}$ consists of:

- A set $\Sigma_\mathcal{G}$ of strategy profiles
- For every $\sigma : \Sigma_\mathcal{G}$, a lens $\mathcal{G}(\sigma) : \binom{X}{S} \to \binom{Y}{R}$
- For every context $(h, k) : \mathbb{V}\binom{X}{S} \times \mathbb{K}\binom{Y}{R}$, a set $\mathbf{E}_\mathcal{G}(h, k) \subseteq \Sigma_\mathcal{G}$ of Nash equilibria

Things that have been abstracted away: players, moves, payoffs, maximisation

Defining open games

An open game $\mathcal{G} : \binom{X}{S} \to \binom{Y}{R}$ consists of:

- A set $\Sigma_\mathcal{G}$ of strategy profiles
- For every $\sigma : \Sigma_\mathcal{G}$, a lens $\mathcal{G}(\sigma) : \binom{X}{S} \to \binom{Y}{R}$
- For every context $(h, k) : \mathbb{V}\binom{X}{S} \times \mathbb{K}\binom{Y}{R}$, a set $\mathbf{E}_\mathcal{G}(h, k) \subseteq \Sigma_\mathcal{G}$ of Nash equilibria

Things that have been abstracted away: players, moves, payoffs, maximisation

We draw it like this:

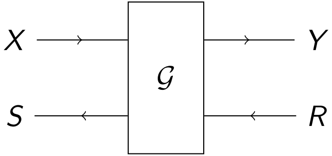

Special open games

A zero player open game has $\Sigma_{\mathcal{G}} = 1$ and $\mathbf{E}_{\mathcal{G}}(h, k) = \{*\}$ for all (h, k)

- Zero-player open games $\binom{X}{S} \to \binom{Y}{R}$ are in bijection with lenses $\binom{X}{S} \to \binom{Y}{R}$

Special open games

A zero player open game has $\Sigma_{\mathcal{G}} = 1$ and $\mathbf{E}_{\mathcal{G}}(h, k) = \{*\}$ for all (h, k)

- Zero-player open games $\binom{X}{S} \to \binom{Y}{R}$ are in bijection with lenses $\binom{X}{S} \to \binom{Y}{R}$

A scalar open game is an open game $\binom{1}{1} \to \binom{1}{1}$

- They are determined by a set of strategy profiles, and a subset of Nash equilibria
- Every ordinary (eg. extensive form) game determines a scalar open game

Sequential play

Suppose we have open games

$$\binom{X}{S} \xrightarrow{\mathcal{G}} \binom{Y}{R} \xrightarrow{\mathcal{H}} \binom{Z}{Q}$$

We define $\mathcal{H} \circ \mathcal{G} : \binom{X}{S} \to \binom{Z}{Q}$ like this:

- $\Sigma_{\mathcal{H} \circ \mathcal{G}} = \Sigma_{\mathcal{G}} \times \Sigma_{\mathcal{H}}$
- $(\mathcal{H} \circ \mathcal{G})(\sigma, \tau) = \mathcal{H}(\tau) \circ \mathcal{G}(\sigma)$
- The magic part:

$$\mathbf{E}_{\mathcal{H} \circ \mathcal{G}}(h, k) = \left\{ (\sigma, \tau) \,\middle|\, \begin{array}{l} \sigma \in \mathbf{E}_{\mathcal{G}}(h, \mathbb{K}(\mathcal{H}(\tau))(k)) \\ \tau \in \mathbf{E}_{\mathcal{H}}(\mathbb{V}(\mathcal{G}(\sigma))(h), k) \end{array} \right\}$$

Example

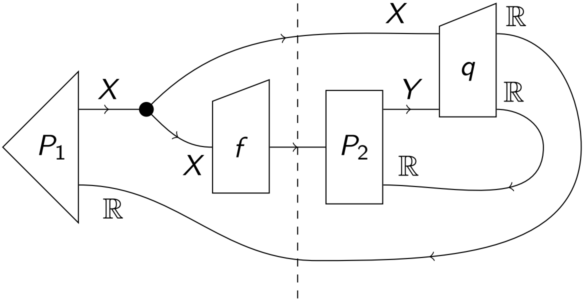

$\mathcal{G} : (1, 1) \to (X \times Z, \mathbb{R})$

- $\Sigma_{\mathcal{G}} = X$
- $v_{\mathcal{G}(x)}(*) = (x, f(x))$
- $\mathbf{E}_{\mathcal{G}}(*, k) = \arg\max_x k(x, f(x))$

$\mathcal{H} : (X \times Z, \mathbb{R}) \to (1, 1)$

- $\Sigma_{\mathcal{H}} = Z \to Y$
- $u_{\mathcal{H}(\sigma)}((x, z), *) = q_1(x, \sigma(z))$
- $\mathbf{E}_{\mathcal{H}}((x, z), *) = \{\sigma \mid \sigma(z) \in \arg\max_y q_2(x, y)\}$

Simultaneous play

. . . is more complicated, cut for time

Finitely generated games

Define an open game $\mathcal{A}_{X,Y} : \binom{X}{1} \to \binom{Y}{\mathbb{R}}$ by

- $\Sigma_{\mathcal{A}_{X,Y}} = X \to Y$
- $v_{\mathcal{A}_{X,Y}(\sigma)} = \sigma$
- $\mathbf{E}_{\mathcal{A}_{X,Y}}(h, k) = \{\sigma \mid \sigma(h) \in \arg\max(k)\}$

It's (a single decision by) an agent

N.B. This is the only place we mention $\mathbb{R}$ or $\arg\max$!

Finitely generated games

Define an open game $\mathcal{A}_{X,Y} : \binom{X}{1} \to \binom{Y}{\mathbb{R}}$ by

- $\Sigma_{\mathcal{A}_{X,Y}} = X \to Y$
- $v_{\mathcal{A}_{X,Y}(\sigma)} = \sigma$
- $\mathbf{E}_{\mathcal{A}_{X,Y}}(h, k) = \{\sigma \mid \sigma(h) \in \arg\max(k)\}$

It's (a single decision by) an agent

N.B. This is the only place we mention $\mathbb{R}$ or $\arg\max$!

Fundamental theorem of compositional game theory: The following are in (sensible) bijective correspondence:

1. Scalar open games finitely generated by zero-player open games, $\mathcal{A}_{X,Y}$, $\circ$ and $\otimes$
2. Strategy profiles & pure Nash equilibria of finite-depth extensive form games of imperfect information

Bimatrix game

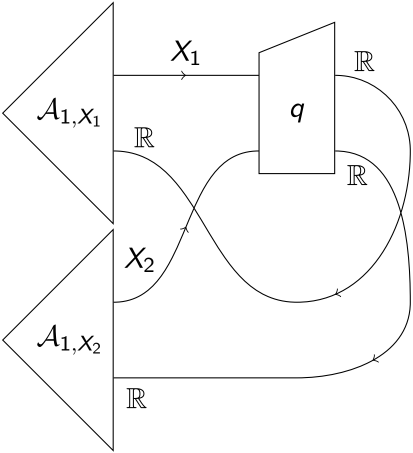

Sequential game of perfect information

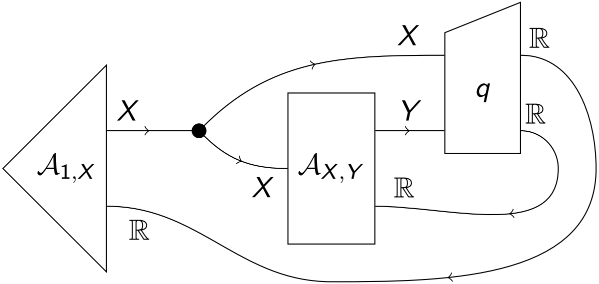

Sequential game of imperfect information

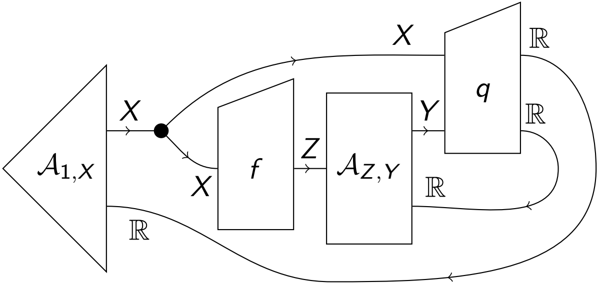

Hybrid sequential-simultaneous game

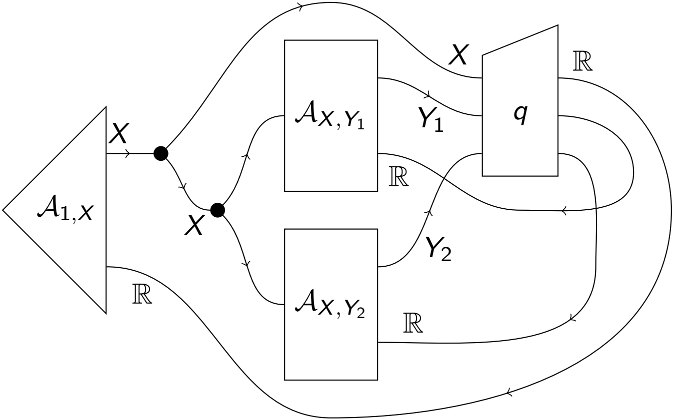

Cool stuff in the past

- Morphisms of open games, version 1:
  - infinitely repeated games are final coalgebras

Cool stuff in the past

- Morphisms of open games, version 1:
  - infinitely repeated games are final coalgebras
- Morphisms between open games, version 2:
  - Nash equilibria are states
  - Subgame perfect equilibria are $\otimes$-separable states
  - Products are external choice

Cool stuff in the past

- Morphisms of open games, version 1:
  - infinitely repeated games are final coalgebras
- Morphisms between open games, version 2:
  - Nash equilibria are states
  - Subgame perfect equilibria are $\otimes$-separable states
  - Products are external choice
- Bayesian open games
  - (not released yet)
  - Unexpectedly hard

Cool stuff in the future

- Compositional economic modelling

Cool stuff in the future

- Compositional economic modelling
- Composing numerical solution methods

Cool stuff in the future

- Compositional economic modelling
- Composing numerical solution methods
- Connections with learning
  - Using deep learning to cheat complexity theory

Cool stuff in the future

- Compositional economic modelling
- Composing numerical solution methods
- Connections with learning
    - Using deep learning to cheat complexity theory
- Open graphical games

Cool stuff in the future

- Compositional economic modelling
- Composing numerical solution methods
- Connections with learning
  - Using deep learning to cheat complexity theory
- Open graphical games
- Getting a compact closed category
  - Version 1: $\mathbf{PC} \hookrightarrow \mathbf{Int}$
  - Version 2: $\mathbf{PC} \hookrightarrow \mathbf{Span}(\mathbf{PC})$