Games with players

Towards categorical foundations of cybernetics

An MSP101 talk

Matteo Capucci (jwww MSP101.ACT)

April 29th, 2021 (Day 488 of the COVID Era)

University of Strathclyde Glasgow

Intro

This talk gathers ideas from the last 6-24 months regarding game theory, machine learning, cybernetics, etc.

Intro

This talk gathers ideas from the last 6-24 months regarding game theory, machine learning, cybernetics, etc.

Most of these are the brainchild of a whole group of people in the group – I just packaged them together with a bow

Intro

This talk gathers ideas from the last 6-24 months regarding game theory, machine learning, cybernetics, etc.

Most of these are the brainchild of a whole group of people in the group – I just packaged them together with a bow

Plan

1. **Game theory 101** – concepts and terminology
2. **Open games with players** – from scratch
3. **Post-credit scene**: cybernetics

Game theory 101

What is a game?

Informal definition:

> *Game theory is the mathematical study of interaction among independent, self-interested agents.*
>
> *– Essentials of Game Theory [LS08]*

Examples:

1. Tic-tac-toe, chess, Monopoly, etc.
2. Economic games (includes/are included in: ecological games)
3. Social dilemmas (PD, 'tragedy of the commons', etc.)
4. Proof theory, model theory, etc.
5. Machine learning
6. **etc.**

Representing games

1. **Normal form**: A set of **players** P, an indexed set of **actions** $A : P \to$ **Set**, a **utility function** $u : \prod_{p \in P} A\ p \to (P \to R)$
2. **Extensive form**: A set of **players** P, a **tree** representing the unfolding of the game. Nodes are assigned to players and grouped in **information sets**. Branches are called **moves**. A **utility vector** assigned to each leaf.

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

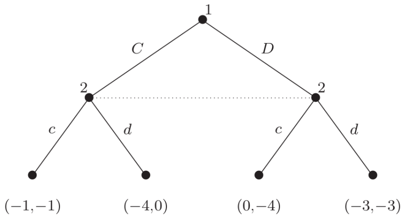

Extensive → normal

One can always convert an extensive form game into normal form:

1. Define
$$A\ p = \prod_{x \in p\text{'s nodes}} \text{moves at } x$$
2. Define
$$u(a_1, \ldots, a_n) \quad = \quad \text{leaf at the end of the path}$$
$$\text{root} \to a_1 \to a_2(a_1) \to \ldots \to a_n(\cdots a_2(a_1))$$

The converse is not always possible since normal-form games have too little structural information.

Solving games

Pre-formal definition: A **solution concept** is a notion of 'optimality' for ways to play a game.

A 'way to play' for a player $p \in P$ is called **strategy**:

$$\Omega\ p = \prod_{x \in p\text{'s nodes}} \text{moves at } x$$

Strategies are a **comprehensive plan of action**: for each **state** of the game, no matter how unlikely, we plan an **action**.

A choice of strategy for each player is a **strategy profile**:

$$S = \prod_{p \in P} \Omega\ p$$

Nash equilibrium

The most important (and general) solution concept is **Nash equilibrium**:

Definition

A strategy profile $s \in S$ is a Nash equilibrium if no player has interest in unilaterally deviating its strategy.

e.g. for utility-maximizing players:

$$\forall p \in P, \forall s'_p \in \Omega\ p \quad u_i(s[s_p/s'_p]) \leq u_p(s)$$

It's not the only one: SGP, ESS, ϵ-Nash, trembling hand, etc.
Afaik, all are **refinements** of Nash.

Nash equilibrium: example

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

Pros and cons

Problems with classical game theory:

1. Games are treated **monolithically**
2. Stuck in **early 20th century mathematical language**
3. Denotations are quite disappointing: normal form is too opaque, extensive form is too... extended

Open games are a proposed improvement:

1. Games are defined **compositionally**, including **equilibria**
2. Mathematically more sophisticated (grounded in **category theory**)
3. Denoted by **string diagrams**: halfway between normal and extensive form

It follows the ACT tradition of 'opening up' systems: *always consider a system as part of an environment it interacts non-trivially with*

Open games ‘2.0’

Open games

Warning: Open games are compositional structures, hence the single building blocks do not make much sense from a classical standpoint – you have to put them together to get something meaningful!

Intuitively, an 'atomic' open game is a forest of bushes:

Back and forth

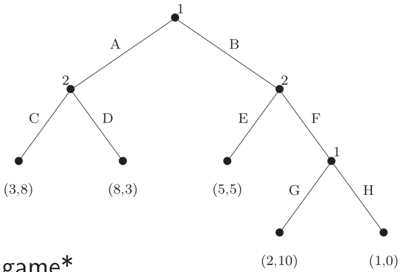

There's two phases in a game*

1. The '**forward phase**'

   Players take turns and make their own decisions until a leaf is reached

2. The '**backward phase**'

   Payoffs propagate back to players along the tree

   $\longrightarrow$ **backward induction**

*handwaving important philosophical point here

Lenses and bidirectional information flows

A **lens** models exactly this bidirectional information flow:

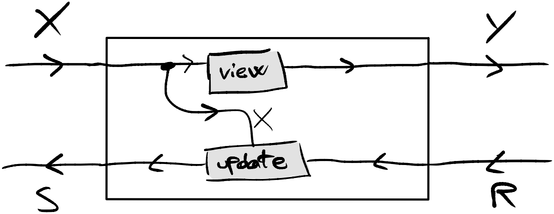

Definition

Let **C** be a cartesian category (think: sets & functions).

A **lens** $(X, S) \to (Y, R)$ is a pair of maps

$$\text{view} : X \to Y, \qquad \text{update} : X \times R \to S$$

They can be generalized greatly, see **optics** [Ril18]

Games as lenses

A game can be represented naively as a lens

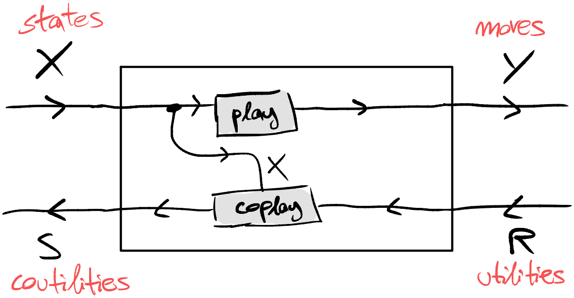

What's coutility?

> *For a given node x in a game, a player's continuation value (also called* **continuation payoff***) is the payoff that this* **player will eventually get** *contingent on the path of play* **passing through node** *x. – Strategy [Wat02]*

The coplay function takes on this job. In classical game theory it's hidden in backward induction.

It doesn't *have to* be trivial, but it often is.

Sequential composition

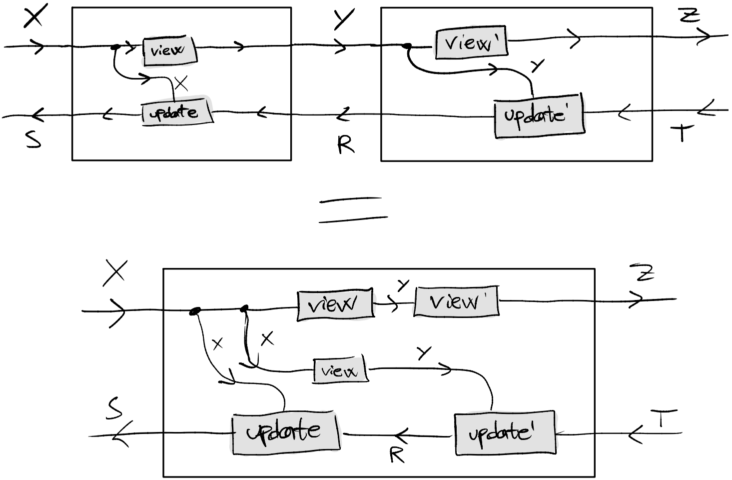

e.g. chess

Parallel composition

Play simultaneously

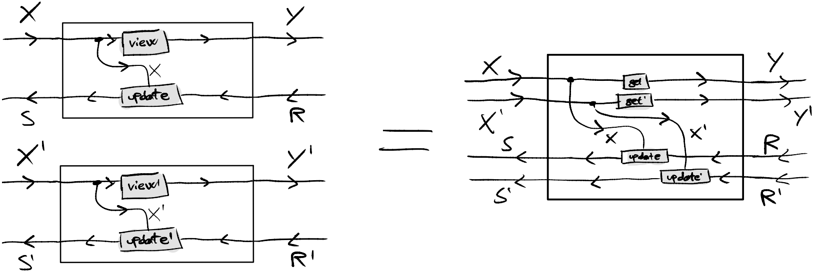

e.g. PD

Closing a game

We can use sequential composition to give a lens a **context**, i.e. an initial **state** and a **payoff function**

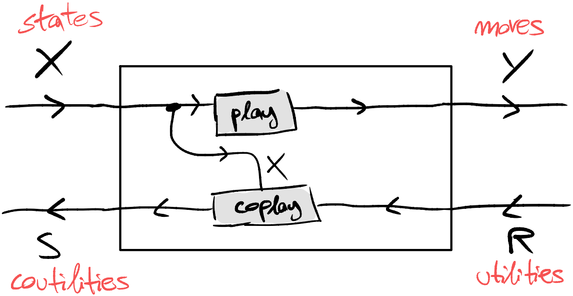

Closing a game

We can use sequential composition to give a lens a **context**, i.e. an initial **state** and a **payoff function**

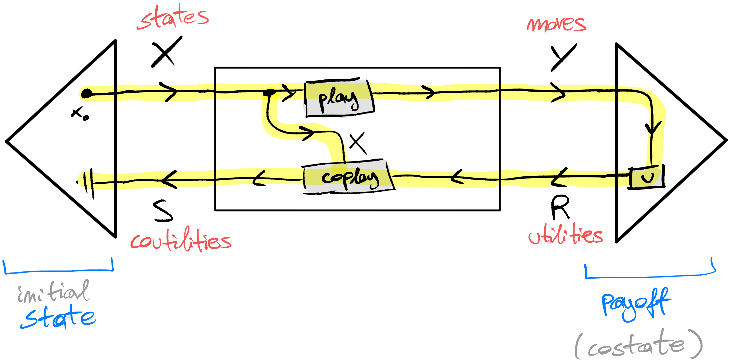

Slogan: '**Time flows clockwise**'

Open games

Recall:

> *Game theory is the mathematical study of* **interaction** *among independent, self-interested* **agents**.
>
> *– Essentials of Game Theory [LS08]*

A game factors in two parts

1. An **arena**, which models the interaction patterns in the game
2. A set of **agents**, i.e. the players, which make **decisions** at different points of a game

Without (2) a game would be only a **dynamical system**, whose dynamic is fixed. Instead, in a game **agents can vary the dynamics in response to the observed unfolding of the interaction**.

Interlude I: the Para construction

Originally from [FST19], but expanded greatly in the last year

Definition

When **C** is symmetric monoidal, **Para**(**C**) is the category of parametrized morphisms of **C**:

1. objects are the same,
2. a morphism $A \to B$ is given by a choice of parameter $P : \mathbf{C}$ and a choice of morphism $f : P \otimes A \to B$ in **C**:

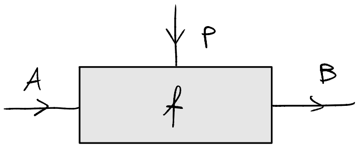

Interlude I: the Para construction

Para(C) is again symmetric monoidal:

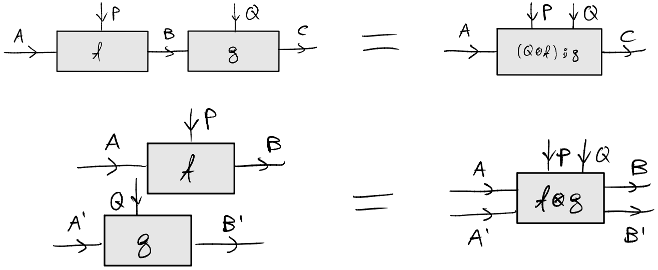

Interlude I: the Para construction

Most importantly, it's a **bicategory**:

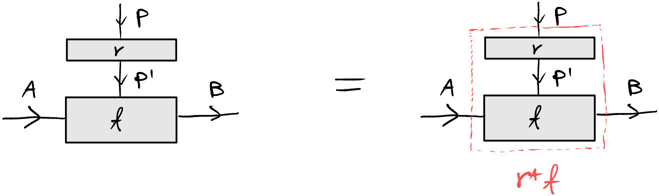

Interlude I: the Para construction

If our morphisms are ‘bidirectional’, we get an even more interesting picture:

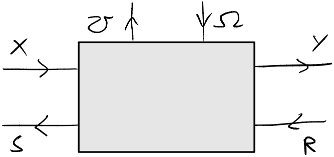

If we peek inside, we can see the new information flow:

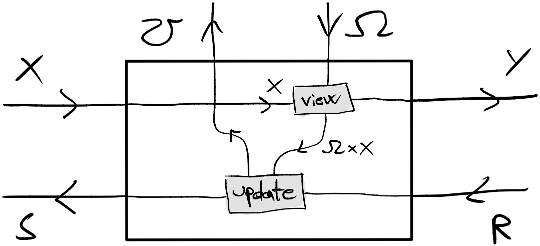

Putting players in a game

Let's go back to games...

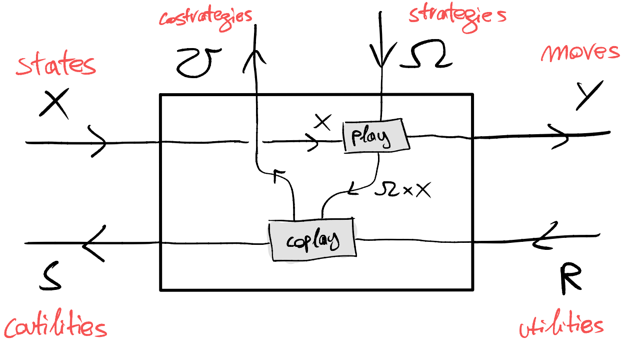

Open games

Slogan: **agents live in the parametrization direction**

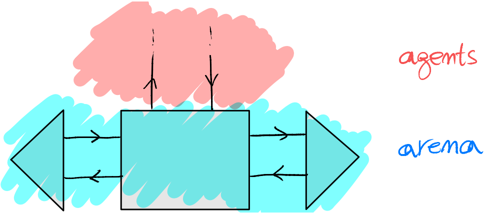

The **arena** plays the role of a costate to them:

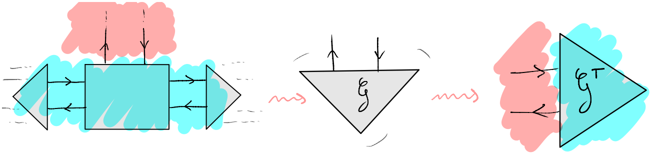

$-^\top$ (**transposition**) is the **arena lifting** (dynamics $\to$ agents) operation.

Open games

Arenas can be defined '**locally**': it's just information plumbing. Agents can't: an agent might observe and interact with the arena at multiple, causally 'distant' points.

We take advantage of the 2-cells in **Para**(**Lens**) to handle this:

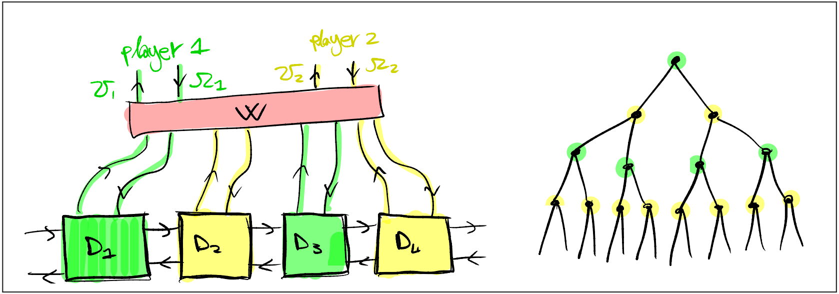

Agency in open games

> *What does it mean to say that agents are self-interested? [...] It means that each agent has* **their own description** *of which* **states of the world they like**—*which can include good things happening to other agents—and that* **they act** *in an attempt to bring about these states of the world.*
>
> *– Essentials of Game Theory [LS08]*

Agents have:

1. a way to **act in the world**
2. a way to **observe the world**
3. a way to **evaluate the world**

Agency in open games

> *What does it mean to say that agents are self-interested? [...] It means that each agent has* **their own description** *of which* **states of the world they like**—*which can include good things happening to other agents—and that* **they act** *in an attempt to bring about these states of the world.*
>
> *– Essentials of Game Theory [LS08]*

Agents have:

1. a way to **act in the world** $\rightsquigarrow$ **strategies**
2. a way to **observe the world** $\rightsquigarrow$ **costrategies**
3. a way to **evaluate the world**

Agency in open games

> *What does it mean to say that agents are self-interested? [...] It means that each agent has* **their own description** *of which* **states of the world they like**—*which can include good things happening to other agents—and that* **they act** *in an attempt to bring about these states of the world.*
>
> *– Essentials of Game Theory [LS08]*

Agents have:

1. a way to **act in the world** $\rightsquigarrow$ **strategies**
2. a way to **observe the world** $\rightsquigarrow$ **costrategies**
3. a way to **evaluate the world** $\rightsquigarrow$ **selection function**

Interlude II: selection functions

Definition

A **continuation** on a object X with scalars an object R is a map

$$K_R(X) = (X \to R) \to R$$

It's a 'generalized quantifier': max, min, $\exists$, $\forall$

If the ambient category is cartesian closed, K_R defines a monad.

Interlude II: selection functions

Selection functions ‘realize’ quantifiers:

Definition

Def. A **selection function** on a object X with scalars an object R is a map

$$J_R(X) = (X \to R) \to X$$

Examples: argmax, argmin, Hilbert's ε

If the ambient category is cartesian closed, J_R defines a monad.

Interlude II: selection functions

Notice: often quantifiers are realized by multiple elements...

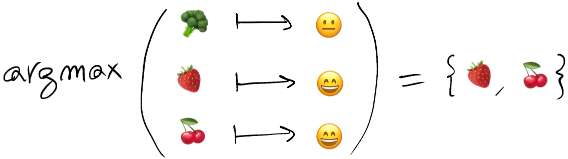

So a better type for selection functions is

$$(X \to R) \to \mathsf{P}X$$

where P is the powerset monad.

Interlude II: selection functions

Also notice:

$$(X \to R) \longrightarrow \mathrm{P}X$$

costates of (X, R) $\qquad$ P(states of (X, R))

so we arrive to a general definition:

Definition

Let **C** be a monoidal category. Then the selection functions functor is given by

$$\mathbf{Sel} : \mathbf{C} \longrightarrow \mathbf{Cat}$$
$$X \longmapsto \mathbf{C}(X, I) \to \mathrm{P}\mathbf{C}(I, X)$$

The codomain is **Cat** since this set is ordered by pointwise inclusion:

$$\varepsilon \leq \varepsilon' \quad \underline{\text{iff}} \quad \forall k \in \mathbf{C}(X, I),\ \varepsilon(k) \subseteq \varepsilon'(k)$$

Interlude II: selection functions

Sel is a functor because it also acts on morphism by **pushforward**:

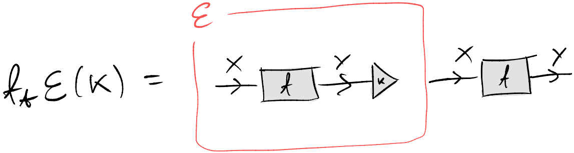

Idea: selection functions are a relation between states and costates.

(**Probably better idea**: selection functions are predicates on *contexts*)

Interlude II: selection functions

Finally, **Sel** is **lax monoidal** with **Nash product**:

$$- \boxtimes - : \mathbf{Sel}(X) \times \mathbf{Sel}(Y) \to \mathbf{Sel}(X \otimes Y)$$

$$(\varepsilon \boxtimes \eta)(k) = \{x \otimes y \in (X \otimes y)_* \varepsilon(k) \cap (x \otimes Y)_* \eta(k)\}$$

Best ‘unilateral deviations’:

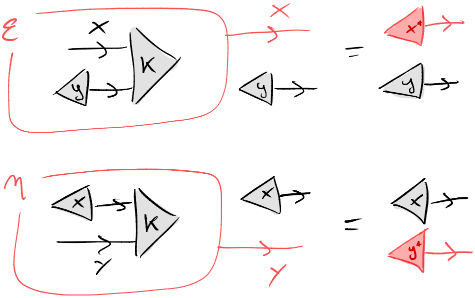

Agency in open games

Idea: An agent's interest is embodied by a selection function on their arenas:

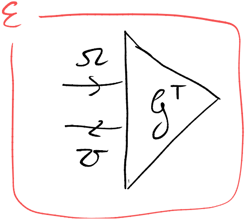

Multiple agents ⊠ together their selection to select common arenas:

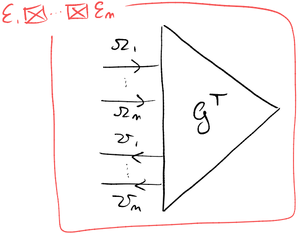

Open games

So an open game is

Definition

1. A **parametrized lens**:

$$\mathcal{G} : (X, S) \xrightarrow{(\Sigma, \Sigma')} (Y, R)$$

2. A set of **selection functions** indexed by player P:

$$\forall p \in P, \ \varepsilon_p \in \textbf{Sel}(\Omega_p, \mho_p)$$

3. A **wiring 2-cell**:

$$w : \prod_{p \in P} (\Omega_p, \mho_p) \longrightarrow (\Sigma, \Sigma')$$

Open games

As anticipated, equilibria are given by

$$\mathrm{eq}_{\mathcal{G}}(x, u) = (\boxtimes_{p \in P}\, \varepsilon_p)(w \,⨾\, (x \,⨾\, \mathcal{G} \,⨾\, u)^\top) \qquad \subseteq \prod_{p \in P} \Omega_p$$

Still compositional wrt sequential and parallel composition of arenas:

$$\mathrm{eq}_{\mathcal{G} ⨾ \mathcal{H}}(x, u) = \{\omega \otimes \xi \mid \omega \in \mathrm{eq}_{\mathcal{G}}(x, \mathcal{H}(\xi ⨾ w') ⨾ u) \wedge \xi \in \mathrm{eq}_{\mathcal{H}}(x ⨾ \mathcal{G}(\omega ⨾ w), u)\}$$

$$\mathrm{eq}_{\mathcal{G} \otimes 1}((x, \_), u) = \mathrm{eq}_{\mathcal{G}}(x, u)^*$$

It recovers Nash equilibria as a solution concept, which justifies calling $\boxtimes$ **Nash product**.

*harder to generalize to monoidal categories / optics

Open games

This solves long-standing problems with open games:

1. We finally compute the **right set of Nash equilibria** (instead of one-shot deviations)
2. We can better handle situations of **imperfect information**
3. We can define **internal choice** (perhaps: approach to **cooperative game theory**)

A lot to explore!

Example I: cooperative Prisoner's Dilemma

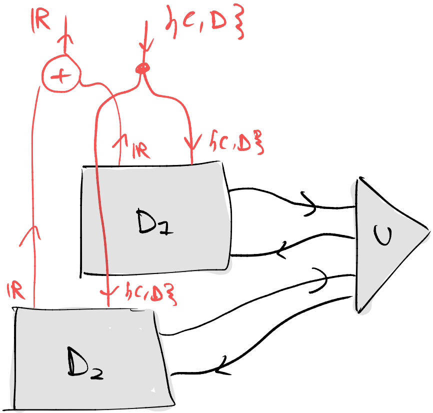

Example II: stopped game

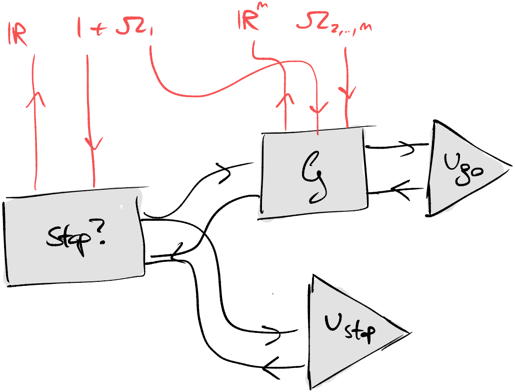

Example III: imperfect recall

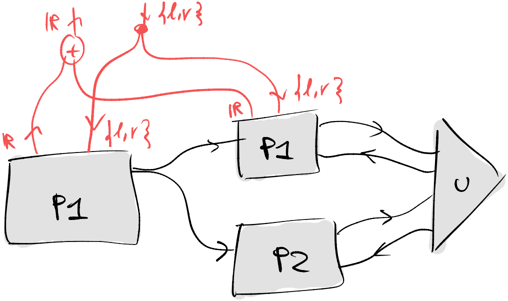

Example IV: internal choice

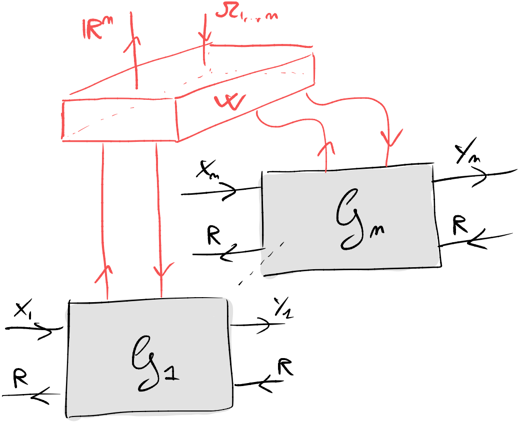

Thanks for your attention!

Questions?

$\longrightarrow$ **post-credit scene**

Cybernetics

Cybernetics

From $\kappa\upsilon\beta\varepsilon\rho\nu\alpha\omega$ (*to govern, to steer*):

> *Science concerned with the study of systems of any nature which are capable of receiving, storing and processing information so as to use it for control.* – *Kolmogorov*

Lenses model dynamical systems (see: [Mye20])

Parametrized lenses (+ decorations) model cybernetic systems!

The missing bits are **storage** and **feedback**.

$\underbrace{\text{Parametrized} \cdots \text{parametrized}}_{n \text{ times}}$ lenses model **...?**

Hierarchical agency?

Higher-order cybernetics

Agents

1. act in the arena,
2. then observe the result of their behaviour,
3. then change their action accordingly,

until an equilibrium is reached.

Higher-order cybernetics

Agents

1. act in the arena,
2. then **observe** the result of their behaviour,
3. then change their action accordingly,

until an equilibrium is reached.

	1st	**2nd**	**3rd**	$\cdots$
non-trivial observation		yes	yes	$\cdots$
non-trivial analysis			yes	$\cdots$
$\vdots$				$\ddots$

Periodic table of cybernetic types

Reasoning negatively:

	-2nd	-1st	0th	1st	2nd	···
non-trivial system		yes	yes	yes	yes	···
non-trivial context			yes	yes	yes	···
non-trivial interaction				yes	yes	···
non-trivial observation					yes	···
⋮						⋱

Games vs. learners

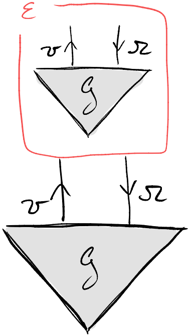

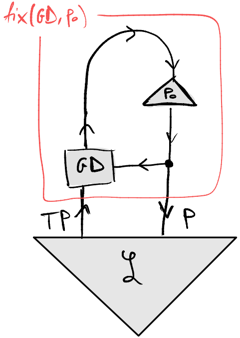

A learner produces its own selection function as a fixpoint:

$$p \in \mathrm{eq}(x, u) \quad \underline{\text{iff}} \quad p = p \mathbin{⨟} GD \mathbin{⨟} w \mathbin{⨟} (x \mathbin{⨟} \mathcal{L} \mathbin{⨟} u)^{\top}$$

2nd order cybernetics

Given a parameter $p \in P$, a learner can only observe $\mathcal{L}^\top$ on an infinitesimal neighbourhood of p

$\rightsquigarrow$ **2nd-order cybernetic systems**

We can consider **Para**(**Lens**(**Smooth**)) a **2nd-order cybernetic doctrine** (terminology borrowed from [Mye])

This is actually a strength:

1. We can encode the selection in the parameter dynamics
2. We can analyze locally and iteratively
   (vs. games 'global and one-step')

References I

Kevin Leyton-Brown and Yoav Shoham. "Essentials of game theory: A concise multidisciplinary introduction". In: *Synthesis lectures on artificial intelligence and machine learning* 2.1 (2008), pp. 1–88.

Mitchell Riley. "Categories of Optics". In: *arXiv: Category Theory* (2018).

Joel Watson. *Strategy: an introduction to game theory*. Vol. 139. WW Norton New York, 2002.

Brendan Fong, David Spivak, and Rémy Tuyéras. "Backprop as Functor: A compositional perspective on supervised learning". In: *2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. IEEE. 2019, pp. 1–13.

References II

David Jaz Myers. "Double Categories of Open Dynamical Systems". In: *arXiv preprint arXiv:2005.05956* (2020).

David Jaz Myers. *Categorical Systems Theory*. URL: https://github.com/DavidJaz/DynamicalSystemsBook/tree/master/book.

Neil Ghani et al. "Compositional game theory". In: *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science*. 2018, pp. 472–481.

G.S.H. Cruttwell et al. "Categorical Foundations of Gradient-Based Learning". In: *arXiv preprint arXiv:2103.01931* (2021).

References III

Martin Escardó and Paulo Oliva. "Selection functions, bar recursion and backward induction". In: *Math. Struct. Comput. Sci.* 20.2 (2010), pp. 127–168.