

Diegetic representation of feedback in open games

Matteo Capucci

MSP Group, University of Strathclyde, Scotland

We improve the framework of open games with agency [10, 4] by showing how the players' counterfactual analysis giving rise to Nash equilibria can be described in the dynamics of the game itself (hence *diegetically*), getting rid of devices such as equilibrium predicates. This new approach overlaps almost completely with the way gradient-based learners [6] are specified and trained. Indeed, we show feedback propagation in games can be seen as reverse-mode differentiation, with a crucial difference explaining the distinctive character of the phenomenology of non-cooperative games. We outline a functorial construction of arena of games, show players form a subsystem over it, and prove that their 'fixpoint behaviours' are Nash equilibria.

1 Motivation

In narratology, *diegetic* is what exists or occurs within the world of a narrative [8] (such as dialog, thoughts, etc.), as opposed to *extra-diegetic elements* which happens outside that world (such as voiceovers, soundtrack, etc.). Open games represent the situations of classical game theory in a compositional and purportedly 'diegetic' way, i.e. explicitly codifying the development of the game actions and payoff distribution phases in their specification. Hedges proposed a framework in [13] which evolved first by adopting the language of lenses, and then that of parametric lenses to describe the bidirectional flow of information in games. In their last iteration [4, 3], *open games with agency* are defined to be given by three functions (for concreteness, we assume to work in **Set**):

$$\text{play}_{\mathcal{G}} : \Omega \times X \rightarrow Y, \quad \text{coplay}_{\mathcal{G}} : \Omega \times X \times R \rightarrow S \times \mathfrak{U}, \quad \varepsilon_{\mathcal{G}} : (\Omega \rightarrow \mathfrak{U}) \rightarrow P\Omega. \quad (1.1)$$

The set Ω represent *strategies*, X and Y *states* of the game, while R and S *utility* and '*coutility*', respectively. Finally, $\varepsilon_{\mathcal{G}}$ is a *selection function* that encodes a player's preferences: given a valuation of strategies in \mathfrak{U} (called *costrategies* or *intrinsic utility*), $\varepsilon_{\mathcal{G}}$ returns the subset of strategies with

satisfactory outcome. This data defines a parametric lens [4]:

$$\mathcal{G} = (\Omega, \mathcal{U}, \varepsilon_{\mathcal{G}}, \text{play}_{\mathcal{G}}, \text{coplay}_{\mathcal{G}}) : (X, S) \rightleftarrows (Y, R). \quad (1.2)$$

However, the dynamics of payoff distribution in open games sets the latter aside from other examples of cybernetic systems¹ described with parametric optics [3], chiefly learners [6] and partly Bayesian reasoners [20, 2]. This is evident in the way the computation of Nash equilibria is defined for open games, which relies on ‘packing up’ the arena² of the game and then running the players’ characteristic counterfactual analysis *extra-diegetically*, thus outside the arena, not recognizing it as part of the game itself.




Figure 1

This issue grows into a serious conceptual flaw when we realize that according to the very notion of ‘system with agency’ proposed by the author and his collaborators in [3], ‘open games with agency’ have no agents! In fact, agents are supposed to be systems modelled as *morphisms* in the vertical systems theory³ of a category of parametric optics, while in open games with agency players’ preferences are embodied in the parameters, which are *objects* in the vertical systems theory (Figure 1a). Contrast this with gradient-based learners (Figure 1b), where gradient descent (which is truly the agent’s learning process) is explicitly represented *in* the system.

The lack of a principled payoff dynamics in open games shows in the way coutilities, costategies and utilities are all different in theory but very rarely in practice, and coplay is very often simply an

¹Here we call ‘cybernetic’ those systems exhibiting agency and interactive feedbacks.

²Borrowing terminology from [4], the *arena* of an open game is the parametric lens left by taking away the equilibrium predicate/selection function.

³Parametric optics representing systems with agency organize in bicategories, and by ‘vertical systems theory’ we mean the 2-morphisms thereof. An example is the box GD in Figure 1b.

identity or, even worse, a discard map, which makes it hard to motivate the existence of a backward pass at all (see e.g. the translation process explained in [4]).

Contributions. In this work we correct the aforementioned problems by describing the entirety of play, payoff distribution and players' counterfactual analysis diegetically, that is, in the dynamics of the game system itself, thereby making open games look much more like gradient-based learners or Bayesian reasoners than they do now.

We achieve this by making two fundamental changes.

First, we observe that feedback propagating in a strategic game has to contain information about the entirety of the payoff function of the game, hence we replace S and R in Figure 1a with P^X and P^Y , where P is a specified payoff object. This allows to define coplay functorially from play as precomposition with a partially-evaluated play. This simple mechanism is enough to reproduce the information on payoffs available at each stage of a sequential or concurrent game. Moreover, we recognize the crucial role of the lax monoidal structure of this functor, which can be blamed for the complexity of even small game-theoretic situations.

Secondly, we describe how players are embodied inside the game by their selection functions, which are now expressed as parts of a 'reparameterisation' describing each player's optimization dynamic. This fully realizes what was already intuited in [14] ('agents are their selection function') and in the drawings in [3, §6], as well as, morally, the idea behind open games with agency introduced in [4]. In fact we find out that the workhorse of open games with agency, the Nash product of selection functions, is decomposed in three more elementary parts, the key one being 'just' tensor product of lenses.

We then show how this story shares many formal analogies with a refinement on the description of gradient-based learners. There is a formal analogy between covectors and payoff vectors, reverse derivatives and coplays, 'raising indices' and selection functions. Ultimately this suggests that data describing how to functorially produce feedback-processing (dependent) lenses from their forward parts is a fundamental constituent of the definition of a 'theory of cybernetic systems' which experiences backpropagation-like feedback.

Acknowledgements. We thank the rest of MSP group for their invaluable feedback. In particular, we thank Jules Hedges, Philipp Zahn, Neil Ghani and Bruno Gavranović for their helpful suggestions and enthusiasm towards this work. We thank the ACT22 reviewers for their patience in reviewing an early version of this manuscript.

2 Diegetic open games

We start by describing our proposed notion of *diegetic open games*. As anticipated, the key idea is to recognize that in a strategic game, players have to observe the entirety of their payoff functions with other players' actions taken into account. This is done by fixing utility, costility and intrinsic utility types to be of the form P^Y , P^X and P^Ω , representing entire payoff functions. Then such functions are propagated through the game in a way which is formally identical to backpropagation in learners, and thus amenable to the same mathematical treatment. Thus $\text{coplay}_{\mathcal{G}}$ is actually functorially determined from $\text{play}_{\mathcal{G}}$, as a kind of reverse derivative.

2.1 Preliminaries

Fix a finitely complete category \mathcal{S} . The category $\mathbf{DLens}(\mathcal{S})$ of *dependent lenses* over \mathcal{S} has objects given by pairs of an object $Y : \mathcal{S}$ and a map $p : R \rightarrow Y$, and maps given by diagrams of the form:

$$\begin{array}{ccc} S & \xleftarrow{f^\sharp} & R \times_Y X & \longrightarrow & R \\ \downarrow & & \downarrow f^*(p) & \lrcorner & \downarrow p \\ X & \xlongequal{\quad} & X & \xrightarrow{f} & Y \end{array} \quad (2.1)$$

In the internal language of \mathcal{S} [19], these maps can be denoted as $f : X \rightarrow Y$ and $f_x^\sharp : (x : X) \times Rf(x) \rightarrow Sx$. The full subcategory of $\mathbf{DLens}(\mathcal{S})$ spanned by those p which are projections is the category of *simple lenses* over \mathcal{S} , $\mathbf{Lens}(\mathcal{S})$. The f^\sharp part of simple lenses has a simpler type, namely $X \times R \rightarrow S$.

Dependent lenses can be built from any indexed category $F : \mathcal{S}^{\text{op}} \rightarrow \mathbf{Cat}$, in which case we denote them by $\mathbf{DLens}(F)$. A detailed definition and intuition is given in [21].

The 2-category $\mathbf{Para}(\mathcal{S})$ [3, §2] is the strictification of the bicategory whose objects are given by objects of \mathcal{S} , morphisms X to Y by a choice of parameter $\Omega : \mathcal{S}$ and a map $f : \Omega \times X \rightarrow Y$, and 2-morphisms $(\Omega, f) \Rightarrow (\Xi, g) : X \rightarrow Y$ by maps $\Omega \rightarrow \Xi$ making the obvious triangle commute (see *loc. cit.*, though we have reversed the direction of 2-cells here), which are called *reparameterisations*. Composition of morphisms $(\Omega, f) : X \rightarrow Y$ and $(\Xi, g) : Y \rightarrow Z$ is given by

$$(\Xi \times \Omega, \Xi \times \Omega \times X \xrightarrow{\Xi \times f} \Xi \times Y \xrightarrow{g} Z) \quad (2.2)$$

This makes it associative only up to coherent isomorphism, hence the strictification. Same applies to the identities, which are given by $(1, 1 \times X \xrightarrow{\pi_X} X)$.

Notice the construction of $\mathbf{Para}(\mathcal{S})$ only used the cartesian monoidal structure of \mathcal{S} . In fact such a construction is functorial over cartesian monoidal categories. Given a lax monoidal functor [15, Definition 1.2.14] $F : \mathcal{S} \rightarrow \mathcal{T}$, with laxators $\ell_{X,Y} : F(X) \times F(Y) \rightarrow F(X \times Y)$, we get a lax 2-functor [15, Definition 4.1.2] $\mathbf{Para}(F) : \mathbf{Para}(\mathcal{S}) \rightarrow \mathbf{Para}(\mathcal{T})$ defined on objects as F and on a morphisms $(\Omega, f) : X \rightarrow Y$ as

$$\mathbf{Para}(F)(\Omega, f) = (F(\Omega), F(\Omega) \times F(X) \xrightarrow{\ell_{\Omega, X}} F(\Omega \times X) \xrightarrow{F(f)} Y). \quad (2.3)$$

Since $\ell_{\Omega, X}$ is, in principle, not invertible, this means $\mathbf{Para}(F)$ preserves composition only up to coherent non-invertible morphism. Explicitly, there is a reparameterisation $\mathbf{Para}(F)(\Omega, f) \circledast \mathbf{Para}(F)(\Xi, g) \Rightarrow \mathbf{Para}(F)((\Omega, f) \circledast (\Xi, g))$, given by $\ell_{\Xi, \Omega}$. Likewise applies to preservation of identities. The well-definedness of these reparameterisations follow from the axioms of lax monoidal structure ℓ [15, Diagram 1.2.14].

2.2 Building arenas

Let us position ourselves in **Set**. We now describe the most simple form of games, deterministic, complete information games, with our new machinery. We remark later on how we can abstract away most of the assumptions we make here.

Fixing a *payoff object* P (often $P = \mathbb{R}^N$, with N the number of players), to a map $f : X \rightarrow Y$ we can associate the map $P^f : P^Y \rightarrow P^X$ given by precomposition with f . This defines a functor $P^{(-)} : \mathbf{Set} \rightarrow \mathbf{Set}^{\text{op}}$, which we can lift to a lax monoidal functor

$$P^* : \mathbf{Set} \longrightarrow \mathbf{Lens}(\mathbf{Set}) \quad (2.4)$$

sending $f : X \rightarrow Y$ to $(f, \pi_2 \circledast P^f) : (X, P^X) \rightleftarrows (Y, P^Y)$. Abusing notation, we'll denote by P^*f both this lens and its backward part, and same with objects: $P^*X := P^X$. Notice landing in lenses is crucial to give P^* a lax monoidal structure: while its unitor $\eta : (1, 1) \rightleftarrows (1, P)$, given by $(1, !_P)$ would be definable anyway; the laxator $(1_{X,Y}, \mathbf{n}_{X,Y}) : (X, P^*X) \otimes (Y, P^*Y) \rightleftarrows (X \times Y, P^*(X \times Y))$, which we call *Nashator*, is defined by partial evaluation at the residuals:

$$\begin{aligned} \mathbf{n}_{X,Y} : X \times Y \times P^*(X \times Y) &\longrightarrow P^*X \times P^*Y \\ (\bar{x}, \bar{y}, u) &\longmapsto \langle u(-, \bar{y}), u(\bar{x}, -) \rangle \end{aligned} \quad (2.5)$$

Ideally, this functor promotes a play function into a lens obtained by canonically adding a coplay function; but since open games are *parametric* lenses, we need to apply **Para** to P^* to obtain the lax 2-functor

$$\mathbf{Para}(P^*) : \mathbf{Para}(\mathbf{Set}) \longrightarrow \mathbf{Para}(\mathbf{Lens}(\mathbf{Set})) \quad (2.6)$$

so that a play function $(\Omega, \text{play}_{\mathcal{G}}) : X \rightarrow Y$ is turned into a full-blown parametric lens:

$$\mathbf{Para}(P^*)(\Omega, \text{play}_{\mathcal{G}}) = (\Omega, P^*\Omega, (1_{\Omega, X}, \mathbf{n}_{\Omega, X}) \mathbin{\text{\textcircled{;}}} (\text{play}_{\mathcal{G}}, P^*\text{play}_{\mathcal{G}})) \quad (2.7)$$

where the backward part of the right hand side boils down to

$$\begin{aligned} \mathbf{Para}(P^*)(\Omega, \text{play}_{\mathcal{G}})^{\sharp} : \Omega \times X \times P^*Y &\longrightarrow P^*\Omega \times P^*X \\ (\bar{\omega}, \bar{x}, u) &\longmapsto \langle u_{\Omega}, u_X \rangle \text{ where } u_{\Omega} = u(\text{play}_{\mathcal{G}}(\bar{x}, -)) \\ &u_X = u(\text{play}_{\mathcal{G}}(-, \bar{\omega})) \end{aligned} \quad (2.8)$$

This definition is the workhorse of diegetic open games. Notice how u_X encapsulates $\bar{\omega}$ as a fixed parameter, so that an opponent receiving such function later has that strategy fixed. Dually, u_{Ω} has \bar{x} fixed so the player playing at this stage can vary their own strategy but not the state the game is at, something determined, in turn, by other players' strategies.

Example 2.1 (Pure sequential game). Consider a very simple game in which two players make one move each, in succession. The first player has strategy space Ω and play function $(\Omega, \text{play}_{\mathcal{G}}) : X \rightarrow Y$, whereas the second player has strategies Ξ and play $(\Xi, \text{play}_{\mathcal{H}}) : Y \rightarrow Z$:




Figure 2

Figure 2 depicts the parametric lens $\mathbf{Para}(P^*)(\Omega, \text{play}_{\mathcal{G}}) \mathbin{\text{\textcircled{;}}} \mathbf{Para}(P^*)(\Xi, \text{play}_{\mathcal{H}})$. This is what we call the *arena* of the game.

Suppose a $\bar{x} \in X$ and a $u \in P^*Z$ are given, so as to close the open input horizontal wires in **Figure 2**. These two pieces of data amount to a so-called *context* for the game, and mathematically correspond to a further (trivially parameterised) lens $(\bar{x}, !_{P^*X}) : (1, 1) \rightleftarrows (X, P^*X)$ and $(!_Z, u) : (Z, P^*Z) \rightleftarrows (1, 1)$.

Then the remaining parametric lens has type $(\Xi \times \Omega, P^*\Xi \times P^*\Omega, \mathcal{A}) : (1, 1) \rightleftarrows (1, 1)$, which one can easily prove being equivalent to a function $\Xi \times \Omega \rightarrow P^*\Xi \times P^*\Omega$. Following \bar{x} and u around the arena, one can see what this function is given by

$$\begin{aligned} (\bar{\xi}, \bar{\omega}) \mapsto \langle u_{\Xi}, u_{\Omega} \rangle \text{ where } u_{\Xi} &= \lambda \xi . u(\text{play}_{\mathcal{H}}(\xi, \text{play}_{\mathcal{G}}(\bar{x}, \bar{\omega}))) \\ u_{\Omega} &= \lambda \omega . u(\text{play}_{\mathcal{H}}(\bar{\xi}, \text{play}_{\mathcal{G}}(\bar{x}, \omega))) \end{aligned} \quad (2.9)$$

These two functions are thus giving, to each player, all the information needed to compute their optimal strategies *given the other player's strategy*. **Para**(P^*) makes these payoff functions emerge automatically from the information flow of lenses and from the careful use of Nashators.

Payoff costates. As we've seen in the latter example, an arena needs, eventually, to be closed by a context. The data of an initial state is not particularly interesting, but we need to spend a few words on the construction of *payoff costates*. Until now, open games shared the definition of payoff function with traditional strategic games: a payoff costate $(Z, P) \rightleftarrows (1, 1)$ encodes exactly the information of a payoff function $Z \rightarrow P$. Now, however, payoff functions have been internalized so that a costate needs not to emit the payoff corresponding to a given outcome of the game, but the entire payoff function.

The most direct way to do so is to have a payoff function $u : Z \rightarrow P$ being promoted to a costate $\text{const } u : (Z, P^*Z) \rightleftarrows (1, 1)$ in **Lens**(**Set**) by

$$\text{const } u = P^*u \circ (!_P, \text{const id}) \quad (2.10)$$

where $\text{const id} : P \rightarrow P^P$ is the constant map picking the identity of P . This costate effectively ignores the outcome of the game, and returns u regardless. Alternatively, if P has the structure of a group, we can keep the information about the outcome and define

$$\Delta u = P^*u \circ (!_P, \text{curr}(-)) \quad (2.11)$$

where $\text{curr}(-) : P \rightarrow P^P$ is the curried subtraction of P . This effectively composes to the costate corresponding to the function

$$\begin{aligned} \Delta u : Z &\longrightarrow P^*Z \\ \bar{z} &\longmapsto \lambda z . (u(z) - u(\bar{z})). \end{aligned} \quad (2.12)$$

which is a sort of ‘discrete differential’ of u . Eventually this would get to players as a continuation describing their possible *increment* in payoff as a function of their deviation. In traditional game theory Δu is known as *regret* [16, §3.2]. We believe it to be more conceptually convincing than the constant costate, especially as we compare games with other cybernetic systems in [Section 3](#).

2.3 Adding players

Once an arena is built, we can add players in it. At this stage, we only deal with the ‘vertical’ part of a game, i.e. we draw *above* the arena (which constitutes the ‘horizontal’ part of a game). Here’s where we specify how players team up, what they observe about payoffs and, most importantly, how players process this feedback to update their strategies.

The first thing to notice is that, since $\mathbf{Para}(P^*)$ is not strongly functorial, lifting the whole play function to an arena in one go versus lifting it piece by piece makes a difference in how players end up being implicitly separated.

In fact, if $\text{play}_{\mathcal{G}} : X \rightarrow Y$ and $\text{play}_{\mathcal{H}} : Y \rightarrow Z$ are parameterised by Ω and Ξ , respectively, then $\mathbf{Para}(P^*)(\text{play}_{\mathcal{G}} \circ \text{play}_{\mathcal{H}})$ is parameterised by $(\Xi \times \Omega, P^*(\Xi \times \Omega))$ whereas $\mathbf{Para}(P^*)(\text{play}_{\mathcal{G}}) \circ \mathbf{Para}(P^*)(\text{play}_{\mathcal{H}})$ is parameterised by $(\Xi \times \Omega, P^*\Omega \times P^*\Xi)$. Effectively, $\mathbf{Para}(P^*)(\text{play}_{\mathcal{G}} \circ \text{play}_{\mathcal{H}})$ represents a game featuring a coalition of two players with strategy space $\Xi \times \Omega$, while $\mathbf{Para}(P^*)(\text{play}_{\mathcal{G}}) \circ \mathbf{Para}(P^*)(\text{play}_{\mathcal{H}})$ represents a game with two competing players, with strategy spaces, respectively, Ω and Ξ .

The difference lies in the way feedback is received by players, and in their possible deviations. In the first case, the two players can evaluate joint deviations since their feedback has type $\Xi \times \Omega \rightarrow P$. In the second case, the two players can only evaluate unilateral deviations, because they receive two feedbacks $\Omega \rightarrow P$ and $\Xi \rightarrow P$ obtained by fixing either player’s strategy. We turn the first to the latter by reparameterising along the Nashator $\mathbf{n}_{\Xi, \Omega} : (\Xi \times \Omega, P^*\Xi \times P^*\Omega) \Rightarrow (\Xi \times \Omega, P^*(\Xi \times \Omega))$. Thus, when used as a reparameterisation, *the Nashator breaks down coalitions of players*.

Example 2.2 (Sequential game). Suppose we extend [Example 2.1](#) with another move by the first player (decided by the same strategy space Ω , hence the copy in [Figure 3](#). Contrary to the previous case, if we lifted the three play functions separately and then composed, we would have ended up splitting player one into two players: the long-range correlation between the first and third stage of the game forces us to lift the arena monolithically, as depicted in [Figure 3](#).

We then reparameterise along Δ_{Ω} to clone the strategies of the first player into the third stage, and only then use $\mathbf{n}_{\Omega, \Xi}$ to make sure players are split into two different coalitions.




Figure 3

Remark 2.3. Observe coalitions can always be broken canonically, but there's no canonical way to form them. This is actually good, since creating coalitions requires non-canonical agreements on how to distribute payoffs among its members (so-called *imputations* [16, Chapter 8]).

Finally, the last bit of the game specification concerns the process each player uses to turn the feedback they receive, which tells them how their strategic decisions gets rewarded in the arena, into *strategic deviations*.

Usually, payoffs are numerical and players seek to maximize them. A bit more generally, players have some preferences encoded by a selection function $\varepsilon : P^*\Omega \rightarrow \mathcal{P}\Omega$. We warn the reader that $P^*\Omega = P^\Omega$ is the set of P -valued function to Ω , while $\mathcal{P}\Omega$ is the powerset of Ω .

A selection function fits very well in the setting we devised so far, since it has (almost) the type of the backward part of a lens $\text{sel} : (\Omega, \mathcal{P}\Omega) \rightleftarrows (\Omega, P^*\Omega)$. We thus call such a lens a *selection lens*.

Remark 2.4. Notice the object $(\Omega, \mathcal{P}\Omega)$ can be considered the 'state boundary' for the player

system, in the sense of [18], and betrays an implicit non-determinism in the game system. In fact, we can generalize away from sets by replacing the powerset monad $\mathcal{P} : \mathbf{Set} \rightarrow \mathbf{Set}$ with other (commutative) monads, like the Giry monad on measurable spaces (yielding stochastic games) or the tangent space monad on smooth spaces (yielding differential games).

Remark 2.5. The backward part of a selection lens is actually of the form $\text{sel} : \Omega \times P^*\Omega \rightarrow \mathcal{P}\Omega$, hence a *parametric selection function*. This suggests that Ω is even more than a set of strategies, it represents the *epistemic type* of a player in the sense of Harsanyi [12], that is, an element $\omega \in \Omega$ encodes not only the way a player plans to play but also their preferences (for instance, their aversion to risk). Harsanyi's games of incomplete information, at the moment codified in the framework of open games in [1], can potentially benefit a lot from the new ideas we introduced here.

2.4 Games as systems


Let's wrap up the construction we sketched so far. The first step to specify a game is to fix the players involved (N) and their payoff type P . The arena is built canonically from a play function $\text{play}_{\mathcal{G}} : \Omega \times X \rightarrow Y$, where $\Omega = \Omega_1 \times \cdots \times \Omega_N$ is the product of a strategy space per player, X is a type of initial states and Y a type of possible final outcomes of the game. Given this, we apply $\mathbf{Para}(P^*)$ to $\text{play}_{\mathcal{G}}$, and get back a parametric lens $(\Omega, P^*\Omega, \mathcal{A}) : (X, P^*X) \rightleftarrows (Y, P^*Y)$, the arena.

Remark 2.6. One might object that an initial state $\bar{x} \in X$ and a utility function $\text{const } u$ (or Δu) deserve to be part of the arena too, but experience tells this data is something to provide only when we want to move on to the analysis of the game, since closing an arena prematurely hinders further composition. The difference between a closed and an open arena is remindful of the subtle difference between a normal (resp. extensive) form and a normal (resp. extensive) form game: the latter is the data of the first plus a utility function.

Once the game arena has been built, we assemble the system of players over it. Usually, such a lens will be of the form $(\bigotimes_{i=1}^N \text{sel}_i) \natural_{\Omega_1, \dots, \Omega_N}$, where $\text{sel}_i : (\Omega_i, \mathcal{P}\Omega_i) \rightleftarrows (\Omega_i, P^*\Omega_i)$ are N selection lenses. Notice such a lens has domain $(\Omega, \mathcal{P}\Omega_1 \times \cdots \times \mathcal{P}\Omega_N)$, so we precompose it with $\prod_{i=1}^N (-) : (\Omega, \mathcal{P}\Omega) \rightleftarrows (\Omega, \mathcal{P}\Omega_1 \times \cdots \times \mathcal{P}\Omega_N)$, which is the identity on the forward part and cartesian product⁴ in the backward part.

We denote the resulting lens $(\Omega, \mathcal{P}\Omega) \rightleftarrows (\Omega, P^*\Omega)$ as \mathcal{G} , and this constitutes a *diegetic open game* over \mathbf{Set} . Abstractly, we can consider this is a 'system with boundary \mathcal{A} ', and any such system can rightfully be called a game.

⁴Or better, the canonical lax monoidal structure of the powerset endofunctor.



We stress this lens deserves to be called a system since its left (‘top’ in the drawings) boundary has a canonical form: the *deviations* $\mathcal{P}\Omega$ canonically associated to the given strategy profiles Ω (what Myers calls *changes* in [18]).

Nash equilibria. So far, we never mentioned Nash equilibria. We have claimed that the way we have woven together the various pieces of a game reproduces, diegetically and dynamically, the counterfactual analysis players do in a non-cooperative strategic game.

To see why our claim holds, let’s analyze a game system constructed from a normal form (N, Ω) , following the above recipe. Since normal forms dispense completely with dynamical information, the associated arena will be trivial: we set $X = 1, Y = \Omega$ and $\text{play}_{(N, \Omega)} := \pi_{\Omega} : \Omega \times 1 \rightarrow \Omega$.

What interests us is the system of players \mathcal{G} over $\mathcal{A}_{(N, \Omega)} = \mathbf{Para}(\mathbb{R}^{N^*})(\text{play}_{(N, \Omega)})$. In a traditional non-cooperative game, players simply maximize their payoffs, so that player i will adopt the ‘non-parametric’ selection $\text{sel}_i = (1_{\Omega_i}, \pi_{P^* \Omega_i} \circ \text{argmax}(- \circ \pi_i))$. Defining the rest of \mathcal{G} as above, we get the translation of (N, Ω) to the parametric lens $(\Omega, \mathcal{P}\Omega, \mathcal{G}^* \mathcal{A}_{(N, \Omega)}) : (1, 1) \rightleftharpoons (\Omega, \mathbb{R}^N)$.

One can easily check that in this case \mathcal{G} turns out to reproduce the Nash product \boxtimes of the involved selection functions defined in [4].

$$\mathcal{G} = \prod_{i=1}^N (-) \circ \left(\bigotimes_{i=1}^N \text{sel}_i \right) \circ \mathbf{n}_{\Omega_1, \dots, \Omega_N} = (1_{\Omega}, \pi_{P^* \Omega_i} \circ \boxtimes_{i=1}^N \pi_{P^* \Omega_i} \circ \text{argmax}(- \circ \pi_i)). \quad (2.13)$$

Of course this still works if instead of argmax we had chosen selection lenses sel_i that don’t use their residual, and arguably we can extend the definition of \boxtimes to encompass the remaining cases.

The point, however, is that this observation links the new framework with that of open games with agency: the act of assigning a selection function to an arena has been turned into a *bona fide* specification of a system over that arena, thereby fixing the conceptual picture.

We can now prove the following:

Theorem 2.7. *Let $(N, \Omega = \Omega_1 \times \dots \times \Omega_N, u : \Omega \rightarrow \mathbb{R}^N)$ be an N -players, strategic game in normal form [16, Definition 1.2.1]. Let $\mathcal{G}^*(\mathcal{A}_{(N, \Omega)} \circ \text{const } u)$ be its translation to a diegetic open game, as*

described above, where $\text{const } u$ has been defined in (2.10). Let $\mathcal{G}_{(\Omega, u)} : \Omega \rightarrow \mathcal{P}\Omega$ be the set-valued function corresponding to such a closed parametric lens. Then a strategy profile $\omega \in \Omega$ is a Nash equilibrium for (Ω, u) if and only if $\omega \in \mathcal{G}_{(\Omega, u)}(\omega)$.

Proof. Running ω through $\mathcal{G}_{(\Omega, u)}$ yields $(\boxtimes_{i=1}^N \text{argmax}^{\mathbb{R}}(-; \pi_i))(u)$, which in [4, Theorem 1] is proven to coincide with the set of strategy profiles in Nash equilibrium. \square

In forthcoming work, we describe a principled, general framework to extract Nash equilibria as ‘behaviours’ of the system \mathcal{G} over the arena \mathcal{A} , in the style of [17, 18]. Specifically, we show that Nash equilibria coincide, unsurprisingly, with non-deterministic fixpoints of such systems, i.e. simulations of trivial game. Most importantly, from such a characterization we can automatically deduce compositions formulae that reproduce the key strength of open games, that is compositionality of the equilibria computation. The form of these new formulae seem interesting for the practical implementation of such computations.

3 Diegetic feedback as backpropagation

The conceptual story behind the diegetic representation of feedback in games is not at all specific to them. On the contrary, it opens a window on a broader conceptual story linking the categorical description of cybernetic systems featuring a ‘backpropagation-like’ feedback dynamics (which is most of them, notable exception being open servers [23]). Here we outline how gradient-based learners [6] share the same abstract features, in a striking example of category theory enabling a rigorous description of a previously only informal analogy.

In gradient-based learning, a smooth function $X \rightarrow Y$ is learned by optimizing a model $f : \Omega \times X \rightarrow Y$ smoothly parameterised by the variable $\omega \in \Omega$. Conceptually, this is only possible because differential structure leaks information about the loss $\ell : Y \times Y \rightarrow \mathbb{R}$ ‘in a neighbourhood’ of $(y, f(\omega, x))$, and this can be used to evaluate which changes in parameter the learner should implement to improve. Hence it is paramount that ℓ is known ‘locally’, and not just pointwise. In practice, the value of ℓ at $(y, f(\omega, x))$ is not even used! Only the covector $d_{f(\omega, x)}\ell(y, -)$ is needed.

This covector is then backpropagated across the various components of the learner until a covector on Ω is obtained. As for games, this backpropagation mechanism is automatically enabled by deploying the functor

$$T^* : \mathbf{Smooth} \longrightarrow \mathbf{DLens}(\mathbf{Vec}_{\mathbb{R}}) \quad (3.1)$$

sending each manifold X to its *cotangent vector bundle* $T^*X \rightarrow X$ (the fiberwise dual of its tangent bundle) and each map $f : X \rightarrow Y$ to its reverse derivative, i.e. pullback of covectors along f [22], naturally expressed as a dependent lens $(f, T^*f) : (X, T^*X) \rightleftarrows (Y, T^*Y)$.⁵

Remark 3.1. In [6], a functor very similar to T^* is obtained from the structure of *reverse differential category* (RDC) on the base category, but **Smooth** is not such a category. Therefore, in *ibid.* the authors confine themselves to its wide subcategory **Euc** of Euclidean spaces. In light of our findings for games, it seems that considering functors $\mathcal{S} \rightarrow \mathbf{DLens}(\mathcal{S})$ splitting the view fibration to be more fundamental than reverse differential structure in the sense of [5]. Already in [5, §4] and [6, Proposition 2.12], it is shown how reverse differential structures can be encoded as sections of the view fibration of lenses, with extra conditions account for the ‘additivity’ necessary in the framework of RDCs. It seems reasonable, therefore, to reformulate RDCs as particularly nice instances of *section of feedbacks*, dualizing that of *section of changes* defined by Myers in [17, 18].

Remark 3.2. The functor T^* is strong monoidal and thus is associated to a pseudofunctor $\mathbf{Para}(T^*)$ that promotes a smooth parametric function straight into a backpropagating model. Compare this with the functor $\mathbf{Para}(P^*)$, whose laxity is, ultimately, the source of the many interesting phenomena in non-cooperative strategic games. The fact T^* is *not* lax is attributable to the additive structure involved in each fiber of a cotangent bundle, whereby $T^*(X \times Y) \cong T^*(X + Y)$.

Once an arena $\mathcal{L} := \mathbf{Para}(T^*)(\Omega, f)$ has been defined, the dynamic of an agent (which is what really deserves the name of ‘learner’) actually doing the learning is given by a *gradient flow* lens $\text{GF} : (\Omega, T\Omega) \rightleftarrows (\Omega, T^*\Omega)$ which defines a system over \mathcal{L} , by reparameterisation (as in **Figure 1b**). The backward part of such a lens is determined by the data of a fiberwise linear morphism $(-)^{\sharp} : T^*\Omega \rightarrow T\Omega$. The most common way such a morphism arises is when Ω is endowed with a Riemannian metric g , in which case $(-)^{\sharp}$ (known as ‘raising indices’ [22]) selects the direction of steepest ascent associated to a covector, so that u^{\sharp} is $\text{argmax}_{v \in T_{\omega}\Omega} u(v) / \|v\|_g$ for a given $u \in T_{\omega}^*\Omega$.

As highlighted in **Table 1**, $(-)^{\sharp}$ is formally analogous to a selection function $\text{sel} : \Omega \times P^*\Omega \rightarrow \mathcal{P}\Omega$, which indeed has the same role for games. This is corroborated by the type signatures of GF and sel, both going from an object of ‘states and feedbacks’ to an object of ‘states and changes’.

What might look odd is the asymmetry between $\mathcal{P}\Omega$ and Ω in the signature of sel, something not present in $(-)^{\sharp}$. Indeed, if Ω is the set of ‘states’ of a player, then there is a dissimilarity

⁵Specifically, the codomain of T^* is the category of dependent lenses [21] obtained from the indexed category of smooth \mathbb{R} -vector bundles $\mathbf{Vec}_{\mathbb{R}} : \mathbf{Smooth}^{\text{op}} \rightarrow \mathbf{Cat}$.

games	gradient-based learners
strategies Ω	parameters Ω
deviations $\mathcal{P}\Omega$	vectors $T\Omega$
payoff functions $P^*\Omega := P^\Omega$	covectors $T^*\Omega$
precomposition $P^*f : X \times P^*Y \rightarrow P^*X$	reverse derivative $T^*f : f^*(T^*Y) \rightarrow T^*X$
selection function $\text{sel} : \Omega \times P^*\Omega \rightarrow \mathcal{P}\Omega$	sharp (iso)morphism $(-)^{\sharp} : T^*\Omega \rightarrow T\Omega$ (of vector bundles over Ω)

Table 1

between $T^*\Omega$ being the set of \mathbb{R} -valuations of $T\Omega$ and $P^*\Omega$ being the set of valuations on $\mathcal{P}\Omega$. This discrepancy requires a bit more scaffolding to be explained, but intuitively it amounts to observing $T^*\Omega$ is the set of *linear* valuations on $T\Omega$, and likewise, when we consider only maps $f : \mathcal{P}\Omega \rightarrow P$ that satisfy $f(A) = \sum_{a \in A} f(\{a\})$, these are determined by maps $\Omega \rightarrow P$.

Table 1 could be expanded with a third column for Bayesian reasoners, though it's still not completely clear how that would exactly work. In his forthcoming work [2], Braithwaite puts forward an approach to Bayesian lenses for general Markov categories, which gives one way to fill the missing column. On the other hand, at least for classical probability spaces, information geometry seems to suggest Bayesian reasoners are gradient-based learners living on (possibly singular) manifolds of probability measures, moving by pure gradient flow [11].

4 Conclusions

In this work we described a new approach to the specification of compositional games in the style of open games [10, 4]. It corrects some of the conceptual shortcomings of open games with agency, and uncovers deeper analogies with gradient-based learners and, speculatively, a wider range of cybernetic systems.

The new approach provides a way to specify a game using machinery analogous to reverse-mode

automatic differentiation, abstractly given by a functor $P^* : \mathbf{Set} \rightarrow \mathbf{DLens}(\mathbf{Set})$. We observed how the lax monoidal structure of such functor plays a profound role in determining the dynamics of non-cooperative games, by hiding ‘cooperative’ information.

We have shown how classical strategic games can be naturally represented as non-deterministic systems over their arenas, systems given by the dynamics of players observing their payoffs and pondering if and how to deviate from their current strategy. The resulting parametric lens is hence a full realization of the ideas in [14, 4, 3], and brings the framework of categorical cybernetics (born with [3]) closer to that of categorical systems theory (detailed in [17, 18]).

Future directions. The new ideas brought about in this paper are not fully formed yet. In preparing this work, three more follow-up works naturally spawned.

The first, which has already been anticipated at the end of [Section 2](#), concerns laying down a proper general theory of specification and simulation of cybernetic systems, in the wake of Myers’ work on dynamical systems [17, 18]. In the first place, this would allow to extract Nash equilibria from diegetic open games in a principled and compositional way, with practical implications in the way these are computed. Secondly, using analogous tools we would then be able to talk about simulations of games and more generally of non-equilibrium trajectories of game dynamics.

The second work concerns the pure game-theoretic aspects of this new definition. Can we improve the toolset of compositional game theory by leveraging a more accurate reproduction of the dynamics involved? We believe the answer to be yes, with exciting connections to the topic of Bayesian games [12], learning theory for games [9], coarsening and ossification, *etc.*

The third work will be an exploration of the ideas roughly outlined in [Section 3](#), with the aim of clarifying the analogy uniting learners, games, Bayesian reasoners and other examples of ‘backpropagating cybernetic systems’. This would formalize the intuitive picture whereby such systems come with a notion of ‘state space’, ‘bundle of changes’, ‘bundle of scalars’ and a ‘bundle of feedbacks’ obtained as valuations of the first in the latter. A tantalizing consequence is to get past the conceptually awkward act of *fixing* the payoff object P , which doesn’t allow for the number of players to be dynamically determined, by attaching payoff types to states like ‘structure sheaves’. This recovers traditional open games [10] as the description of payoff types and their transformation across the game, while a functor analogous to T^* turns that static picture into a dynamical one.

References

- [1] Joe Bolt, Jules Hedges & Philipp Zahn (2019): *Bayesian open games*. Available at <https://arxiv.org/abs/1910.03656>.
- [2] Dylan Braithwaite (2022): *Dependent Bayesian Lenses: Categories of Bidirectional Markov Kernels with Canonical Bayesian Inversion*. In preparation.
- [3] Matteo Capucci, Bruno Gavranović, Jules Hedges & Eigil Fjeldgren Rischel (2021): *Towards foundations of categorical cybernetics*. In: *Proceedings of Applied Category Theory 2021*, EPTCS. Available at <https://arxiv.org/abs/2105.06332>.
- [4] Matteo Capucci, Neil Ghani, Jérémy Ledent & Fredrik Nordvall Forsberg (2021): *Translating extensive form games to open games with agency*. In: *Proceedings of Applied Category Theory 2021*, EPTCS. Available at <https://arxiv.org/abs/2105.06763>.
- [5] Robin Cockett, Geoffrey Cruttwell, Jonathan Gallagher, Jean-Simon Pacaud Lemay, Benjamin MacAdam, Gordon Plotkin & Dorette Pronk (2020): *Reverse Derivative Categories*. In: *28th EACSL Annual Conference on Computer Science Logic (CSL 2020)*, Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
- [6] Geoff S. H. Cruttwell, Bruno Gavranovic, Neil Ghani, Paul W. Wilson & Fabio Zanasi (2021): *Categorical Foundations of Gradient-Based Learning*. CoRR abs/2103.01931. Available at <https://arxiv.org/abs/2103.01931>.
- [7] David Dalrymple (2019): *Dioptics: a Common Generalization of Open Games and Gradient-Based Learners*. In: *Fifth Symposium on Compositional Structures (SYCO 5)*. Available at <https://www.cl.cam.ac.uk/events/syco/strings3-syco5/papers/dalrymple.pdf>.
- [8] Merriam-Webster.com Dictionary: *Diegetic*. Available at <https://www.merriam-webster.com/dictionary/diegetic>.
- [9] Drew Fudenberg, Fudenberg Drew, David K Levine & David K Levine (1998): *The theory of learning in games*. 2, MIT press.
- [10] Neil Ghani, Jules Hedges, Viktor Winschel & Philipp Zahn (2016): *Compositional game theory*. arXiv e-prints.
- [11] Marc Harper (2009): *The replicator equation as an inference dynamic*. Available at <https://arxiv.org/abs/0911.1763>.
- [12] John C Harsanyi (1967): *Games with incomplete information played by “Bayesian” players, I–III Part I. The basic model*. *Management science* 14(3), pp. 159–182.

- [13] Jules Hedges (2016): *Towards compositional game theory*. Ph.D. thesis, Queen Mary University of London.
- [14] Jules Hedges, Paulo Oliva, Evguenia Shprits, Viktor Winschel & Philipp Zahn (2017): *Higher-order decision theory*. In: *International Conference on Algorithmic Decision Theory*, Springer, pp. 241–254.
- [15] Niles Johnson & Donald Yau (2021): *2-dimensional Categories*. Oxford University Press, USA.
- [16] Kevin Leyton-Brown & Yoav Shoham (2008): *Essentials of Game Theory: A Concise Multidisciplinary Introduction*. Morgan & Claypool.
- [17] David Jaz Myers (2020): *Double categories of open dynamical systems*. Available at <https://arxiv.org/abs/2005.05956>.
- [18] David Jaz Myers (2022): *Categorical Systems Theory*. Available at <http://davidjaz.com/Papers/DynamicalBook.pdf>. (Work in progress).
- [19] Michael Shulman (2021): *Homotopy type theory: The logic of space*. In: *New Spaces in Mathematics: Volume 1: Formal and Conceptual Reflections*, Cambridge University Press.
- [20] Toby St Clere Smithe (2021): *Compositional Active Inference I: Bayesian Lenses*. *Statistical Games*. Available at <https://arxiv.org/abs/2109.04461>.
- [21] David I Spivak (2019): *Generalized Lens Categories via functors $F : \mathcal{C}^{\text{op}} \rightarrow \text{Cat}$* . Available at <https://arxiv.org/abs/1908.02202>.
- [22] Michael Spivak (1973): *A comprehensive introduction to differential geometry*. *Bulletins of the American Mathematical Society* 79, pp. 303–306.
- [23] Andre Videla & Matteo Capucci (2022): *Lenses for Composable Servers*. Available at <https://arxiv.org/abs/2203.15633>.