A Theory for Game Theories

Michel Hirschowitz ⋆, André Hirschowitz ⋆⋆, and Tom Hirschowitz ⋆⋆⋆

Abstract. Game semantics is a valuable source of fully abstract models of programming languages or proof theories based on categories of so-called games and strategies. However, there are many variants of this technique, whose interrelationships largely remain to be elucidated. This raises the question: what is a category of games and strategies?
Our central idea, taken from the first author's PhD thesis [11], is that positions and moves in a game should be morphisms in a base category: playing move m in position f consists in factoring f through m, the new position being the other factor. Accordingly, we provide a general construction which, from a selection of *legal moves* in an almost arbitrary category, produces a category of games and strategies, together with subcategories of *deterministic* and *winning* strategies.
As our running example, we instantiate our construction to obtain the standard category of Hyland-Ong games subject to the switching condition. The extension of our framework to games without the switching condition is handled in the first author's PhD thesis [11].

Keywords: Game semantics, categories.

1 Introduction

1.1 The flavor problem

Game semantics appeared in the early 90's [3, 12] and provided convenient denotational semantics to proof theories and programming languages, including their non functional features [2, 5, 4, 13, 8, 14]. However, game semantics has roughly as many variants as it has authors. Each of these game theories starts from a notion of "arrow" game (with corresponding positions and moves), yielding the natural notion of strategy. The crucial construction is then the composition of strategies, with the crucial feature that various meaningful classes of strategies (deterministic, innocent, winning) are preserved by composition.

All these compositions clearly have a common flavor (sometimes called "compose+hide"). In the present work, we propose an explanation for this common flavor. To this effect, we define, through a single construction, a huge class of game theories where the composition of strategies preserves good properties. This class contains those among existing game theories which respect the

⋆ CEA-LIST **michel.hirschowitz@cea.fr**
⋆⋆ UMR 6621 CNRS-Université de Nice-Sophia-Antipolis **ah@math.unice.fr**
⋆⋆⋆ UMR 5668 CNRS-ENS Lyon-INRIA-UCBL **tom.hirschowitz@ens-lyon.fr**

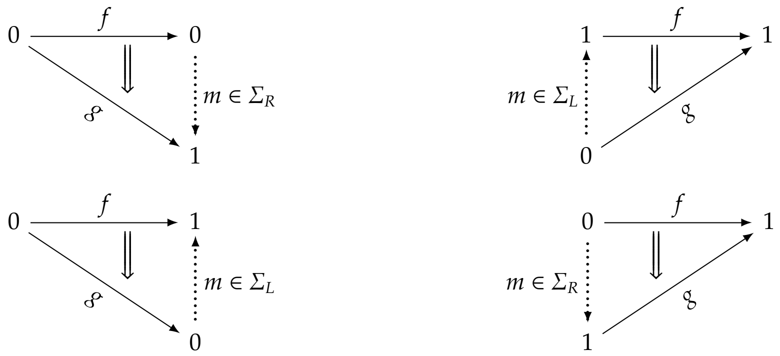

Fig. 1. The four kinds of edges in W_C, from f to g

so-called switching condition [7]. This restriction is only due to the fact that we have chosen to present the simplest version of the construction. Indeed, the more general version [11] involves a serious amount of weak categorical material. Nevertheless, future game models relying on our framework will avoid the burden of re-proving the combinatorial lemmas leading to the category of games and strategies. We now proceed to give a more detailed overview.

1.2 Playing in a one-way category

In our approach, a play may take place in any *one-way* category, which we define to be a category where objects have a sign (1/0) and where morphisms cannot go from a 1-object to a 0-one. Equivalently, a one-way category is a category C, equipped with a functor $\lambda : C \to \mathbb{2}$, where $\mathbb{2}$ is the ordered set $0 \leq 1$.

The crucial part of our construction builds a *wild* game W_C from a one-way category C. This game is wild in the sense that the two players play without any restriction (meaningful restrictions will be considered later). Let us sketch the construction of W_C. It is a directed graph, whose vertices are the morphisms of C. Thus we have one kind (01) of *odd* vertices and two kinds (00, 11) of *even* vertices. We think of these "states" as follows: at an odd vertex, Player has to play and reach an even vertex; at a 11-vertex, Opponent has to play "on the left-hand side" (and reach an odd vertex) , while at a 00-vertex Opponent has to play "on the right-hand side" (and reach an odd vertex). This yields the following diagram of states

$$11 \underset{\mathsf{L}}{\overset{\mathsf{M}L}{\rightleftarrows}} 01 \underset{\mathsf{R}}{\overset{\mathsf{M}R}{\rightleftarrows}} 00.$$

In other words we have four kinds of edges ($\mathsf{M}L$ and $\mathsf{M}R$ for Player's moves, L and R for Opponent's) which we now describe in more detail. The rule is that only one end of the vertex (a morphism in C) changes, and the slogan says that

O composes while P decomposes, as pictured in Figure 1: an edge from f to g, consists of an odd morphism m respectively satisfying the following rule:

Kind of move	R	L	MR	ML
Rule	$g = m \circ f$	$g = f \circ m$	$f = m \circ g$	$f = g \circ m.$

Because each move changes the signs, all the m's above and in Figure 1 have sign $0 \to 1$.

The wild game we have constructed so far offers essentially the complete picture which we want to show, in particular one may define strategies and their composition. On the other hand, as far as meaning is concerned, the wild game is trivial, in the sense that players can easily neutralize each other. Indeed, for instance, assume Opponent moves from the current position f to, say, $m \circ f$ by composing with m. Then, Player may move back to f by decomposing $m \circ f$ into m and f. Thus, in the wild game, all moves are undoable. More meaningful and sufficiently general games are obtained as subgames of the wild game simply by restricting the set of odd morphisms allowed in the process of composition/decomposition. For this reason we define a *game setting* to be a one-way category equipped with two sets of odd morphisms as explained in Section 2.2. In the rest of the paper, we explain the basic theory of plays and strategies in a game setting, and we show how the theory of HO games may be recovered in terms of a game setting.

Related work Cockett and Seely [6] offer another categorical investigation into game semantics. The relationship between their work and ours remains unclear to us. Let us also mention a recent paper [10] which describes a categorical reconstruction of ”pointer” games and *innocent* strategies from ”general” games and strategies. In this sense, they reduce one sophisticated (but efficient) category of games to a much simpler one. Thus they aim at a better understanding of one (very important) category of games, and of the concept of innocence, while we aim at a better understanding of what could be a category of games, and do not consider the concept of innocence.

Organization of the paper In Section 2, we provide the categorical construction which, from a so-called *game setting*, constructs a double category of plays, where vertical composition is sequential composition, while horizontal composition is reminiscent of the usual composition of strategies. We then instantiate our framework in Section 3: after recalling the basics of (a standard variant of) HO games, we exhibit a game setting hidden in it, for which our construction yields the usual notion of plays and arrow arenas. In Section 4, we describe strategies in our abstract framework, as well as their composition, and we show that the obtained notion of HO strategies closely corresponds to the standard one.

2 The abstract framework: building the double category

2.1 Game settings

In order to restrict moves in the game sketched above, we should a priori specify four sets M_{OL}, M_{OR}, M_{PL}, M_{PR} of legal odd morphisms, one for each of the four kinds of moves in Figure 1. However, these restrictions will be compatible with the composition of strategies only if we impose $M_{OL} = M_{PR}$ and $M_{OR} = M_{PL}$. This leads to our definition: a *game setting* $G \triangleq (C, \Sigma_R, \Sigma_L)$ consists of a one-way category C equipped with a pair of sets of odd morphisms: Σ_R is the set of *forward* moves (or *f-moves*; those going downwards in Figure 1); Σ_L is the set of *backward* moves or *b-moves*. The wild game (on C) is obtained by taking as Σ_R and Σ_L the whole set of odd morphisms.

In a game setting G, we view objects as positions in a two-player game, actually a signed graph. Morphisms in Σ_R and Σ_L are Opponent and Player moves, respectively. On 0-labeled objects, Opponent is to play, whilst on 1-labeled ones, it's Player's turn. As illustrated in Figure 2, from some 0-labeled position p, Opponent plays by choosing an f-move $m : p \to q$ with domain p, thereby reaching the 1-labeled position q. Conversely, from such a q, Player plays by choosing a b-move $m' : r \to q$ with codomain q, thereby reaching the position r. This defines a graph whose vertices are the objects of C, which we call the 0-*dimensional game* (0-game for short) of G and denote by $\mathcal{G}_0(G)$. We call the free category over this graph the category of 0-*plays* over G, and denote it by $C_0(G)$.

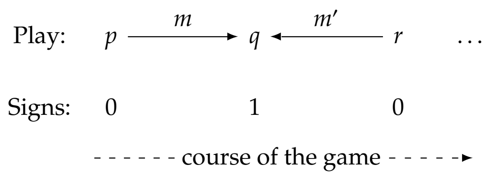

Fig. 2. Example play in the 0-game

Each 0-play v has a predecessor Pred(v) obtained by deleting the last move (if any) .

2.2 The 1-dimensional game

As in standard game semantics, this yields a natural notion of arrow game, also a graph, which we call the 1-*dimensional game* (1-game for short) of G and denote by $\mathcal{G}_1(G)$. We describe the positions of this game first, then its moves, and finally we show how to equip it with signs, in a way that refines the above interpretation of signs in the 0-game. Positions (or vertices) in $\mathcal{G}_1(G)$ are morphisms in C. Given the constraints on signs, there are just three kinds of positions: 00, 01, and 11. Then, moves from f to g in the 1-game are defined to be commutative triangles in C, of one of the four shapes in Figure 1.

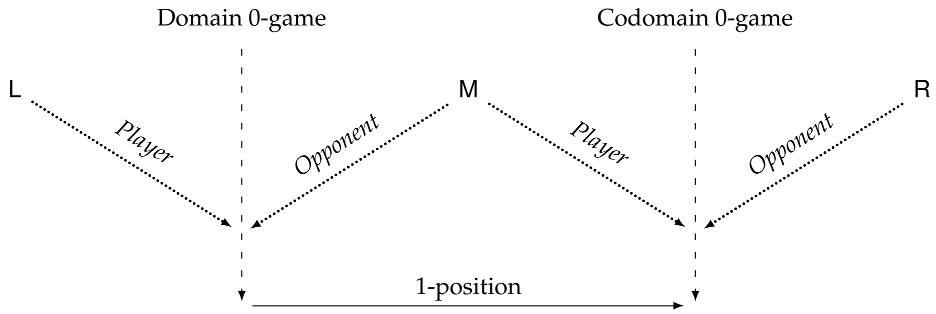

Fig. 3. All agents (L, M, R) act as Player on their rhs and as Opponent on their lhs

The interpretation of signs in 1-games, illustrated in Figure 3, entirely follows from the idea that in 0-games, Player lives on the left-hand side of the position, whilst Opponent lives on its right-hand side. For a 1-position, there is thus one agent M in the middle, and one agent on each side, which we call L and R in the obvious way. M plays Opponent in the domain 0-game, and Player in the codomain 0-game. L plays Player in the domain 0-game, whilst R plays Opponent in the codomain 0-game. This yields the following rule for the 1-game:

Signs of the 1-position	Who's to play?
$0 \longrightarrow 0$	R
$0 \longrightarrow 1$	M
$1 \longrightarrow 1$	L .

We consider the free category over this graph $\mathcal{G}_1(G)$: we call it the category of 1-*plays* over G and denote it by $C_1(G)$. Again each 1-play v has a predecessor $\mathrm{Pred}(v)$ obtained by deleting the last move (if any).

Finally, we define the (horizontal) source and target functors on 1-plays, $\mathsf{s}, \mathsf{t} : C_1(G) \to C_0(G)$, by the obvious induction (or adjunction). We thus have a pullback category $C_1(G)\ {}_{\mathsf{s}}\!\times_{\mathsf{t}} C_1(G))$ of composable pairs of 1-plays.

2.3 The double category associated to a game setting

In this section, we derive a double-categorical structure from our game setting G. For this, we will define a notion of horizontal composition of 1-plays, yielding a category whose objects are 0-plays, and whose morphisms are 1-plays. We start by defining the graph $\mathcal{G}_2(C)$ of *primitive interactions* as follows. As vertices, take composable pairs of morphisms in C, and as edges from the pair $\xrightarrow{f}\ \xrightarrow{g}$

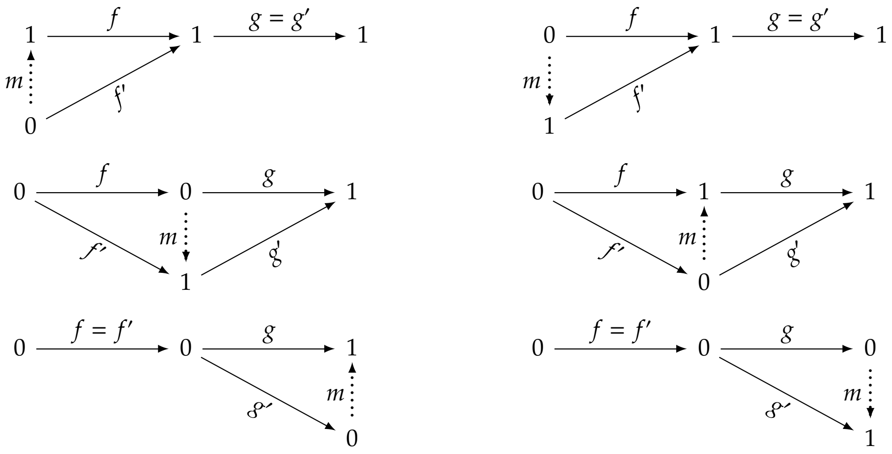

Fig. 4. The six kinds of edges in $\mathcal{G}_2(C)$ (each edge top-down)

to the pair $\xrightarrow{f'} \xrightarrow{g'}$, take all the commutative diagrams as in Figure 4. This gives four kinds of vertices $(000, 001, 011, 111)$ according to the signs of objects, yielding the following state diagram:

$$111 \underset{\mathsf{M}_1 L}{\overset{\mathsf{L}}{\rightleftarrows}} 011 \underset{\mathsf{M}_1 R}{\overset{\mathsf{M}_2 L}{\leftrightarrows}} 001 \underset{\mathsf{R}}{\overset{\mathsf{M}_2 R}{\rightleftarrows}} 000.$$

For $\mathcal{G}_2(C)$, the intuition is that there are two players M_1 and M_2, and two opponents L and R, who interact respectively on the left-hand side with M_1 and on the right-hand side with M_2. Thanks to categorical composition, both players act exactly as if they were facing two opponents. For instance, M_1 interacts with L on the left-hand side, and with M_2 on the right-hand side. Because of sign rules, at most one of M_1 and M_2 may play at a given time, which prevents any conflict to arise.

Next, we let $C_2(G)$ denote the free category generated by $\mathcal{G}_2(G)$, and we call its morphisms *interactions* in G. Accordingly, the edges in $\mathcal{G}_2(G)$ are *primitive* interactions. Let us also deem the primitive interactions of the middle row *internal*, and the other ones *external*. Now a key observation is that the functor

$$C_2(G) \xrightarrow{\langle \pi_1, \pi_2 \rangle} C_1(G) \,{}_{\mathrm{s}}{\times}_{\mathrm{t}}\, C_1(G)$$

which maps a path in $\mathcal{G}_2(G)$ to its left and right borders is an isomorphism, which says altogether that interactions are determined by their projections, and that $C_1(G) \,{}_{\mathrm{s}}{\times}_{\mathrm{t}}\, C_1(G)$ is freely generated by the primitive interactions.

Thanks to this statement, it is enough to define our 1-horizontal composition $Y \bullet X$ on primitive interactions, which is straightforward: for an internal interaction, the 1-horizontal composition is the empty 1-play. Otherwise it is the obvious move from $g \circ f$ to $g' \circ f'$, for each external interaction as in Figure 4.

To construct our horizontal category, we finally define identity morphisms, by mimicking what is standardly called *copycat* in game semantics: let copycat be the unique functor from $\mathcal{G}_0(G)$ to $\mathcal{G}_1(G)$ such that f-moves $m : p \to p'$ and b-moves $m : p' \to p$ are respectively sent to plays

By the standard adjunction between categories and directed graphs, this defines copycat uniquely: on arbitrary plays, copycat simply piles up sequences of such elementary plays.

Proposition 1. *The horizontal composition of* 1*-plays is associative and unital.*

The proof of associativity relies on a freeness result concerning 3-interoctions, completely analogous to our previous freeness result concerning 2-interactions.

This all gives the data for a *double category*. A short definition is as follows: a double category is a category object in the category of categories. A more explicit, elementary definition may be found, e.g., in Melliès [15]. We've already checked all the required properties, except the *interchange law*, which makes $\bullet$ into a functor from the pullback $C_1(G)_{\mathsf{s}}\times_{\mathsf{t}}C_1(G)$ to $C_1(G)$. Explicitly: $(Y_1 \bullet X_1) \circ (Y_2 \bullet X_2) = (Y_2 \circ Y_1) \bullet (X_2 \circ X_1)$. It happens to be satisfied, which entails:

Theorem 1. *For any game setting G, the categories $C_0(G)$ and $C_1(G)$, the domain and codomain functors $\mathsf{s}, \mathsf{t} : C_1(G) \to C_0(G)$, the horizontal composition functor $\bullet : C_1(G) \,_{\mathsf{s}}\times_{\mathsf{t}}\, C_1(G) \to C_1(G)$, and the horizontal identity functor $I : C_0(G) \to C_1(G)$ form a double category.*

3 The one-way category underlying Hyland-Ong games

3.1 A brief review of HO-arenas and HO-plays

We briefly recall some definitions of HO game theory, and refer the reader to Harmer's notes [9] for details.

An *arena* A is a triple $(M_A, \lambda_A, \vdash_A)$, where M_A is a set of *moves*, λ_A gives signs to moves, i.e., is a function from M_A to $\{0, 1\}$, and $\vdash_A$ represents altogether a binary relation (justification) and a predicate (initiality) on M_A, such that:

1. if $\vdash_A m$, then $\lambda_A(m) = 0$ and for all $m' \in M_A$, $m' \nvdash_A m$,
2. if $m \vdash_A m'$, then $\lambda_A(m) \neq \lambda_A(m')$.

Moves m such that $\vdash_A m$ are called *initial*. When $m \vdash_A m'$, we say that m *justifies* m'.

A *position* in an arena A is a pair (s, ρ), where $s = m_1 \dots m_n$ is a sequence of moves of alternate signs in A, and ρ is a function from $\{1 \dots n\}$ to $\{0 \dots n-1\}$ such that for all $i \in \{1 \dots n\}$

1. (priority condition) $\rho(i) < i$,
2. if $\rho(i) = 0$, then m_i is initial,
3. if $\rho(i) = j \neq 0$, then m_j justifies m_i.

We say that n is the length of the position. Our position p also has an *initial part* $Init_p \subset [1, \ldots, n]$ which is the set of indices i for which m_i is initial. Since positions carry their history, they also may be seen as plays, and we freely call them either way. A position of length 0 is called *initial*, and a non initial position p of length n has a predecessor position $\mathrm{Pred}(p)$, of length $n - 1$, obtained by deleting the last move. For simplicity, we define $\mathrm{Pred}(p) \triangleq p$ when p is initial. A set of positions is *prefix-closed* when it is closed under application of Pred.

Given a sign function λ, we write $\overline{\lambda}$ for the opposite one. Given two arenas A and B, one constructs the *arrow arena* $A \multimap B$ by taking $M_A + M_B$ as the set of moves, $[\overline{\lambda_A}, \lambda_B]$ as a sign function, the (injections of) initial moves of B as initial moves, and for the binary $\vdash_{A \multimap B}$, taking the union (up to injection) of $\vdash_A$ and $\vdash_B$, plus the pairs (m, m') with m initial in B and m' initial in A. Note that a position p in an arrow arena $A \multimap B$ determines two projections p_A and p_B which are in general not positions in A and B. Intuitively, this is because Opponent may switch sides, and, when asked a question in A, ask a question in B.

Define a position p in $A \multimap B$ to be *valid* if its two projections are again positions, respectively in A and B. Combinatorially, if n_A and n_B are the lengths of these projections, p determines a *shuffle* $p_S = [1, \ldots, n_A + n_B] \to [1, \ldots, n_B] \amalg [1, \ldots, n_A]$. We say that such a shuffle p_s satisfies the *switching condition*, or is *even* when

- if $n_A + n_B > 0$ then $p_S(1)$ is on the B-side,
- for i satisfying $1 < 2i < n_A + n_B$, $p_S(2i)$ and $p_S(2i + 1)$ are on the same side.

It turns out that p is valid exactly when p_S is even. We note that p determines a restricted justification map $\mathrm{RJ}_p : Init_{p_A} \to Init_{p_B}$. Conversely, given the projections p_A and p_B, a position p is determined by an arbitrary map $\mathrm{RJ} : Init_{p_A} \to Init_{p_B}$ and an even shuffle compatible with RJ (with respect to the priority condition).

Strategies from A to B are defined to be non-empty, prefix-closed sets of valid positions in $A \multimap B$. One then shows that strategies compose and have identities, which yields a category of games and strategies Strat_{HO}.

3.2 The one-way category C_{HO}

Let us now describe the one-way category C_{HO} relevant for HO games. An object $(A, (s, \rho))$ of C_{HO} is merely a position (s, ρ) in a game arena A, while a morphism from $p = (A, (s, \rho))$ to $q = (B, (s', \rho'))$ is a (valid) position $f = (A \multimap B, (t, \tau))$ whose projections respectively give p and q. Thus our morphisms also have predecessors. Note that f and $\mathrm{Pred}(f)$ share one end, but in general not both.

We are especially concerned with two kinds of morphisms. Firstly, for each position $p = (A, (s, \rho))$, we have a *copycat* morphism $\mathrm{copycat}_p : p \to p$, which is defined by induction on the length of p: the empty play on $A \multimap A$ is the copycat of the initial position on A, and for greater lengths, $\mathrm{copycat}_p$ is determined by the requirement that its second predecessor is the copycat of Pred(p): the last two moves are determined by the given projections (p and p). Secondly, we are interested in those morphisms whose predecessor is a copycat, which we call *subcopycat* morphisms. Each subcopycat morphism is also the predecessor of a unique copycat morphism. Thus, for a non initial position p, define Sub_p to be the predecessor of $\mathrm{copycat}_p$. Then, each subcopycat morphism can be written Sub_p in a unique way. Furthermore, if p is even, then Sub_p goes from p to Pred(p) while if p is odd, then Sub_p goes from Pred(p) to p.

Next, we define the composition of our morphisms. Consider two consecutive arrows, i.e., valid positions f in some $A \multimap B$ and g in $B \multimap C$ with the same projection p_B on B. We denote by p_A the projection of f on A, and by p_C the projection of g on C and by n_A, n_B, n_C the corresponding lengths. We will define $h \triangleq g \circ f$ by its restricted justification map RJ_h and its even shuffle h_S. For RJ_h, we take the composition $\mathrm{RJ}_g \circ \mathrm{RJ}_f$. For h_S, we observe that, thanks to the switching condition, there is a unique shuffle $s : [1, \ldots, n_A + n_B + n_C] \to [1, \ldots, n_C] \amalg [1, \ldots, n_B] \amalg [1, \ldots, n_A]$ compatible with f_S and g_S. We view this shuffle as an order on $[1, \ldots, n_C] \amalg [1, \ldots, n_B] \amalg [1, \ldots, n_A]$ and take for h_S its restriction to $[1, \ldots, n_C] \amalg [1, \ldots, n_A]$.

This composition is easily seen to be associative, and it is easily checked that the identity on a position p is the copycat morphism $\mathrm{copycat}_p$.

This altogether gives a category C_{HO}, whose objects may be given a sign as follows: the sign of a position is 0 if Opponent is to play, or equivalently if its length is even, and 1 otherwise. Thus, a priori, we have four kinds of morphisms, $0 \to 0$, $0 \to 1$, $1 \to 0$, $1 \to 1$. However, we easily check that there are no morphisms of type $1 \to 0$. This is a consequence of the switching condition, and the convention that plays always start with a move by Opponent, which furthermore, in the case of arrow arenas, has to be on the right-hand side. Our category may thus be seen as a one-way category.

Remark 1 (Relaxing the switching condition). If we relax the switching condition, and allow Opponent to switch sides in an arrow game, the main new feature is that the horizontal composition of 1-plays is no more well-defined, because interactions are no more determined by their projections. As a consequence, the double category constructed above has to be replaced by some kind of weak double category, to be defined accordingly. This approach has been pursued in the first author's PhD thesis [11], where one eventually recovers a proper category when passing to strategies.

3.3 The game setting G_{HO}

Now we explain how HO-moves may be seen as morphisms in C_{HO}. Playing a move in a position p in A is understood as extending p (with one move in A), yielding a new position q. To this move, we attach the morphism Sub_q. Note that Sub_q goes from p to q if p is even, and from q to p if p is odd. Hence in our view, the set of HO-moves is precisely the set of subcopycat morphisms, which we split into the set RHO of subcopycat morphisms where the length of the codomain exceeds the length of the domain by one, and the set LHO of subcopycat morphisms where the length of the domain exceeds the length of the codomain by one. Thus, standard HO plays are 0-plays starting on an initial position in the game setting $G_{HO} \triangleq (\mathcal{C}_{HO}, \mathrm{LHO}, \mathrm{RHO})$. (In the game setting, we also consider plays starting on non initial positions.)

Now let us see how our view fits with plays in an arrow game: consider a valid position f in the game $A \multimap B$ and its extension to a new valid position g, through a HO-move m (in A or in B). We have four kinds of extensions corresponding to who is playing and where. A careful inspection shows that

- if O plays in B, then we have $g = m \circ f$ (in $\mathcal{C}_{HO}$),
- if O plays in A, then we have $g = f \circ m$,
- if P plays in B, then we have $f = m \circ g$,
- if P plays in A, then we have $f = g \circ m$;

which shows that, indeed, O composes the original position with her move, while P decomposes the original position with her move. Thus, standard HO arrow plays are precisely 1-plays in G_{HO} starting on an initial position.

4 An abstract view of strategies

In this section, we show how some standard results on strategies may be understood abstractly in a game setting $G = (\mathcal{C}, \Sigma_R, \Sigma_L)$. Recall that an object of $\mathcal{C}$ is *even* when its sign is 0 and *odd* otherwise. We say that a 1-position $f : p \to q$ is even when p and q have the same sign, and odd otherwise. We note that f is odd exactly when the middle player M is to play, and even exactly when it's L or R's turn. Let us now define strategies, writing $\cdot$ for concatenation.

Definition 1. *A* 0-strategy *(or* strategy*) σ on a 0-position p is a non empty, prefix-closed set of 0-plays of domain p such that, for any x in σ with even codomain q, and for any move $m : q \to r$ in $\mathcal{G}_0(G)$, $x \cdot m$ is also in σ.*

A 1-strategy *(or* strategy*) Σ on a 1-position f is a non empty prefix-closed set of 1-plays of domain f, such that, for any X in Σ with even codomain g, and for any move $M : g \to h$ in $\mathcal{G}_1(G)$, $X \cdot M$ is also in Σ.*

We use S to range over 0 or 1-strategies (or both), leaving the context to disambiguate. Given $f : p \to q$ and $g : q \to r$, we define the horizontal composition of strategies σ and σ' (respectively on f and g) to be the set of all plays on $g \circ f$ of the form $Y \bullet X$ for some (horizontally) composable $X \in \sigma$ and $Y \in \sigma'$. We easily prove that this definition is sensible:

Proposition 2. *A composition of 1-strategies is again a 1-strategy.*

Proposition 3. *The composition of 1-strategies is associative.*

The proof of the latter statement is an easy consequence of the associativity of our horizontal composition of plays.

We define the copycat strategy on an identity 1-position p as the smallest strategy containing the copycat 1-plays (as defined above) starting at p. These copycat strategies are neutral for our composition. We thus have a category Strat(G) whose objects are 0-positions, and morphisms are pairs of a 1-position and a strategy for it.

In the case of our running example G_{HO}, this new category fits with the "classical" one, up to the fact that we also consider non empty plays as objects in the new category.

Theorem 2. *The map sending an arena to the corresponding initial play yields a full embedding* $\mathrm{Strat}_{HO} \longrightarrow \mathrm{Strat}(G_{HO})$.

Next, we show that two crucial properties of strategies are stable under composition. A strategy is *deterministic* iff it does not contain two plays ending on an even position and sharing all their proper prefixes.

Proposition 4. *The composition of deterministic* 1*-strategies is again deterministic.*

A play is *final* in a strategy S when it has no extension in S. A strategy is *complete* iff its final plays all end on an even position. In other words, a complete strategy is one which never gets stuck. However, this definition is a bit loose w.r.t. potential infinite plays. Indeed, a complete strategy may contain infinite plays, and the composition of two complete strategies may not be complete. Intuitively, it may get lost in infinite internal "chattering" between M_1 and M_2. Thus, we refine the picture as follows. We deem a strategy *noetherian* iff it contains only finite plays, and *winning* iff it is noetherian and complete. This yields the following:

Proposition 5. *The composition of two winning* 1*-strategies is again winning.*

The previous notion of a winning strategy is not totally satisfactory. For instance, we would like copycat strategies to be winning. This somehow forces to consider some kind of non noetherian strategies. Anyway, we also wish to handle infinite plays in the spirit of Abramsky [1], but this is beyond the scope of the present work.

5 Conclusion

We have designed a notion of game theory. This is not one more category whose objects are new kinds of arenas. Rather we have shown how to build such a category from a very minimal set of data: a (one-way) category and two sets of morphisms therein. We have sketched how our composition of strategies has the desired stability properties (but we did not consider innocence). We hope that our framework will help in the design of new, helpful game semantics. We believe that it can be extended in various ways in order to encompass most of existing game semantics, and plan to explore some of these extensions in the near future.

Acknowledgements We thank Vincent Danos for having advised the first author's PhD thesis, Pierre-Louis Curien for his constant benevolence and assistance, and Martin Hyland for encouraging us to write the present work.

References

1. Samson Abramsky. *Semantics and Logics of Computation*, chapter Semantics of interaction, pages 1–31. Cambridge University Press, 1997.
2. Samson Abramsky, Kohei Honda, and Guy McCusker. A fully abstract game semantics for general references. In *Proceedings of the thirteenth annual symposium on Logic In Computer Science*, pages 334–344. IEEE Computer Society Press, 1998.
3. Samson Abramsky, Radha Jagadeesan, and Pasquale Malacaria. Full abstraction for PCF. *Information and Computation*, 163(2):409–470, 2000.
4. Samson Abramsky and Guy McCusker. Full abstraction for Idealized Algol with passive expressions. *Theoretical Computer Science*, 227:3–42, 1999.
5. Samson Abramsky and Paul-André Melliès. Concurrent games and full completeness. In *Proceedings of the fourteenth annual symposium on Logic In Computer Science*, pages 431–442, Trento, 1999. IEEE Computer Society Press.
6. Robin Cockett and Robert Seely. Polarized category theory, modules, and game semantics. *Theory and Applications of Categories*, 18:4–101, 2007.
7. Vincent Danos and Laurent Regnier. The structure of multiplicatives. *Archive for Mathematical Logic*, 28:181–203, 1989.
8. Russ Harmer. *Games and Full Abstraction for Nondeterministic Languages*. Ph.D. thesis, Imperial College and University of London, 1999.
9. Russ Harmer. Innocent game semantics, 2005. Course notes.
10. Russ Harmer, Martin Hyland, and Paul-André Melliès. Categorical combinatorics for innocent strategies. Technical report, Paris 7 University, 2007.
11. Michel Hirschowitz. *Jeux abstraits et composition catégorique*. Thèse de doctorat, Université Paris VII, 2004.
12. Martin Hyland and Luke Ong. On full abstraction for PCF. *Information and Computation*, 163(2):285–408, 2000.
13. James Laird. Full abstraction for functional languages with control. In *Proceedings of the twelfth annual symposium on Logic In Computer Science*, pages 58–67, Warsaw, 1997. IEEE Computer Society Press.
14. Olivier Laurent. Sémantique des jeux. Course notes (Paris VII), 2004.
15. Paul-André Melliès. Double categories: a modular model of multiplicative linear logic. *Mathematical Structures in Computer Science*, 12:449–479, 2002.