A Theory for Game Theories

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

Abstract. Game semantics is concerned with providing game models to programming languages or proof theories. Despite their success in defining fully abstract denotational semantics, game models lack count many variants, whose interrelationships largely remain to be elucidated. This raises the question: what is a game model?
This paper defines an abstract notion of game model, and provides a general constrution which, from such a game model, produces a category of games and strategies. A game model comes equipped with notions of position and move, and the construction yields the familiar notions of play and strategy. Among positions there are winning and losing ones, and among strategies, there are deterministic and winning ones. Composition of strategies has the following mandatory properties. It is associative and has the expected neutral elements (the so-called *copycat* strategies). Moreover, it preserves winning as well as deterministic strategies.
As our running example, we instantiate our construction on a particular game model, obtaining the standard category of Hyland-Ong games and strategies with the switching condition. Extensions of our framework to game models without the switching condition are handled elsewhere in the first author's PhD thesis [10].

Keywords: Game semantics, categories.

1 Introduction

The flavor problem Game semantics appeared in the early 90's [2, 11] and provided convenient denotational semantics to proof theories and programming languages, including their non functional features [1, 5, 4, 12, 8, 13]. However, game semantics counts 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, complete, 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

⋆ CEA-LIST **michel.hirschowitz@cea.fr**
⋆⋆ UNSA **ah@math.unice.fr**
⋆ ⋆ ⋆ UMR 5668 CNRS-ENS Lyon-INRIA-UCBL **tom.hirschowitz@ens-lyon.fr**

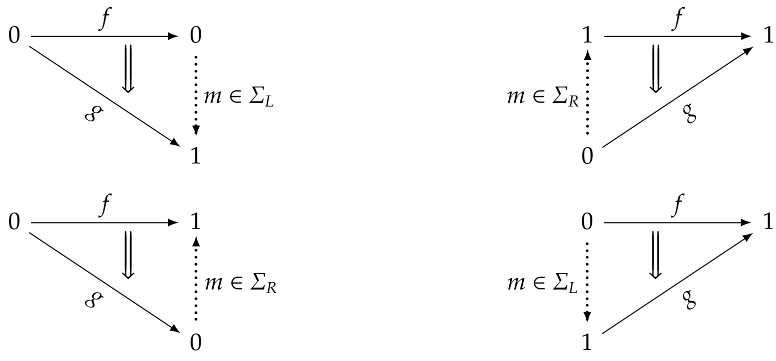

Fig. 1. The four kinds of edges in $\mathcal{G}_1(C)$, from f to g

class of game theories where the composition of strategies preserves good properties. This class contains those among existing game theories which respect the 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 [10] 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.

Playing in a one-way category In our approach, a play may take place in any "one-way" category: by a one-way category, we mean a category where objects have a sign (1/0) and where morphisms cannot go from a 1-object to a 0-one. The crucial part of our construction builds a naive game $\mathcal{G}_1(C)$ from a one-way category C. By a naive game, we mean a directed bipartite graph. Let us sketch this construction: the vertices of $\mathcal{G}_1(C)$ 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

$$11 \underset{\mathsf{L}}{\overset{\mathsf{M}L}{\rightleftarrows}} 01 \underset{\mathsf{R}}{\overset{\mathsf{M}R}{\rightleftarrows}} 00. \quad (1)$$

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	$\mathsf{M}R$	$\mathsf{M}L$
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 signs $0 \to 1$.

Interacting in a one-way category We have a similar graph $\mathcal{G}_2(C)$ for *interactions* between two plays in $\mathcal{G}_1(C)$. The vertices of $\mathcal{G}_2(C)$ are now composable pairs of morphisms in C. Hence, we have four kinds of vertices $(000, 001, 011, 111)$ according to the signs of objects, yielding the following state diagram:

$$111 \underset{\mathsf{M}_1 L}{\overset{\mathsf{L}}{\rightleftarrows}} 011 \underset{\mathsf{M}_1 R}{\overset{\mathsf{M}_2 L}{\leftrightarrows}} 001 \underset{\mathsf{R}}{\overset{\mathsf{M}_2 R}{\rightleftarrows}} 000. \qquad (2)$$

The six kinds of transitions correspond to six kinds of commutative diagrams in C, as displayed in Figure 2. In $\mathcal{G}_2(C)$, there are two players M_1 and M_2, and one opponent, who interacts on the left with M_1 and on the right with M_2. Thanks to categorical composition, both players act exactly as if they were facing one opponent (see Figure 2). Actually, M_1 interacts with Opponent on the left, and with M_2 on the right. 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.

Using these interactions, it is straightforward to define the "natural" composition of strategies for $G_1(C)$.

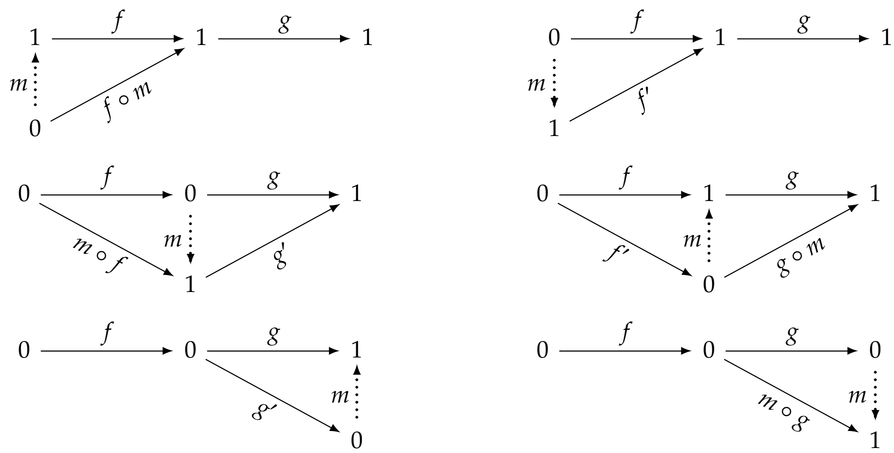

Fig. 2. The six kinds of edges in $\mathcal{G}_2(C)$ (each edge top-down)

Game settings The class of game theories we have constructed so-far, starting from one-way categories, have a perfectly natural composition of strategies. But the corresponding games are desperately balanced, because players can easily neutralize each other. More meaningful game theories are obtained simply by restricting the set of odd morphisms allowed in the process of composition/decomposition. A priori we should specify four sets M_{OL}, M_{OR}, M_{PL}, M_{PR}

of legal odd morphisms, one for each of our four kinds of moves. However, these restrictions are compatible with interactions only if we impose $M_{OL} = M_{PR}$ and $M_{OR} = M_{PL}$. Hence our definition: a *game setting* (C, Σ_L, Σ_R) consists of a one-way category C equipped with a pair of sets (or classes) of odd morphisms: Σ_L is the set of *forward* moves (those going downwards in Figure 1); accordingly, Σ_R is the set of *backward* moves.

Our notion of game setting comes equipped with a rich structure on the set (or class) of plays. Indeed, plays are cells in a double-category where vertical composition corresponds to concatenation and horizontal composition corresponds to interaction.

Related and further work Cockett and Seely [**?**] offer another categorical investigation into game semantics. The relationship between their work and ours remains unclear to us. A recent preprint by Harmer, Hyland, and Melliès [**?**] explores a categorical characterization of innocent strategies.

Organization of the paper Starting from the concrete, we first exhibit a double category (of plays) hidden in a standard variant of Hyland and Ong's (HO) game theory. We then provide the categorical construction which, from a so-called *game setting*, constructs such a double category, where we describe strategies and their composition. We show that the obtained notion of strategies closely corresponds to the standard one.

2 The one way category underlying Hyland-Ong games

2.1 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 $m' \in M_A$, $m' \nvdash_A m$, and (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 $s = m_1 \ldots m_n$ is a sequence of moves of alternate signs in A, and ρ is a function from $\{1 \ldots n\}$ to $\{0 \ldots n-1\}$ such that for all $i \in \{1 \ldots n\}$ (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.

A position of length 0 is said *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.

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' 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 *valid* positions p of $A \multimap B$ to be those whose projections are indeed positions 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 this shuffle satisfies the *switching condition* when if $n_A + n_B > 0$ then $p_S(1)$ is on the B-side, and, for i satisfying $1 < 2i < n_A + n_B$, $p_S(2i)$ and $p_S(2i + 1)$ are on the same side. We say that p_S is an *even* shuffle. 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).

It is then standard to define strategies from A to B to be non-empty, prefix-closed sets of valid plays in $A \multimap B$ containing at least one answer to each move by Opponent (a bit more formally: containing at least one 1-extension of each odd play in itself). One then shows that strategies compose and have identities, which makes them into a (locally small) *category of games and strategies* Strat.

2.2 The one-way category C_{HO}

Let us now describe the (locally small) "one-way" category C_{HO} relevant for HO games. An object $(A, (s, \rho))$ of C_{HO} is merely a position (s, ρ) in game arena A, while a morphism from $p = (A, (s, \rho))$ to $q = (B, (s', \rho'))$ is a 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. Also note that since morphisms from p to q have positions (p and q) as projections, we could simply have required them to be pre-positions. Not doing it is a matter of readability. Conversely, a pre-position in $A \multimap B$ determines a morphism exactly when its projections are positions.

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 $\mathrm{Pred}(p)$. 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 $\mathrm{Pred}(p)$ while if p is odd, then Sub_p goes from $\mathrm{Pred}(p)$ to p.

Next, we define the composition of our morphisms. Assume two consecutive arrows, i.e., 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. It is easily checked that the identity on a position p is the copycat morphism $\mathrm{copycat}_p$.

This altogether gives a category $\mathcal{C}_{HO}$, whose objects are moreover signed: 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 remark the following fact.

Proposition 1. *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. Thus, for example, an initial arrow position is always of kind $0 \to 0$.

Our category may thus be seen as equipped with a functor to $\mathbf{2}$, the ordered set $0 \leq 1$, in other words it is a one-way category.

Remark 1 (Relaxing the switching condition). If we were to relax the switching condition, and allow Opponent to switch sides in an arrow game, instead of functors to $\mathbf{2}$, we would simply consider functors to $\mathbb{2}$, the preordered set with $0 \leq 1$ and $1 \leq 0$. This approach has been pursued in the first author's PhD thesis [10], and gives rise, instead of a double category of plays, to a kind of weak double category of plays. One eventually recovers a proper category when passing to strategies.

2.3 The game setting for HO

Now we explain how HO moves may be seen as morphisms in $\mathcal{C}_{HO}$. 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 class of HO-moves is precisely the class of subcopycat morphisms, which we split into the class of RHO-moves, where the length of the codomain exceeds the length of the domain by one, and the class of LHO-moves, where the length of the domain exceeds the length of the codomain by one.

Thus, a play p of length n in A may be seen as an "anti-composable" sequence of HO-moves in the sense just defined, namely if $p_0, \ldots, p_n$ is the corresponding sequence of positions, our sequence of moves is $\mathrm{Sub}_{p_1}, \ldots, \mathrm{Sub}_{p_n}$; and instead of being composable (which is anyway impossible for odd morphisms) our sequence of moves is *anti-composable* in the following sense: the sequence $(m_1, \ldots, m_n)$ is anti-composable, if for $1 \leq i < n$, m_i and m_{i+1} have the same domain (for i even) or the same codomain (for i odd), as displayed in Figure 3.

$$
\begin{array}{llllll}
\text{Moves} & p_0 \xrightarrow{m_1} & p_1 & \xleftarrow{m_2} p_2 & \ldots & p_n \\
\text{Signs} & 0 \longrightarrow & 1 & \longleftarrow 0 & \ldots & O \text{ or } 1.
\end{array}
$$

Fig. 3. Anti-composable sequences of HO moves

Now let us see how our view fits with plays in an arrow game: let us consider a position f in the game $A \multimap B$ and its extension to a new position g, through a HO-move m (in A or in B, not in $A \multimap B$). 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 $g = m \circ f$,
- 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.

This analysis of arrow plays as sequences of compositions and decompositions of arrow positions in C_{HO} may in fact be carried out independently from the underlying game setting, as we show in the next section. Furthermore, plays in arrow games determine the category of games and strategies, which allows our construction to recover the category of games and strategies Strat of Section 2.1 from the triple $G_{HO} \triangleq (C_{HO}, \text{LHO-moves}, \text{RHO-moves})$.

3 The Abstract Framework: building the double category

3.1 The 0-dimensional game

As announced, we define a *game setting* G to be a tuple $(C_\lambda, \Sigma_L, \Sigma_R)$, where C_λ is a *one-way* category, and Σ_L and Σ_R are classes of morphisms of C_λ. Recall that a one-way category is a category C, equipped with a functor $\lambda : C \to \mathbb{2}$, where $\mathbb{2}$ is the ordered set $0 \leq 1$. This amounts to labeling each object of C with a sign (0 or 1), and requiring C to not have arrows from 1-labeled objects to 0-labeled ones. Furthermore, the morphisms in Σ_L and Σ_R, respectively called *0*- and *1*-*moves*, are required to change sign (i.e., to be from 0 to 1).

In this setting, we view objects as positions in a two-players game, actually a signed graph. Morphisms in Σ_L and Σ_R 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 4, from some 0-labeled position p, Opponent plays by choosing a 0-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 1-move $m' : r \to q$ with codomain q, thereby reaching the position r. This defines a graph between the objects of $\mathbb{C}$, which we call the 0-*dimensional game* (0D for short) of G and write $\mathcal{G}_0(G)$. We call the free category over this graph the category of 0*D plays* over G, and write it $\mathbb{C}_0(G)$.

$$\begin{array}{lccccccc} \text{Play:} & p & \xrightarrow{\ m\ } & q & \xleftarrow{\ m'\ } & r & & \ldots \\ \text{Signs:} & 0 & & 1 & & 0 & & \end{array}$$

- - - - - - course of the game - - - - - ▸

Fig. 4. Example play in the 0-game

As in the HO setting, each 0D play v has a predecessor $\mathrm{Pred}(v)$ obtained by deleting the last move (if any) .

3.2 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* (1D for short) of G and write $\mathcal{G}_1(G)$. 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 0D game. Positions (or vertices) in $\mathcal{G}_1(G)$ are morphisms in $\mathbb{C}$. Given the constraints on signs, there are just three kinds of positions:

$$0 \longrightarrow 0 \qquad\qquad 0 \longrightarrow 1 \qquad\qquad 1 \longrightarrow 1.$$

Then, moves from f to g in the 1D game are defined to be commutative triangles in $\mathbb{C}$, of one of the shapes in Figure 1.

The interpretation of signs in 1D games, illustrated in Figure 5, entirely follows from the idea that in 0D games, Player lives on the left of the position, whilst Opponent lives on its right. For a 1D 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 0D game, and Player in the codomain 0D game. L plays Player in the domain 0D game, whilst R plays Opponent in the codomain 0D game. This forces on us the following rule for the 1D game:

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

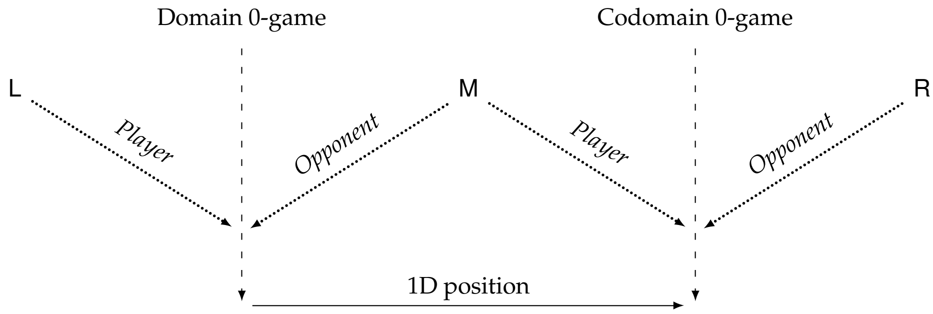

Fig. 5. All agents (L, M, R) act as Player on their right and as Opponent on their left

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

Finally, we define the source and target, respectively written s and t, of 1D plays as functors $C_1(G) \to C_0(G)$ by the obvious induction. We thus have a notion of a composable pair of 1D plays, which we call *interaction* (actually a morphism in the pullback category $C_1(G)_{\mathsf{s}}\times_{\mathsf{t}} C_1(G)$). We call *length* of an interaction the sum of the lengths of its two components.

3.3 1D horizontal composition

We know define 1D horizontal composition, which takes an interaction and returns the composite play. First, we define $\mathcal{G}_2(G)$ exactly as in Fig. 2, and remark that each edge in $\mathcal{G}_2(G)$, i.e., each configuration in Fig. 2, is indeed an interaction, and we call them *primitive interactions*. Let us also deem the primitive interactions of the middle row *internal*, and the other ones *external*. The key fact is then that:

Lemma 1. *Interactions are in bijection with sequences of primitive interactions, i.e., paths in* $\mathcal{G}_2(G)$*.*

We then define the 1D horizontal composition $Y \bullet X$, or $\bullet(X, Y)$, of an interaction (X, Y) by induction on the length of this decomposition γ: if γ is the empty path on $g \circ f$, then return the empty 1D play on $g \circ f$; otherwise $\gamma = \gamma' \cdot I$ for some γ' and primitive interaction I. If I is internal, then $Y \bullet X \triangleq \bullet(\gamma')$. Otherwise, $Y \bullet X$ is $\bullet(\gamma')$, followed by the obvious move from $g \circ f$ to $g \circ f \circ m$, $m \circ g \circ f$, $g \circ f'$, or $g' \circ f$, for each respective external interaction in Fig. 2.

3.4 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 want to show that horizontal composition of 1D plays yields a category, whose objects are 0D plays, and whose morphisms are 1D plays. The missing data is an identity morphism, which we define 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 0-moves $m : p \to p'$ and 1-moves $m : p' \to p$ are respectively sent to plays

By the standard adjunction between categories and directed graphs, these two elementary copycat plays define the functor copycat uniquely. On arbitrary plays, easy calculations show that copycat simply piles up sequences of such elementary plays.

Proposition 2. *Horizontal composition of* 1D *plays is associative and unital.*

The proof of associativity is analogous to our definition of horizontal composition. It relies on the analysis of the graph of 3-interactions: its vertices are triples of composable 1D plays. Concerning 3-interactions, we obtain a decomposition result completely analogous to Lemma 1.

This all gives the data for a *double category*. A short definition is as follows: a double category is a category object in **Cat**, the category of small categories. Here, for size reasons, we should rather define it as a category object in **CAT**, the category of not necessarily small categories. A more explicit, elementary definition may be found, e.g., in Melliès [14]. We've checked all the required properties, but 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.*

3.5 An abstract view of strategies

In this section, we show that some results on strategies may be understood abstractly in our framework. Recall that an object of C is *even* when its sign is 0 and *odd* otherwise. For 1-positions and say that $f : p \to q$ 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 all 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 all* X *in* Σ *with even codomain* g*, and for all 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 disambiguate. Given $f : p \to q$ and $g : q \to r$, we define horizontal composition of strategies σ and σ' (respectively on f and g) to be all plays obtained by composition of plays from σ and σ'. More formally, we take the set of all plays on $g \circ f$ equal to $Y \bullet X$ for some (horizontally) composable $X \in \sigma$ and $Y \in \sigma'$. We easily prove that this definition is sensible:

Proposition 3. *A composition of* 1*-strategies is again a* 1*-strategy.*

Proposition 4. *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 plays as defined above. 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. This brings us back to where we started:

Theorem 2. *There is a full embedding* Strat $\longrightarrow$ Strat(G_{HO}).

Next, we show that two crucial properties of strategies are stable under composition.

A strategy is said *deterministic* whenever it does not contain two plays ending on an even position and sharing all their proper prefixes.

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

A play is said *final* in a strategy S when it has no extension in S.

Definition 2. *A* complete *strategy is then one whose all final plays 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 ping-pong plays 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 6. *The composition of two winning* 1*-strategies is again winning.*

Noetherian and complete strategies define a notion of a winning strategy, which is not really satisfactory. For instance, we would like identity positions 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 [3], but this is beyond the scope of the present work.

4 Conclusion

We have designed a new 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 properties of stability. 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.

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, 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, Indianapolis, June 1998. IEEE, IEEE Computer Society Press.
2. Samson Abramsky, Radha Jagadeesan, and Pasquale Malacaria. Full abstraction for PCF. *Information and Computation*, 163(2):409–470, December 2000.
3. Samson Abramsky and Guy McCusker. Game semantics. In Helmut Schwichtenberg and Ulrich Berger, editors, *Logic and Computation: Proceedings of the 1997 Marktoberdorf Summer School*. Springer, 1998.
4. Samson Abramsky and Guy McCusker. Full abstraction for idealized algol with passive expressions. *Theoretical Computer Science*, 227:3–42, September 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, July 1999. IEEE, 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. Russel Harmer. *Games and Full Abstraction for Nondeterministic Languages*. Ph.D. thesis, Imperial College and University of London, 1999.
9. Russel Harmer. Innocent game semantics. http://www.pps.jussieu.fr/ russ/GS.pdf, 2005. Course notes.
10. Michel Hirschowitz. *Un cadre catégorique abstrait pour la sémantique des jeux*. Thèse de doctorat, Université Paris VII, December 2004.
11. Martin Hyland and Luke Ong. On full abstraction for PCF. *Information and Computation*, 163(2):285–408, December 2000.
12. 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, June 1997. IEEE, IEEE Computer Society Press.
13. Olivier Laurent. Sémantique des jeux. Notes de cours du DEA de Programmation (Paris VII), April 2004.
14. Paul-André Melliès. Double categories: a modular model of multiplicative linear logic. *Mathematical Structures in Computer Science*, 12:449–479, 2002.