



Non-angelic Concurrent Game Semantics

Simon Castellan^{1(✉)}, Pierre Clairambault², Jonathan Hayman³,
and Glynn Winskel³

¹ Imperial College London, London, UK
`simon@phis.me`

² Univ Lyon, CNRS, ENS de Lyon, UCB Lyon 1, LIP, Lyon, France

³ Computer Laboratory, University of Cambridge, Cambridge, UK

Abstract. The *hiding* operation, crucial in the compositional aspect of game semantics, removes computation paths not leading to observable results. Accordingly, games models are usually biased towards *angelic* non-determinism: diverging branches are forgotten.

We present here new categories of games, not suffering from this bias. In our first category, we achieve this by avoiding hiding altogether; instead morphisms are *uncovered* strategies (with neutral events) up to *weak bisimulation*. Then, we show that by hiding only certain events dubbed *inessential* we can consider strategies up to *isomorphism*, and still get a category – this partial hiding remains sound up to weak bisimulation, so we get a concrete representations of programs (as in standard concurrent games) while avoiding the angelic bias. These techniques are illustrated with an interpretation of affine nondeterministic PCF which is adequate for weak bisimulation; and may, must and fair convergences.

1 Introduction

Game semantics represents programs as strategies for two player games determined by the types. Traditionally, a strategy is simply a collection of execution traces, each presented as a play (a structured sequence of events) on the corresponding game. Beyond giving a compositional framework for the formal semantics of programming languages, game semantics proved exceptionally versatile, providing very precise (often fully abstract) models of a variety of languages and programming features. One of its rightly celebrated achievements is the realisation that combinations of certain effects, such as various notions of state or control, could be characterised via corresponding conditions on strategies (innocence, well bracketing, . . .) in a single unifying framework. This led Abramsky to propose the *semantic cube* programme [1], aiming to extend this success to further programming features: concurrency, non-determinism, probabilities, etc. . .

However, this elegant picture soon showed some limitations. While indeed the basic category of games was successfully extended to deal with concurrency [10, 13], non-determinism [11], and probabilities [9] among others, these extensions (although fully abstract) are often incompatible with each other, and really, incompatible as well with the central condition of innocence. Hence a semantic

hypercube encompassing all these effects remained out of reach. It is only recently that some new progress has been made with the discovery that some of these effects could be reconciled in a more refined, more intensional games framework. For instance, in [6, 16] innocence is reconciled with non-determinism, and in [15] with probabilities. In [7], innocence is reconciled with concurrency.

But something is still missing: the works above dealing with non-deterministic innocence consider only *may-convergence*; they ignore execution branches leading to divergence. To some extent this seems to be a fundamental limitation of the game semantics methodology: at the heart of the composition of strategies lies the *hiding* operation that removes unobservable events. Diverging paths, by nature non-observable, are forgotten by hiding. Some models of must-testing do exist for particular languages, notably McCusker and Harmer’s model for non-deterministic Idealized Algol [11]; the model works by annotating strategies with *stopping traces*, recording where the program may diverge. But this approach again mixes poorly with other constructions (notably innocence), and more importantly, is tied to may and must equivalences. It is not clear how it could be extended to support richer notions of convergence, such as *fair-testing* [2].

Our aim is to present a basis for non-deterministic game semantics which, besides being compatible with innocence, concurrency, *etc.*, is not biased towards may-testing; it is *non-angelic*. It should not be biased towards must-testing either; it should in fact be *agnostic* with respect to the testing equivalence, and support them all. Clearly, for this purpose it is paramount to remember the non-deterministic branching information; indeed in the absence of that information, notions such as *fair-testing* are lost. In fact, there has been a lot of activity in the past five years or so around games model that *do* observe the branching information. It is a feature of Hirschowitz’s work presenting strategies as presheaves or sheaves on certain categories of cospans [12]; of Tsukada and Ong’s work on nondeterministic innocence via sheaves [16]; and of our own line of work presenting strategies as certain event structures [5, 7, 14].

But observing branching information is not sufficient. Of the works mentioned above, those of Tsukada and Ong and our own previous work are still angelic, because they rely on hiding for composition. On the other hand, Hirschowitz’s work gets close to achieving our goals; by refraining from hiding altogether, his model constructs an agnostic and precise representation of the operational behaviour of programs, on which he then considers fair-testing. But by not considering hiding he departs from the previous work and methods of game semantics, and from the methodology of denotational semantics. In contrast, we would like an agnostic games model that still has the categorical structure of traditional semantics. A games model with partial hiding was also recently introduced by Yamada [18], albeit for a different purpose: he uses partial hiding to represent normalization steps, whereas we use it to represent fine-grained nondeterminism.

Contributions. In this paper, we present the first category of games and strategies equipped to handle non-determinism, but agnostic with respect to the notion of convergence (including fair convergence). We showcase our model by interpreting \mathbf{APCF}_+ , an affine variant of non-deterministic PCF: it is the

simplest language featuring the phenomena of interest. We show adequacy with respect to may, must and fair convergences. The reader will find in the first author’s PhD thesis [3] corresponding results for full non-deterministic PCF (with detailed proofs), and an interpretation of a higher-order language with shared memory concurrency. In [3], the model is proved compatible with our earlier notions of innocence, by establishing a result of full abstraction for may equivalence, for nondeterministic PCF. We have yet to prove full abstraction in the fair and must cases; finite definability does not suffice anymore.

Outline. We begin Sect. 2 by introducing \mathbf{APCF}_+ . To set the stage, we describe an angelic interpretation of \mathbf{APCF}_+ in the category \mathbf{CG} built in [14] with strategies up to isomorphism, and hint at our two new interpretations. In Sect. 3, starting from the observation that the cause of “angelism” is hiding, we omit it altogether, constructing an *uncovered* variant of our concurrent games, similar to that of Hirschowitz. Despite not hiding, when restricting the location of non-deterministic choices to internal events, we can still obtain a category up to *weak bisimulation*. But weak bisimulation is not perfect: it does not preserve must-testing, and is not easily computed. So in Sect. 4, we reinstate some hiding: we show that by hiding all synchronised events except some dubbed *essential*, we arrive at the best of both worlds. We get an agnostic category of games and strategies *up to isomorphism*, and we prove our adequacy results.

2 Three Interpretations of Affine Nondeterministic PCF

2.1 Syntax of \mathbf{APCF}_+

The language \mathbf{APCF}_+ extends affine PCF with a nondeterministic boolean choice, **choice**. Its types are $A, B ::= \mathbb{B} \mid A \multimap B$, where $A \multimap B$ represents affine functions from A to B . The following grammar describes terms of \mathbf{APCF}_+ :

$$M, N ::= x \mid MN \mid \lambda x. M \mid \mathbf{tt} \mid \mathbf{ff} \mid \mathbf{if} M N_1 N_2 \mid \mathbf{choice} \mid \perp$$

Typing rules are standard, we show application and conditionals. As usual, a conditional eliminating to arbitrary types can be defined as syntactic sugar.

$$\frac{\Gamma \vdash M : A \multimap B \quad \Delta \vdash N : A}{\Gamma, \Delta \vdash MN : B} \quad \frac{\Gamma \vdash M : \mathbb{B} \quad \Delta \vdash N_1 : \mathbb{B} \quad \Delta \vdash N_2 : \mathbb{B}}{\Gamma, \Delta \vdash \mathbf{if} M N_1 N_2 : \mathbb{B}}$$


The first rule is *multiplicative*: Γ and Δ are disjoint. The operational semantics is that of PCF extended with the (only) two nondeterministic rules $\mathbf{choice} \rightarrow \mathbf{tt}$ and $\mathbf{choice} \rightarrow \mathbf{ff}$.

2.2 Game Semantics and Event Structures

Game semantics interprets an open program by a strategy, recording the behaviour of the program (Player) against the context (Opponent) in a 2-player game. Usually, the executions recorded are represented as *plays*, *i.e.* linear

sequences of computational events called *moves*; a strategy being then a set of such plays. For instance, the nondeterministic boolean would be represented as the (even-prefix closure of the) set of plays $\{q^- \cdot \mathbf{tt}^+, q^- \cdot \mathbf{ff}^+\}$ on the game for booleans. In the play $q^- \cdot \mathbf{tt}^+$, the context starts the computation by asking the value of the program (q^-) and the program replies (\mathbf{tt}^+). Polarity indicates the origin (Program (+) or Opponent/Environment (-)) of the event.


Being based on sequences of moves, traditional game semantics handles concurrency via interleavings [10]. In contrast, in concurrent games [14], plays are generalised to partial orders which can express concurrency as a primitive. For instance, the execution of a parallel implementation of **and** against the context $(\mathbf{tt}, \mathbf{tt})$ gives the following partial order:



In this picture, the usual chronological linear order is replaced by an explicit partial order representing **causality**. Moves are concurrent when they are incomparable (as the two Player questions here). Following the longstanding convention in game semantics, we show which component of the type a computational event corresponds to by displaying it under the corresponding occurrence of a ground type. For instance in this diagram, Opponent first triggers the computation by asking the output value, and then **and** concurrently evaluates his two arguments. The arguments having evaluated to \mathbf{tt} , **and** can finally answer Opponent’s initial question and provide the output value.

In [7], we have shown how deterministic pure functional parallel programs can be interpreted (in a *fully abstract* way) using such representations.

Partial-Orders and Non-determinism. To represent nondeterminism in this partial order setting, one possibility is to use sets of partial orders [4]. This representation suffers however from two drawbacks: firstly it forgets the point of non-deterministic branching; secondly, one cannot talk of an *occurrence* of a move independently of an execution. Those issues are solved by moving to *event structures* [17], where the nondeterministic boolean can be represented as:




The wiggly line (\sim) indicates *conflict*: the boolean values cannot coexist in an execution. Together this forms an *event structure*, defined formally later.

2.3 Interpretations of APCF_+ with Event Structures


Let us introduce informally our interpretations by showing which event structures they associate to certain terms of APCF_+ .

Angelic Covered Interpretation. Traditional game semantics interpretations of nondeterminism are angelic (with exceptions, see *e.g.* [11]); they only describe what terms may do, and forget where they might get stuck. The interpretation of $M = (\lambda b. \text{if } b \text{ tt } \perp)$ choice for instance, in usual game semantics is the same as that of tt . This is due to the nature of composition which tends to forget paths that do not lead to a value. Consider the strategy for the function $\lambda b. \text{if } b \text{ tt } \perp$:



The interpretation of M arises as the *composition* of this strategy with the nondeterministic boolean. Composition is defined in two steps: interaction (Fig. 1a) and then hiding (Fig. 1b). Hiding removes intermediate behaviour which does not correspond to visible actions in the output type of the composition.

Hiding is crucial in order for composition to satisfy basic categorical properties (without it, the identity candidate, copycat, is not even idempotent). Strategies on event structures are usually considered *up to isomorphism*, which is the strongest equivalence relation that makes sense. Without hiding, there is no hope to recover categorical laws up to isomorphism. However, it turns out that, treating events in the middle as τ -transitions ($*$ in Fig. 1a), weak bisimulation equates enough strategies to get a category. Following these ideas, a category of *uncovered* strategies up to *weak bisimilarity* is built in Sect. 3.



(a) Interp. before hiding

(b) Interp. after hiding

(c) Interp. with partial hiding

Fig. 1. Three interpretations of $(\lambda b. \text{if } b \text{ tt } \perp)$ choice

Interpretation with Partial Hiding. However, considering uncovered strategies up to weak bisimulation blurs their concrete nature; *causal information* is lost, for instance. Moreover checking for weak bisimilarity is computationally expensive, and because of the absence of hiding, a term evaluating to **skip** may yield a very large representative. However, there is a way to cut down the strategies to reach a compromise between hiding *no* internal events, or hiding *all* of them and collapsing to an angelic interpretation.

In our games based on event structures, having a non-ambiguous notion of an occurrence of event allows us to give a simple definition of the internal events we need to retain (Definition 9). Hiding other internal events yields a strategy still weakly bisimilar to the original (uncovered) strategy, while allowing us to get a category *up to isomorphism*. The interpretation of M in this setting appears in Fig. 1c. As before, only the events under the result type (not labelled $*$) are now *visible*, *i.e.* observable by a context. But the events corresponding to the argument evaluation are only partially hidden; those remaining are considered *internal*, treated like τ -transitions. Because of their presence, the partial hiding performed loses no information (*w.r.t.* the uncovered interpretation) up to weak bisimilarity. But we have hidden enough so that the required categorical laws between strategies hold *w.r.t.* isomorphism. The model is more precise and concrete than that of weak bisimilarity, preserves causal information and preserves must-convergence (unlike weak bisimilarity).

Following these ideas, a category of partially covered strategies up to iso (the target of our adequacy results) is constructed in Sect. 4.

3 Uncovered Strategies up to Weak Bisimulation

We now construct a category of “uncovered strategies”, up to weak bisimulation. Uncovered strategies are very close to the *partial strategies* of [8], but [8] focused on connections with operational semantics rather than categorical structure.

3.1 Preliminaries on Event Structures

Definition 1. An *event structure* is a triple $(E, \leq_E, \text{Con}_E)$ where (E, \leq_E) is a partial-order and Con_E is a non-empty collection of finite subsets of E called consistent sets *subject to the following axioms*:

- If $e \in E$, the set $[e] = \{e' \in E \mid e' \leq e\}$ is finite,
- For all $e \in E$, the set $\{e\}$ is consistent,
- For all $Y \in \text{Con}_E$, for all $X \subseteq Y$, then $X \in \text{Con}_E$.
- If $X \in \text{Con}_E$ and $e \leq e' \in X$ then $X \cup \{e\}$ is consistent.

A down-closed subset of events whose finite subsets are all consistent is called a **configuration**. The set of finite configurations of E is denoted $\mathcal{C}(E)$. If $x \in \mathcal{C}(E)$ and $e \notin x$, we write $x \xrightarrow{e} x'$ when $x' = x \cup \{e\} \in \mathcal{C}(E)$; this is the **covering relation** between configurations, and we say that e gives an **extension** of x .

Two extensions e and e' of x are **compatible** when $x \cup \{e, e'\} \in \mathcal{C}(E)$, **incompatible** otherwise. In the latter case, we have a **minimal conflict** between e and e' **in context** x (written $e \rightsquigarrow_x e'$).

These event structures are based on *consistent sets* rather than the more commonly-encountered binary *conflict* relation. Consistent sets are more general, and more handy mathematically, but throughout this paper, event structures concretely represented in diagrams will only use *binary conflict*, i.e. the relation $e \rightsquigarrow_x e'$ does not depend on x , meaning $e \rightsquigarrow_y e'$ whenever y extends with e , and with e' – in which case we only write $e \rightsquigarrow e'$. Then consistent sets can be recovered as those finite $X \subseteq E$ such that $\neg(e \rightsquigarrow e')$ for all $e, e' \in X$. Our diagrams display the relation \rightsquigarrow , along with the *Hasse diagram* of \leq_E , called **immediate causality** and denoted by \rightarrow_E . All the diagrams above denote event structures. The missing ingredient in making the diagrams formal is the *names* accompanying the events ($\mathbf{q}, \mathbf{tt}, \mathbf{ff}, \dots$). These will arise as annotations by events from *games*, themselves event structures, representing the types.

The **parallel composition** $E_0 \parallel E_1$ of event structures E_0 and E_1 has for *events* $(\{0\} \times E_0) \cup (\{1\} \times E_1)$. The *causal order* is given by $(i, e) \leq_{E_0 \parallel E_1} (j, e')$ when $i = j$ and $e \leq_{E_i} e'$, and *consistent sets* by those finite subsets of $E_0 \parallel E_1$ that project to consistent sets in both E_0 and E_1 .

A **(partial) map of event structures** $f : A \rightarrow B$ is a (partial) function on events which (1) maps any finite configuration of A to a configuration of B , and (2) is locally injective: for $a, a' \in x \in \mathcal{C}(A)$ and $fa = fa'$ (both defined) then $a = a'$. We write \mathcal{E} for the category of event structures and total maps and \mathcal{E}_\perp for the category of event structures and partial maps.

An **event structure with partial polarities** is an event structure A with a map $pol : A \rightarrow \{-, +, *\}$ (where events are labelled “negative”, “positive”, or “internal” respectively). It is a **game** when no events are internal. The dual A^\perp of a game A is obtained by reversing polarities. Parallel composition naturally extends to games. If x and y are configurations of an event structure with partial polarities we use $x \subseteq^p y$ where $p \in \{-, +, *\}$ for $x \subseteq y$ & $pol(y \setminus x) \subseteq \{p\}$.

Given an event structure E and a subset $V \subseteq E$ of events, there is an event structure $E \downarrow V$ whose events are V and causality and consistency are inherited from E . This construction is called the **projection** of E to V and is used in [14] to perform hiding during composition.

3.2 Definition of Uncovered Pre-strategies

As in [14], we first introduce *pre-strategies* and their composition, and then consider *strategies*, those pre-strategies well-behaved with respect to copycat.

Uncovered Pre-strategies. An **uncovered pre-strategy** on a game A is a partial map of event structures $\sigma : S \rightarrow A$. Events in the domain of σ are called **visible** or **external**, and events outside **invisible** or **internal**. Via σ , visible events inherit polarities from A .

Uncovered pre-strategies are drawn just like the usual strategies of [14]: the event structure S has its events drawn as their labelling in A if defined or $*$ if

undefined. The drawing of Fig. 1a is an example of an uncovered pre-strategy. From an (uncovered) pre-strategy, one can get a pre-strategy in the sense of [14]: for $\sigma : S \rightarrow A$, define $S_{\downarrow} = S \downarrow \text{dom}(\sigma)$ where $\text{dom}(\sigma)$ is the domain of σ . By restriction σ yields $\sigma_{\downarrow} : S_{\downarrow} \rightarrow A$, called a **covered pre-strategy**. A configuration x of S can be decomposed as the disjoint union $x_{\downarrow} \cup x_*$ where x_{\downarrow} is a configuration of S_{\downarrow} and x_* a set of internal events of S .

A pre-strategy **from a game A to a game B** is a (uncovered) pre-strategy on $A^{\perp} \parallel B$. An important pre-strategy from a game A to itself is the **copycat pre-strategy**. In $A^{\perp} \parallel A$, each move of A appears twice with dual polarity. The copycat pre-strategy α_A simply waits for the negative occurrence of a move a before playing the positive occurrence. See [5] for a formal definition.

Isomorphism of strategies [14] can be extended to uncovered pre-strategies:

Definition 2. *Pre-strategies $\sigma : S \rightarrow A, \tau : T \rightarrow A$ are **isomorphic** (written $\sigma \cong \tau$) if there is an iso $\varphi : S \cong T$ s.t. $\tau \circ \varphi = \sigma$ (equality of partial maps).*

Interaction of Pre-strategies. Recall that in the covered case, composition is performed first by interaction, then hiding; where interaction of pre-strategies is described as their pullback in the category of *total maps* [14]. Even though \mathcal{E}_{\perp} has pullbacks, those pullbacks are inadequate to describe interaction. In [8], uncovered strategies are seen as total maps $\sigma : S \rightarrow A \parallel N$, and their interaction as a pullback involving these. This method has its awkwardness so, instead, here we give a direct universal construction of interaction, replacing pullbacks.

We start with the simpler case of a **closed** interaction of a pre-strategy $\sigma : S \rightarrow A$ against a counter pre-strategy $\tau : T \rightarrow A^{\perp}$. As in [5] we first describe the expected *states* of the closed interaction in terms of *secured bijections*, from which we construct an event structure; before characterising the whole construction via a universal property.

Definition 3 (Secured bijection). *Let \mathbf{q}, \mathbf{q}' be partial orders and $\varphi : \mathbf{q} \simeq \mathbf{q}'$ be a bijection between the carrier sets (non necessarily order-preserving). It is **secured** when the following relation \triangleleft_{φ} on the graph of φ is acyclic:*

$$(s, \varphi(s)) \triangleleft_{\varphi} (s', \varphi(s')) \text{ iff } s \rightarrow_{\mathbf{q}} s' \vee \varphi(s) \rightarrow_{\mathbf{q}'} \varphi(s') \quad (1)$$

If so, the resulting partial order $(\triangleleft_{\varphi})^$ is written \leq_{φ} .*

Let $\sigma : S \rightarrow A$ and $\tau : T \rightarrow A$ be partial maps of event structures (we dropped polarities, as the construction is completely independent of them). A pair $(x, y) \in \mathcal{C}(S) \times \mathcal{C}(T)$ such that $\sigma_{\downarrow}x = \tau_{\downarrow}y \in \mathcal{C}(A)$, induces a bijection $\varphi_{x,y} : x \parallel y_* \simeq x_* \parallel y$ defined by local injectivity of σ and τ :

$$\begin{aligned} \varphi_{x,y}(0, s) &= (0, s) & (s \in x_*) \\ \varphi_{x,y}(0, s) &= (1, \tau^{-1}(\sigma s)) & (s \in x_{\downarrow}) \\ \varphi_{x,y}(1, t) &= (1, t) \end{aligned}$$

The configurations x and y have a partial order inherited from S and T . Viewing y_* and x_* as discrete orders (the ordering relation is the equality), $\varphi_{x,y}$

is a bijection between carrier sets of partial orders. An **interaction state** of σ and τ is $(x, y) \in \mathcal{C}(S) \times \mathcal{C}(T)$ with $\sigma \downarrow x = \tau \downarrow y$ for which $\varphi_{x,y}$ is secured. As a result (the graph of) $\varphi_{x,y}$ is naturally partial ordered. Write $\mathcal{S}_{\sigma,\tau}$ for the set of interaction states of σ and τ . As usual [5], we can recover an event structure:

Definition 4 (Closed interaction of uncovered pre-strategies). *Let A be an event structure, and $\sigma : S \rightarrow A$ and $\tau : T \rightarrow A$ be partial maps of event structures. The following data defines an event structure $S \wedge T$:*

- events: those interaction states (x, y) such that $\varphi_{x,y}$ has a top element,
- causality: $(x, y) \leq_{S \wedge T} (x', y')$ iff $x \subseteq x'$ and $y \subseteq y'$,
- consistency: a finite set of interaction states $X \subseteq S \wedge T$ is consistent iff its union $\bigcup X$ is an interaction state in $\mathcal{S}_{\sigma,\tau}$.

This event structure comes with partial maps $\Pi_1 : S \wedge T \rightarrow S$ and $\Pi_2 : S \wedge T \rightarrow T$, analogous to the usual projections of a pullback: for $(x, y) \in S \wedge T$, $\Pi_1(x, y)$ is defined to $s \in S$ whenever the top-element of $\varphi_{x,y}$ is $((0, s), w_2)$ for some $w_2 \in x_* \parallel y$. The map Π_1 is undefined only on events of $S \wedge T$ corresponding to internal events of T (i.e. (x, y) with top element of $\varphi_{x,y}$ of the form $((1, t), (1, t))$). The map Π_2 is defined symmetrically, and undefined on events corresponding to internal events of S . We write $\sigma \wedge \tau$ for $\sigma \circ \Pi_1 = \tau \circ \Pi_2 : S \wedge T \rightarrow A$.

Lemma 1. *Let $\sigma : S \rightarrow A$ and $\tau : T \rightarrow A$ be partial maps. Let $(X, f : X \rightarrow S, g : X \rightarrow T)$ be a triple such that the following outer square commutes:*

$$\begin{array}{ccccc}
 & & X & & \\
 & f \swarrow & \vdots \langle f, g \rangle & \searrow g & \\
 S & \xrightarrow{\Pi_1} & S \wedge T & \xrightarrow{\Pi_2} & T \\
 & \searrow \sigma & \downarrow \sigma \wedge \tau & \swarrow \tau & \\
 & & A & &
 \end{array}$$

If for all $p \in X$ with $f p$ and $g p$ defined, $\sigma(f p) = \tau(g p)$ is defined, then there exists a unique $\langle f, g \rangle : X \rightarrow S \wedge T$ making the two upper triangles commute.

From this closed interaction, we define the open interaction as in [14]. Given two pre-strategies $\sigma : S \rightarrow A^\perp \parallel B$ and $\tau : T \rightarrow B^\perp \parallel C$, their interaction

$$\tau \otimes \sigma : (S \parallel C) \wedge (A \parallel T) \rightarrow A^\perp \parallel C$$

is defined as the composite partial map $(S \parallel C) \wedge (A \parallel T) \rightarrow A \parallel B \parallel C \rightarrow A \parallel C$, where the “pullback” is first computed ignoring polarities – the codomain of the resulting partial map is $A^\perp \parallel C$, once we reinstate polarities.

Weak Bisimulation. To compare uncovered pre-strategies, we cannot use isomorphisms as in [14], since as hinted earlier, $\alpha_A \otimes \sigma$ comprises synchronised events not corresponding to those in σ . To solve this, we introduce weak bisimulation between uncovered strategies:

Definition 5. Let $\sigma : S \rightarrow A$ and $\tau : T \rightarrow A$ be uncovered pre-strategies. A weak bisimulation between σ and τ is a relation $\mathcal{R} \subseteq \mathcal{C}(S) \times \mathcal{C}(T)$ containing (\emptyset, \emptyset) , such that for all $x \mathcal{R} y$, we have:

- If $x \xrightarrow{s} c x'$ such that s is visible, then there exists $y \subseteq^* y' \xrightarrow{t} c y''$ with $\sigma s = \tau t$ and $x' \mathcal{R} y''$ (and the symmetric condition for τ)
- If $x \xrightarrow{s} c x'$ such that s is internal, then there exists $y \subseteq^* y'$ such that $x' \mathcal{R} y'$ (and the symmetric condition for τ)

Two uncovered pre-strategies σ, τ are weakly bisimilar (written $\sigma \simeq \tau$) when there is a weak bisimulation between them.

Associativity of interaction (up to isomorphism, hence up to weak bisimulation) follows directly from Lemma 1. Moreover, it is straightforward to check that weak bisimulation is a congruence (*i.e.* compatible with composition).

Composition of Covered Strategies. From interaction, we can easily define the composition of covered strategies. If $\sigma : S \rightarrow A^\perp \parallel B$ and $\tau : T \rightarrow B^\perp \parallel C$ are covered pre-strategies, their composition (in the sense of [14]) $\tau \odot \sigma$ is defined as $(\tau \otimes \sigma)_\downarrow$. The operation \downarrow is well-behaved with respect to interaction:

Lemma 2. For σ, τ composable pre-strategies, $(\tau \otimes \sigma)_\downarrow \cong \tau_\downarrow \odot \sigma_\downarrow$.

3.3 A Compact-Closed Category of Uncovered Strategies

Although we have a notion of morphism (pre-strategies) between games and an associative composition, we do not have a category up to weak bisimulation yet. Unlike in [14], races in a game may cause copycat on this game to not be idempotent (see [3] for a counterexample), which is necessary for it to be an identity. To ensure that, we restrict ourselves to **race-free** games: those such that whenever a configuration x can be extended by a_1, a_2 of distinct polarities, the union $x \cup \{a_1, a_2\}$ is consistent. From now on, games are assumed race-free.

Lemma 3. For a race-free game A , $\alpha_A \otimes \alpha_A \simeq \alpha_A$.

Proof. It will follow from the forthcoming Lemma 4.

Uncovered Strategies. Finally, we characterise the pre-strategies invariant under composition with copycat. The two ingredients of [5, 14], receptivity and courtesy (called *innocence* in [14]) are needed, but this is not enough: we need another condition as witnessed by the following example.

Consider the strategy $\sigma : \oplus_1 \rightsquigarrow \oplus_2$ on the game $A = \oplus_1 \oplus_2$ playing non-deterministically one of the two moves. Then the interaction $\alpha_A \otimes \sigma$ is:

$$\begin{array}{ccc} A^* & & A \\ *_1 & \longrightarrow & \triangleright \oplus_1 \\ \} & & \\ *_2 & \longrightarrow & \triangleright \oplus_2 \end{array}$$

It is not weakly bisimilar to σ : $\alpha_A \otimes \sigma$ can do $*_1$, an internal transition, to which σ can only respond by not doing anything. Then σ can still do \oplus_1 and \oplus_2 whereas $\alpha_A \otimes \sigma$ cannot: it is committed to doing \oplus_1 . To solve this problem, we need to force strategies to decide their nondeterministic choices *secretly*, by means of internal events – so σ will not be a valid uncovered strategy, but $\alpha_A \otimes \sigma$ will. Indeed, $\alpha_A \otimes (\alpha_A \otimes \sigma)$ below is indeed weakly bisimilar to $\alpha_A \otimes \sigma$.

$$\begin{array}{ccccc} A^* & & A^* & & A \\ *_1 & \longrightarrow & \triangleright *_1 & \longrightarrow & \triangleright \oplus_1 \\ \} & & & & \\ *_2 & \longrightarrow & \triangleright *_2 & \longrightarrow & \triangleright \oplus_2 \end{array}$$

Definition 6. An (uncovered) strategy is a pre-strategy $\sigma : S \rightarrow A$ satisfying:

- receptivity: if $x \in \mathcal{C}(S)$ is such that $\sigma x \xrightarrow{a} \text{C}$ with $a \in A$ negative, then there exists a unique $x \xrightarrow{s} \text{C}$ with $\sigma s = a$.
- courtesy: if $s \rightarrow s'$ and s is positive or s' is negative, then $\sigma s \rightarrow \sigma s'$.
- secrecy: if $x \in \mathcal{C}(S)$ extends with s_1, s_2 but $x \cup \{s_1, s_2\} \notin \mathcal{C}(S)$, then s_1 and s_2 are either both negative, or both internal.

Receptivity and courtesy are stated exactly as in [14]. As a result, hiding the internal events of an uncovered strategy yields a strategy σ_{\downarrow} in the sense of [14].

For any game A , α_A is an uncovered strategy: it satisfies secrecy as its only minimal conflicts are inherited from the game and are between negative events.

The Category \mathbf{CG}_{\otimes} . Our definition of uncovered strategy does imply that copycat is neutral for composition.

Lemma 4. Let $\sigma : S \rightarrow A$ be an uncovered strategy. Then $\alpha_A \otimes \sigma \simeq \sigma$.

The result follows immediately:

Theorem 1. Race-free games and uncovered strategies up to weak bisimulation form a compact-closed category \mathbf{CG}_{\otimes} .

3.4 Interpretation of Affine Nondeterministic PCF

From now on, strategies are by default considered uncovered. We sketch the interpretation of \mathbf{APCF}_+ inside \mathbf{CG}_{\otimes} . As a compact-closed category, \mathbf{CG}_{\otimes} supports an interpretation of the linear λ -calculus. However, the empty game 1 is not terminal, as there are no natural transformation $\epsilon_A : A \rightarrow 1$ in \mathbf{CG}_{\otimes} .

The negative category \mathbf{CG}_{\otimes}^- . We solve this issue as in [4], by looking at negative strategies and negative games.

Definition 7. An event structure with partial polarities is **negative** when all its minimal events are negative.

A strategy $\sigma : S \multimap A$ is negative when S is. Copycat on a negative game is negative, and negative strategies are stable under composition:

Lemma 5. There is a subcategory \mathbf{CG}_{\otimes}^- of \mathbf{CG}_{\otimes} consisting in negative race-free games and negative strategies. It inherits a monoidal structure from \mathbf{CG} in which the unit (the empty game) is terminal.

Moreover, \mathbf{CG}_{\otimes}^- has products. The **product** $A \& B$ of two games A and B , has events, causality, polarities as for $A \parallel B$, but consistent sets restricted to those of the form $\{0\} \times X$ or $\{1\} \times X$ with X consistent in A or B . The **projections** are $\varpi_A : \mathbb{C}_A \rightarrow (A \& B)^\perp \parallel A$, and $\varpi_B : \mathbb{C}_B \rightarrow (A \& B)^\perp \parallel B$.

Finally, the **pairing** of negative strategies $\sigma : S \multimap A^\perp \parallel B$ and $\tau : T \multimap A^\perp \parallel C$ is the obvious map $\langle \sigma, \tau \rangle : S \& T \multimap A^\perp \parallel B \& C$, and the laws for the cartesian product are direct verifications.

We also need a construction to interpret the function space. However, for A and B negative, $A^\perp \parallel B$ is not usually negative. To circumvent this, we introduce a negative variant $A \multimap B$, the linear arrow. To simplify the presentation, we only define it in a special case. A game is **well-opened** when it has at most one initial event. When B is well-opened, we define $A \multimap B$ to be 1 if $B = 1$; and otherwise $A^\perp \parallel B$ with the exception that every move in A depends on the single minimal move in B . As a result \multimap preserves negativity. We get:

Lemma 6. If B is well-opened, $A \multimap B$ is well-opened and is an exponential object of A and B .

In other words, well-opened games are an exponential ideal in \mathbf{CG}_{\otimes}^- . We interpret types of \mathbf{APCF}_+ inside well-opened games of \mathbf{CG}_{\otimes}^- :

$$\llbracket \mathbf{com} \rrbracket = \begin{array}{c} \mathbf{run}^- \\ \downarrow \\ \mathbf{done}^+ \end{array} \quad \llbracket \mathbb{B} \rrbracket = \begin{array}{c} q^- \\ \swarrow \quad \searrow \\ \mathbf{tt}^+ \quad \sim \quad \mathbf{ff}^+ \end{array} \quad \llbracket A \multimap B \rrbracket = \llbracket A \rrbracket \multimap \llbracket B \rrbracket$$

Interpretation of Terms. Interpretation of the affine λ -calculus in \mathbf{CG}_{\otimes}^- follows standard methods. First, the primitives $\mathbf{tt}, \mathbf{ff}, \perp, \mathbf{if}$ are interpreted as:

$$\begin{array}{llll} \llbracket \mathbf{tt} \rrbracket : \mathbb{B} & \llbracket \mathbf{ff} \rrbracket : \mathbb{B} & \llbracket \perp \rrbracket : \mathbb{B} & \llbracket \mathbf{if} \rrbracket : \mathbb{B} \multimap (\mathbb{B} \& \mathbb{B}) \multimap \mathbb{B} \\ \begin{array}{c} q \\ \downarrow \\ \mathbf{tt} \end{array} & \begin{array}{c} q \\ \downarrow \\ \mathbf{ff} \end{array} & \begin{array}{c} q \\ \downarrow * \\ \downarrow * \\ \downarrow * \\ \vdots \end{array} & \begin{array}{c} \text{Diagram showing the interpretation of } \mathbf{if} \text{ as a game tree. A root node } q \text{ branches into } q \text{ and } \mathbf{ff}. \text{ The } q \text{ node branches into } q \text{ and } \mathbf{tt}. \text{ The } \mathbf{ff} \text{ node branches into } q \text{ and } \mathbf{ff}. \text{ The } q \text{ nodes branch into } b \text{ and } b'. \text{ Dashed lines indicate dependencies.} \end{array} \end{array}$$

A non-standard point is the interpretation of \perp : usually interpreted in game semantics by the minimal strategy simply playing q (as will be done in the next section), our interpretation here reflects the fact that \perp represents an infinite computation that never returns. Conditionals are implemented as usual:

$$\llbracket \text{if } M N N' \rrbracket_{\otimes} = \text{if} \otimes (\llbracket M \rrbracket_{\otimes} \parallel \langle \llbracket N \rrbracket_{\otimes}, \llbracket N' \rrbracket_{\otimes} \rangle).$$

Soundness and Adequacy. We now prove adequacy for various notions of convergence. First, we build an uncovered strategy from the operational semantics.

Definition 8 (The operational tree). *Let M be a closed term of type \mathbb{B} . We define the pre-strategy $\mathfrak{t}(M)$ on \mathbb{B} as follows:*

Events: *An initial event \perp plus one event per derivation $M \rightarrow^* M'$.*

Causality: *\perp is below other events, and derivations are ordered by prefix*

Consistency: *A set of events is consistent when its events are comparable.*

Labelling: *\perp has label q , a derivation $M \rightarrow^* b$ where $b \in \{\mathfrak{tt}, \mathfrak{ff}\}$ is labelled by b . Other derivations are internal.*

As a result, $\mathfrak{t}(M)$ is a tree. Our main result of adequacy can now be stated:

Theorem 2. *For a term $\vdash M : \mathbb{B}$, $\mathfrak{t}(\text{if } M \mathfrak{tt} \mathfrak{ff})$ and $\llbracket M \rrbracket_{\otimes}$ are weakly bisimilar.*

We need to consider $\mathfrak{t}(\text{if } M \mathfrak{tt} \mathfrak{ff})$ and not simply $\mathfrak{t}(M)$ to ensure secrecy. From this theorem, adequacy results for may and fair convergences arise:

Corollary 1. *For any term $\vdash M : \mathbb{B}$, we have:*

May: *$M \rightarrow^* \mathfrak{tt}$ if and only if $\llbracket M \rrbracket_{\otimes}$ contains a positive move*

Fair: *For all $M \rightarrow^* M'$, M' can converge, if and only if all finite configurations of $\llbracket M' \rrbracket_{\otimes}$ can be extended to contain a positive move.*

However, we cannot conclude adequacy for must equivalence from Theorem 2. Indeed, must convergence is not generally stable under weak bisimilarity: for instance, (the strategies representing) \mathfrak{tt} and $Y(\lambda x. \text{if choice } \mathfrak{tt} x)$ are weakly bisimilar but the latter is not must convergent. To address this in the next section we will refine the interpretation to obtain a closer connection with syntax.

4 Essential Events

The model presented in the previous section is very operational; configurations of $\llbracket M \rrbracket_{\otimes}$ can be seen as derivations for an operational semantics. The price, however, is that besides the fact that the interpretation grows dramatically in size, we can only get a category up to weak bisimulation, which can be too coarse (for instance for must convergence). We would like to remove all events that are not relevant to the behaviour of terms up to weak bisimulation. In other words, we want a notion of *essential internal events* that (1) suffices to recover all behaviour with respect to weak bisimulation, but which (2) is not an obstacle to getting a category up to isomorphism (which amounts to $\alpha_A \circ \sigma \cong \sigma$).

4.1 Definition of Essential Events

As shown before, the loss of behaviours when hiding is due to the disappearance of events participating in a conflict. A neutral event may not have visible consequences but still be relevant if in a minimal conflict; such events are *essential*.

Definition 9. *Let $\sigma : S \rightarrow A$ be an uncovered pre-strategy. An **essential event** of S is an event s which is either visible, or (internal and) involved in a minimal conflict (that is such that we have $s \sim_x s'$ for some s', x).*

Write E_S for the set of essential events of σ . Any pre-strategy $\sigma : S \rightarrow A$ induces another pre-strategy $\mathcal{E}(\sigma) : \mathcal{E}(S) = S \downarrow E_S \rightarrow A$ called **the essential part** of σ . The following proves that our definition satisfies (1): no behaviour is lost.

Lemma 7. *An uncovered pre-strategy $\sigma : S \rightarrow A$ is weakly bisimilar to $\mathcal{E}(\sigma)$.*

This induces a new notion of (associative) composition only keeping the essential events. For $\sigma : A^\perp \parallel B$ and $\tau : B^\perp \parallel C$, let $\tau \odot \sigma = \mathcal{E}(\tau \otimes \sigma)$. We observe that $\mathcal{E}(\tau \otimes \sigma) \cong \mathcal{E}(\tau) \odot \mathcal{E}(\sigma)$.

Which pre-strategies compose well with copycat with this new composition?

4.2 Essential Strategies

We now can state property (2): the events added by composition with copycat are inessential, hence hidden during composition:

Theorem 3. *Let $\sigma : S \rightarrow A$ be an uncovered strategy. Then $\alpha_A \odot \sigma \cong \mathcal{E}(\sigma)$.*

This prompts the following definition. An uncovered pre-strategy σ is **essential** when it is a strategy, and if, equivalently: (1) all its events are essential, (2) $\sigma \cong \mathcal{E}(\sigma)$. We obtain a characterisation of strategies in the spirit of [14]:

Theorem 4. *A pre-strategy $\sigma : S \rightarrow A$ is essential if and only if $\alpha_A \odot \sigma \cong \sigma$.*

As a result, we get:

Theorem 5. *Race-free games, and essential strategies up to isomorphism form a compact-closed category \mathbf{CG}_\odot .*

Relationship Between \mathbf{CG} and \mathbf{CG}_\odot . Covered strategies can be made into a compact-closed category [5, 14]. Remember that the composition of $\sigma : S \rightarrow A^\perp \parallel B$ and $\tau : T \rightarrow B^\perp \parallel C$ in \mathbf{CG} is defined as $\tau \odot \sigma = (\tau \otimes \sigma)_\downarrow$.

Lemma 8. *The operation $\sigma \mapsto \sigma_\downarrow$ extends to an identity-on-object functor $\mathbf{CG}_\odot \rightarrow \mathbf{CG}$.*

In the other direction, a strategy $\sigma : A$ might not be an essential strategy; in fact it might not even be an uncovered strategy, as it may fail secrecy. Sending σ to $\alpha_A \odot \sigma$ delegates the non-deterministic choices to internal events and yields an essential strategy, but this operation is not functorial.

Relationship Between \mathbf{CG}_{\odot} and \mathbf{CG}_{\otimes} . The forgetful operation mapping an essential strategy σ to itself, seen as an uncovered strategy, defines a functor $\mathbf{CG}_{\odot} \rightarrow \mathbf{CG}_{\otimes}$. Indeed, if two essential strategies are isomorphic, they are also weakly bisimilar. Moreover, we have that $\tau \otimes \sigma \simeq \mathcal{E}(\tau \otimes \sigma) = \tau \odot \sigma$. However the operation $\mathcal{E}(\cdot)$ does not extend to a functor in the other direction even though $\mathcal{E}(\tau) \odot \mathcal{E}(\sigma) \cong \mathcal{E}(\tau \otimes \sigma)$, as it is defined only on concrete representatives, not on equivalence classes for weak bisimilarity.

4.3 Interpretation of \mathbf{APCF}_+

We now show that this new category also supports a sound and adequate interpretation of \mathbf{APCF}_+ for various testing equivalences, including must. As before, we need to construct the category of negative games and strategies.

Lemma 9. *There is a cartesian symmetric monoidal category \mathbf{CG}_{\odot}^- of negative race-free games and negative essential strategies up to isomorphism. Well-opened negative race-free games form an exponential ideal of \mathbf{CG}_{\odot}^- .*

We keep the same interpretation of types of affine nondeterministic PCF. Moreover, the strategy **if** is essential. As a result, we let:

$$\llbracket \perp \rrbracket_{\odot} = q : \mathbb{B} \quad \llbracket \mathbf{if} \ M \ N \ N' \rrbracket_{\odot} = \mathbf{if} \odot (\llbracket M \rrbracket_{\odot} \parallel \langle \llbracket N \rrbracket_{\odot}, \llbracket N' \rrbracket_{\odot} \rangle)$$

Using $\mathcal{E}(\sigma \otimes \tau) = \mathcal{E}(\sigma) \odot \mathcal{E}(\tau)$, one can prove by induction that for any term M we have $\llbracket M \rrbracket_{\odot} = \mathcal{E}(\llbracket M \rrbracket_{\otimes})$. Furthermore, this interpretation permits a stronger link between the operational and the denotational semantics:

Theorem 6. *For all terms $\vdash M : \mathbb{B}$, $\mathcal{E}(\mathbf{t}(M)) \cong \llbracket M \rrbracket_{\odot}$.*

Theorem 6 implies Theorem 2. It also implies adequacy for must:

Corollary 2. *The interpretation $\llbracket \cdot \rrbracket_{\odot}$ is adequate for may, and fair, and must: $\vdash M : \mathbb{B}$ has no infinite derivations if and only if all (possibly infinite) maximal configurations of $\llbracket M \rrbracket_{\odot}$ have a positive event.*

This result also implies that $\llbracket \cdot \rrbracket_{\otimes}$ is adequate for must.

5 Conclusion

We have described an extension of the games of [14] to uncovered strategies, composed without hiding. It has strong connections with operational semantics, as the interpretations of terms of base type match their tree of reductions. It also forms a compact-closed category up to weak bisimulation, and is adequate for the denotational semantics of programming languages. Identifying the *inessential* events as those responsible for the non-neutrality of copycat, we remove them to yield a compact closed category up to isomorphism. Doing so we obtain our sought-after setting for the denotational semantics of programming languages, one *agnostic w.r.t.* the chosen testing equivalence. The work blends well with the technology of [7] (symmetry, concurrent innocence) dealing with non-affine languages and characterising strategies corresponding to pure programs; these developments appear in the first author's PhD thesis [3].

Acknowledgements. We gratefully acknowledge the support of the ERC Advanced Grant ECSYM, EPSRC grants EP/K034413/1 and EP/K011715/1, and LABEX MILYON (ANR-10-LABX-0070) of Université de Lyon, within the program “Investissements d’Avenir” (ANR-11-IDEX-0007) operated by the ANR.

References

1. Abramsky, S.: Game semantics for programming languages. In: Prívvara, I., Ružička, P. (eds.) MFCS 1997. LNCS, vol. 1295, pp. 1–4. Springer, Heidelberg (1997). <https://doi.org/10.1007/BFb0029944>
2. Brinksma, E., Rensink, A., Vogler, W.: Fair testing. In: Lee, I., Smolka, S.A. (eds.) CONCUR 1995. LNCS, vol. 962, pp. 313–327. Springer, Heidelberg (1995). https://doi.org/10.1007/3-540-60218-6_23
3. Castellan, S.: Concurrent structures in game semantics. Ph.D. thesis, ENS Lyon, France (2017)
4. Castellan, S., Clairambault, P.: Causality vs. interleaving in game semantics. In: CONCUR 2016 - Concurrency Theory (2016)
5. Castellan, S., Clairambault, P., Rideau, S., Winskel, G.: Games and strategies as event structures. *Log. Methods Comput. Sci.* **13** (2017)
6. Castellan, S., Clairambault, P., Winskel, G.: Symmetry in concurrent games. In: Henzinger, T.A., Miller, D. (eds.) CSL-LICS 2014, Vienna, Austria, July 14–18, 2014, p. 28. ACM (2014)
7. Castellan, S., Clairambault, P., Winskel, G.: The parallel intensionally fully abstract games model of PCF. In: LICS 2015. IEEE Computer Society (2015)
8. Castellan, S., Hayman, J., Lasson, M., Winskel, G.: Strategies as concurrent processes. *Electr. Notes Theor. Comput. Sci.* **308**, 87–107 (2014)
9. Danos, V., Harmer, R.: Probabilistic game semantics. In: 15th Annual IEEE Symposium on Logic in Computer Science, Santa Barbara, California, USA, 26–29 June 2000, pp. 204–213 (2000)
10. Ghica, D.R., Murawski, A.S.: Angelic semantics of fine-grained concurrency. *Ann. Pure Appl. Log.* **151**(2–3), 89–114 (2008)
11. Harmer, R., McCusker, G.: A fully abstract game semantics for finite nondeterminism. In: 14th Annual IEEE Symposium on Logic in Computer Science, Trento, Italy, 2–5 July 1999, pp. 422–430 (1999)
12. Hirschowitz, T.: Full abstraction for fair testing in CCS. In: Heckel, R., Milius, S. (eds.) CALCO 2013. LNCS, vol. 8089, pp. 175–190. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40206-7_14
13. Laird, J.: A game semantics of idealized CSP. *Electr. Notes Theor. Comput. Sci.* **45**, 232–257 (2001)
14. Rideau, S., Winskel, G.: Concurrent strategies. In: LICS, pp. 409–418. IEEE Computer Society (2011)
15. Tsukada, T., Luke Ong, C.-H.: Innocent strategies are sheaves over plays - deterministic, non-deterministic and probabilistic innocence. CoRR, abs/1409.2764 (2014)
16. Tsukada, T., Luke Ong, C.-H.: Nondeterminism in game semantics via sheaves. In: 30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2015, Kyoto, Japan, 6–10 July 2015, pp. 220–231 (2015)

17. Winskel, G.: Event structures. In: Brauer, W., Reisig, W., Rozenberg, G. (eds.) ACPN 1986. LNCS, vol. 255, pp. 325–392. Springer, Heidelberg (1987). https://doi.org/10.1007/3-540-17906-2_31
18. Yamada, N., Abramsky, S.: Dynamic games and strategies. CoRR, abs/1601.04147 (2016)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

