

A Universal Innocent Game Model for the Böhm Tree Lambda Theory

Andrew D. Ker, Hanno Nickau, and C.-H. Luke Ong

Computing Laboratory, Parks Road, Oxford OX1 3QD, UK,
{Andrew.Ker, Hanno.Nickau, Luke.Ong}@comlab.ox.ac.uk.

Abstract. We present a game model of the untyped λ -calculus, with equational theory equal to the Böhm tree λ -theory \mathcal{B} , which is universal (i.e. every element of the model is definable by some term). This answers a question of Di Gianantonio, Franco and Honsell. We build on our earlier work, which uses the methods of innocent game semantics to develop a universal model inducing the maximal consistent sensible theory \mathcal{H}^* . To our knowledge these are the first syntax-independent universal models of the untyped λ -calculus.

1 Introduction

We aim to construct a *universal* model (i.e. every element of the model is the denotation of some term) of the pure untyped λ -calculus which induces the Böhm tree λ -theory \mathcal{B} , by building on the game models presented in [4]. Although the general approach is innocent in the sense of [3] and [7], the two-player games we use are simpler and can be considered a special case where moves are neither questions nor answers but simply “declarations”. A notable feature of game semantics is that the λ -definable strategies are effective methods for copying moves uniformly from one “component” of the game to another. For example, the identity strategy on an arena $A \Rightarrow A$ is *everywhere copycat* i.e. P always plays back every O-move (but in the opposite component of A). The key idea is that the innocent strategies definable by untyped λ -terms are, what we call, *effectively almost-everywhere copycat* (EAC). Informally this means that at every position, except in response to finitely many possible O-moves, the strategy is constrained to behave, from that point onwards, uniformly in an everywhere-copycat fashion, just like the identity strategy. Effectively here means that (not only is the strategy itself recursive but also) at every position, the boundary of that finite part of the game tree in which the strategy is *not* forced to play copycat must be computable.

We find it convenient to introduce innocent strategies in a concrete setting whereby (tree) arenas are defined as subsets of \mathbb{N}^* of a certain kind, and this we do in Sect. 2. Section 3 introduces the EAC strategies which give rise to a universal λ -model \mathcal{D}_{EAC} whose theory is the maximal consistent sensible λ -theory \mathcal{H}^* . The definition of such strategies uses an efficient encoding of innocent strategies, as partial functions from \mathbb{N}^* to tuples of numbers, which we call

economical form. Sections 2 and 3 should be regarded as a survey of [4], and this paper is a sequel to that work. The notion of EAC strategies has a natural extension to *explicitly and effectively almost-everywhere copycat*. However finding an ambient cartesian closed category for these strategies to inhabit proved to be a painful process as we briefly show in Sect. 4 – the natural analogues fail to work quite as intended. Once this has been overcome we use a reflexive object to describe a λ -algebra which we call \mathcal{M} . We formulate a new version of the powerful *Exact Correspondence Theorem* of the earlier work, with which we can show that \mathcal{M} is a universal λ -model which induces the intended equational theory. To our knowledge, \mathcal{D}_{EAC} and \mathcal{M} are the first syntax-independent universal λ -models.

In [2], Di Gianantonio et al. have obtained game models of the untyped λ -calculus using history-free strategies. They show that all their models induce the same λ -theory \mathcal{H}^* and have asked for “new techniques for overcoming this apparent rigidity of game λ -models”. This paper answers that question by constructing a universal game model for the Böhm tree *lambda*-theory.

2 Arenas and Innocent Strategies

This section and the next give a quick introduction to the basic ideas underpinning the main result of the paper. We refer the reader to [4] for further details and to [3] and [5] for proofs of all results quoted. We define an *arena* to be a finite tuple of nonempty trees of *moves*. The root of each tree is called an *initial move*. Our trees are considered “upside-down” with the root at the top, rather like family trees. We can also refer to the child of a node, and say that one node inherits from another, in the same vein. We say that moves at an even depth of the trees (including the roots at depth 0) are *O-moves*, and moves at an odd depth are *P-moves*. O-moves are often denoted by \bullet and P-moves by \circ .

We will only be interested in countably branching, countably deep trees. Thus we can encode each tree of the arena as a subset of \mathbb{N}^{*1} by inductively labelling the root as ε and the n^{th} child of the move s as $s \cdot n$ (we use the notation s, t etc. to denote sequences). Hence each move of each tree is associated uniquely with a sequence of natural numbers. Conversely, given any subset $A \subseteq \mathbb{N}^*$ which is prefix-closed and has the property that whenever $s \cdot n \in A$ we have $s \cdot m \in A$ for each $m \leq n$, we can form an arena of one tree where the moves are the elements of A , ordered by prefix. *Henceforth by arenas, we shall always mean arenas in sequence-subset (of \mathbb{N}^*) form.* For example, the empty sequence $\langle \rangle$ is an arena, which we call the *empty arena*; $\langle \{\varepsilon\} \rangle$ is the minimal one-tree arena consisting of a root node; the maximal one-tree arena, consisting of an infinitely deep, infinitely branching tree, is $\langle \mathbb{N}^* \rangle$. As the empty arena, the minimal and maximal one-tree arenas are important, we shall name them E, M and U respectively.

There are two major constructions for forming arenas. Suppose $A = \langle A_1, \dots, A_m \rangle$ and $B = \langle B_1, \dots, B_n \rangle$ are arenas.

- The *product arena* $A \times B$ is the “disjoint union” of the trees of A and B , the concatenation of their tuples. Formally $A \times B = \langle A_1, \dots, A_m, B_1, \dots, B_n \rangle$.

¹ We do not include 0 in the set \mathbb{N} , and write \mathbb{N}_0 for $\mathbb{N} \cup \{0\}$.

- The *function space arena* $A \Rightarrow B$ is constructed as follows: the initial moves of $A \Rightarrow B$ are those of B ; and to the tree below each such initial move, we graft onto it a copy of A . More precisely $A \Rightarrow B = \langle C_1, \dots, C_n \rangle$ where

$$C_i = \{ \varepsilon \} \cup \{ a \cdot \mathbf{s} \mid 1 \leq a \leq m \wedge \mathbf{s} \in A_a \} \cup \{ (a+m) \cdot \mathbf{s} \mid a \cdot \mathbf{s} \in B_i \}.$$

The reader may wish to check that $M \Rightarrow M = \langle \{ \varepsilon, 1 \} \rangle$, and that $U \Rightarrow U$ and U are equal as arenas.

A *justified sequence* of an arena A is a sequence of moves of which each element except the first, which must be an initial move, is equipped with a pointer to some previous move. We call the pointer a *justification pointer* and if the move m^- is pointed to by m we say that m^- *justifies* m . We say that a move m^- in a justified sequence *hereditarily justifies* m if one can reach m^- from m by repeatedly following justification pointers. A justified sequence \mathbf{s} is said to be *well-formed* if elements of \mathbf{s} alternate between P-moves and O-moves and if $m \in \mathbf{s}$ is justified by m^- then the move m is directly beneath m^- in the tree of the arena². Henceforth all justified sequences are assumed to be well-formed.

The *P-view* of a justified sequence \mathbf{s} , written $\lceil \mathbf{s} \rceil$, is given recursively by:

$$\begin{aligned} \lceil \varepsilon \rceil &= \varepsilon && \text{for initial moves } \varepsilon \\ \lceil \mathbf{s} \cdot m \rceil &= \lceil \mathbf{s} \rceil \cdot m && \text{for } m \text{ a P-move} \\ \lceil \mathbf{s} \cdot m^- \cdot \mathbf{t} \cdot m \rceil &= \lceil \mathbf{s} \rceil \cdot m^- \cdot m && \text{for } m \text{ an O-move justified by } m^- \end{aligned}$$

The definition of *O-view*, $\lfloor \mathbf{s} \rfloor$, is given analogously.

A *legal position* of an arena A is a well-formed justified sequence \mathbf{s} satisfying the *visibility condition*: for each non-initial P-move m justified by m^- , say $\mathbf{s} = \mathbf{t}_1 \cdot m^- \cdot \mathbf{t}_2 \cdot m \cdot \mathbf{t}_3$, we have that $m^- \in \lceil \mathbf{t}_1 \cdot m^- \cdot \mathbf{t}_2 \cdot m \rceil$. Similarly all O-moves are justified by P-moves appearing in the O-view up to that point. Then if \mathbf{s} is a legal position then so are $\lceil \mathbf{s} \rceil$ and $\lfloor \mathbf{s} \rfloor$. By a *P-view* of an arena A , we shall mean a justified sequence which is the P-view of some legal position of A .

Lemma 1 (View Characterisation). *A justified sequence of an arena A is a P-view if and only if it is well-formed and every non-initial O-move is justified by the immediately preceding P-move.*

Within arenas there are games played out between P and O. A *P-strategy* σ for a single-tree arena A consists of a prefix-closed subset of legal positions of A which is deterministic (if $\mathbf{s} \cdot m \in \sigma$ and $\mathbf{s} \cdot m' \in \sigma$ for P-moves m and m' then $m = m'$) and such that if $\mathbf{s} \cdot m \in \sigma$ for a P-move m and $\mathbf{s} \cdot m \cdot m'$ is a legal position of A then $\mathbf{s} \cdot m \cdot m' \in \sigma$. An *O-strategy* is defined analogously. However we are more often interested in P-strategies which we will usually just refer to as strategies. For a general arena $A = \langle A_1, \dots, A_n \rangle$ a P-strategy is an n -tuple of P-strategies, one for each tree. In contrast, an O-strategy is just a single O-strategy on one of the trees, together with information which selects that tree.

² The no-dangling-question-mark condition in [3, 7] (equivalently the well-bracketing condition) is redundant for our arenas.

If we have strategies σ and τ on arenas $A \Rightarrow B$ and $B \Rightarrow C$ respectively then we can form their composite strategy $\sigma; \tau$ on $A \Rightarrow C$. Informally we do this by identifying O/P-moves of the B component of $A \Rightarrow B$ with P/O-moves of the B component of $B \Rightarrow C$, and then hiding all the moves in B . This is reminiscent of CSP’s “parallel composition” and “hiding” operators; see [4] for a formal definition. Similar ideas extend to arenas of multiple trees. An essentially straightforward result, although tedious in proof, is that composition is well-defined and associative. We will only be interested in strategies with a property called *innocence*.

A P-strategy σ is *innocent* if for odd-length legal positions \mathbf{s} and \mathbf{t} and P-moves m , $\mathbf{s} \cdot m \in \sigma \wedge \mathbf{t} \in \sigma \wedge \lceil \mathbf{s} \rceil = \lceil \mathbf{t} \rceil \rightarrow \mathbf{t} \cdot m \in \sigma$ and the moves m are justified by moves which are identical in the P-view $\lceil \mathbf{s} \rceil = \lceil \mathbf{t} \rceil$. i.e. P’s next move, and its justification, at each stage depends only on the P-view up to that point. An important fact is that composition of innocent strategies is well-defined (for a proof see [3, §5.3]): if σ is an innocent strategy on $A \Rightarrow B$ and τ an innocent strategy on $B \Rightarrow C$ then $\sigma; \tau$ is an innocent strategy on $A \Rightarrow C$.

The property of innocence means that such a strategy is determined by a partial function from odd-length P-views to *justified P-moves* i.e. P-moves equipped with a justification pointer back into the P-view. In fact, given an innocent strategy σ , we can define a canonical such function, which we write f_σ , that defines it. A function constructed in this way is called *innocent* and we can formalise such functions. We say that f is an *innocent function* if f is a partial function from odd-length P-views in A to justified P-moves of A such that $\text{dom}(f)$ is closed under odd-length prefix, and if $\mathbf{s} \cdot m \cdot m' \in \text{dom}(f)$ then $f(\mathbf{s}) = m$. We note that such a function can only encode an innocent strategy. The conditions given are required to make the function “strategic”, i.e. the set of legal positions it describes are prefix-closed, deterministic and made up of properly justified sequences of moves. These conditions are sufficient to allow us to define the reverse construction of a unique strategy σ_f from an innocent function f such that the construction is invertible (i.e. $f_{\sigma_f} = f$ and $\sigma_{f_\sigma} = \sigma$) and it preserves and reflects inclusion (i.e. $f \subseteq f' \iff \sigma_f \subseteq \sigma_{f'}$). Thus we can identify the representation by innocent function and that by subset of legal positions.

An innocent strategy is said to be *compact* if the graph of its innocent function is finite (i.e. is defined on finitely many P-views). It is said to be *recursive* if the innocent function representing it is recursive. It is easy to see that the composition of two recursive innocent strategies is itself recursive.

Definition 1. *Objects of the Category of Arenas and Innocent Strategies, \mathbb{A} , are arenas (in sequence-subset form); morphisms $f : A \rightarrow B$ are innocent strategies on the function space arena $A \Rightarrow B$. Composition of morphisms is composition as strategies. The Category of Arenas and Recursive Innocent Strategies, \mathbb{A}_{REC} , has recursive arenas as objects and recursive innocent strategies as morphisms.*

Theorem 1. *\mathbb{A} and \mathbb{A}_{REC} are both cartesian closed.*

The terminal object $\mathbf{1}$ of both \mathbb{A} and \mathbb{A}_{REC} is the empty arena E , and the categorical constructions of product and function space are exactly the respective

arena constructs. The category \mathbb{A} is enriched over dI-domains. One cannot say the same of \mathbb{A}_{REC} , because the computable partial functions do not form a cpo. For example, one can “approximate” the Halting Problem by computable functions.

Scott has observed that every λ -algebra arises from a reflexive object R in some cartesian closed category \mathbb{C} ; that is, there exist morphisms $Fun : R \rightarrow [R \Rightarrow R]$ and $Gr : [R \Rightarrow R] \rightarrow R$ such that $Gr ; Fun = \text{id}_{[R \Rightarrow R]}$. Thus we may specify a λ -algebra by a 4-tuple $\langle \mathbb{C}, R, Fun, Gr \rangle$ (it is in fact a $\lambda\eta$ -algebra if $Fun ; Gr = \text{id}_R$); the underlying set of the λ -algebra is the set $\mathbb{C}(\mathbf{1}, R)$ of global sections. If the reflexive object R has enough points (i.e. $\forall f, g : R \rightarrow R. [\forall r : \mathbf{1} \rightarrow R. r ; f = r ; g] \rightarrow f = g$) then $\langle \mathbb{C}, R, Fun, Gr \rangle$ is a λ -model (i.e. a weakly extensional λ -algebra). We refer the reader to [1] for a comprehensive treatment of the model theory of the untyped λ -calculus.

Recall that the arena U has the key property that $U = U \Rightarrow U$ so that in this case the morphisms Fun and Gr are both the identity on U . We can now define the first two of our game λ -algebras (which are both $\lambda\eta$ -algebras): $\langle \mathbb{A}, U, \text{id}_U, \text{id}_U \rangle$ which we shall write simply as \mathcal{D} , and $\langle \mathbb{A}_{\text{REC}}, U, \text{id}_U, \text{id}_U \rangle$ which we shall write as \mathcal{D}_{REC} . By abuse of notation, we shall use \mathcal{D} and \mathcal{D}_{REC} to denote the respective underlying sets. Clearly $\mathcal{D}_{\text{REC}} \subset \mathcal{D}$. By a method of approximation we can show that both the $\lambda\eta$ -algebras are *sensible* i.e. all unsolvable λ -terms have the same denotation which in this case is given by the everywhere undefined innocent function.

3 Effectively Almost-Everywhere Copycat Strategies

There are three properties that allow for a more compact representation of an innocent strategy:

- (i) Each non-initial O-move in any P-view must be a child of the previous move, and the initial move must be ε .
- (ii) Given only the O-moves of a P-view and the value of the innocent function on strictly shorter P-views we can reconstruct the original P-view entirely.
- (iii) The P-move to which this P-view is mapped must be a child of the move justifying it.

In view of these redundancies, we encode innocent strategies σ , over any single-tree arena, as (partial) maps from \mathbb{N}^* to $\mathbb{N} \times \mathbb{N}_0$ (where $\mathbb{N}_0 = \{0, 1, 2, \dots\}$). We call this encoding the *economical form* of σ and sometimes write it e_σ (quite often we abuse notation and write it f_σ too). It is defined as follows:


$$e_\sigma : \langle v_1, \dots, v_n \rangle \mapsto (i, p) \text{ if and only if}$$

$$f_\sigma : \begin{array}{ccccccccccc} \bullet & \circ & \bullet & \circ & \bullet & \dots & \circ & \bullet & \dots & \circ & \bullet \\ \varepsilon & s_1 & s_1 v_1 & s_2 & s_2 v_2 & \dots & s_{n-p} & s_{n-p} v_{n-p} & s_n & s_n v_n & s_{n-p}(v_{n-p} i) \end{array}$$

Justification pointers in the P-view can be deduced from the behaviour of f_σ on shorter P-views, and so have been omitted. Note that each s_i is a sequence of natural numbers.

Furthermore, we can expand any partial function $f : \mathbb{N}^* \rightarrow \mathbb{N} \times \mathbb{N}_0$ which has prefix-closed domain and satisfies $f(\mathbf{v}) = (i, p) \rightarrow 0 \leq p \leq |\mathbf{v}|$ into an innocent strategy on U . Depending on the function, we might not need the whole of U to contain the strategy. We could extend this idea for multiple-tree arenas, but since we will not use it except on the arena U there is no need to do so.

Example 1. The following is the innocent function of the “copycat” strategy id_U :



Here \mathbf{s} range over sequences of appropriate parity, a and b over positive natural numbers. The reader is invited to check that the economical form of this strategy is given by: $\varepsilon \mapsto (1, 0)$, $i \mapsto (i + 1, 1)$ and for nonempty sequences \mathbf{v} , $\mathbf{v}i \mapsto (i, 1)$.

A principle of the λ -calculus is that a term can be applied successively to any other term. So the term $\lambda x.x$ (say) is really more like “ $\lambda x z_0 z_1 z_2 \dots \bullet x z_0 z_1 z_2 \dots$ ” (we use a large dot \bullet to make the “end” of the infinite chain of abstractions really clear). Thus there is some notion of infinite η -expansion. If we think about the denotation of $\lambda x.x$ in the game models, it is similarly expanded — it copies the whole of the first subtree to the rest of the arena, as if copying not only the x variable but also all of its arguments. This correspondence turns out to be general, and can be made precise by relating innocent strategies in economical form to a kind of (infinitely) η -expanded Böhm trees first studied by Nakajima in [6]. We call a formal connexion of this form an *Exact Correspondence Theorem*.

For a λ -term s the *Nakajima tree* of s , written $\text{NT}(s)$, is (informally) the countably branching, countably deep tree labelled as follows. If s is unsolvable then $\text{NT}(s) = \perp$, the empty tree. If s has HNF $\lambda x_1 \dots x_n \bullet y s_1 \dots s_m$ then

$$\text{NT}(s) = \begin{array}{c} \lambda x_1 \dots x_n z_0 z_1 \dots \bullet y \\ \swarrow \quad \downarrow \quad \searrow \\ \text{NT}(s_1) \quad \dots \quad \text{NT}(s_m) \quad \text{NT}(z_0) \quad \text{NT}(z_1) \quad \dots \end{array}$$

where z_0, z_1, \dots are countably many fresh variables. The process of finding such fresh variables given in [6] is quite complicated. In [4] we propose a *variable-free* representation of Nakajima trees so that for a closed term s , $\text{NT}(s)$ is represented as $\text{VFF}(s)$, a partial function from \mathbb{N}^* to $\mathbb{N} \times \mathbb{N}_0$. Note that the “infinitely nested” λ -abstractions of the form $\lambda z_1 z_2 \dots \bullet y$, which label the nodes of a Nakajima tree (of a closed term), can be coded as a pair (i, r) whereby the head variable y is the i^{th} in the infinite list of variables bound by the λ -abstraction situated r levels up in the tree. The map $\text{VFF}(s)$ is just a function that maps occurrences (of nodes) to such labels encoded as pairs of numbers.

The theorem of key importance in [4] is the *Exact Correspondence Theorem*, which states that for every closed λ -term s , the innocent strategy denoting s (in both \mathcal{D} and \mathcal{D}_{REC}) given in economical form is exactly $\text{VFF}(s)$, the Nakajima tree of s in variable free form.

Example 2. We now introduce example terms and strategies which we will use repeatedly to illustrate many of the concepts in the rest of the paper. Consider the terms $I = \lambda x.x$ and $1 = \lambda xy.xy$. The reader may wish to verify that the following represents the first two levels of the Nakajima trees of those terms:

$$\begin{array}{ccc} \text{NT}(I) = \lambda x z_0 z_1 \dots \bullet x & & \text{NT}(1) = \lambda x y z_0 z_1 \dots \bullet x \\ \begin{array}{c} \diagup \quad \diagdown \\ \lambda \mathbf{u} \bullet z_0 \quad \lambda \mathbf{v} \bullet z_1 \quad \lambda \mathbf{w} \bullet z_2 \dots \end{array} & & \begin{array}{c} \diagup \quad \diagdown \\ \lambda \mathbf{u} \bullet y \quad \lambda \mathbf{v} \bullet z_0 \quad \lambda \mathbf{w} \bullet z_1 \dots \end{array} \end{array}$$

After renaming of bound variables, these are the same. Since I and 1 differ only by η -conversion, this should be no surprise. Thus we can calculate their common variable-free form, the first two levels of which is:

$$\begin{array}{c} (1, 0) \\ \diagup \quad \diagdown \\ (2, 1) \quad (3, 1) \quad (4, 1) \quad \dots \end{array}$$

For example, the node labelled $(2, 1)$ means that the head variable of the corresponding node in the Nakajima tree is found as the second in the list of variables abstracted at the node one level above. The Exact Correspondence Theorem tells us that $\llbracket I \rrbracket = \llbracket 1 \rrbracket$ has the economical form which is given (in part) by $\varepsilon \mapsto (1, 0)$, $\langle 1 \rangle \mapsto (2, 1)$, $\langle 2 \rangle \mapsto (3, 1)$ and so on.

We say that a λ -algebra is *universal* if every element is the denotation of some λ -term. By the Exact Correspondence Theorem, it is easy to see that neither \mathcal{D} nor \mathcal{D}_{REC} is universal, since no *non-trivial* compact innocent strategy can be the denotation of any λ -term (note that the only finite Nakajima tree is the single-node tree \perp). Our aim in the rest of this section is to characterise the definable parts of \mathcal{D}_{REC} , and we shall do so by capturing the right ambient CCC.

Notation For tree-like $A \subseteq \mathbb{N}^*$ (i.e. those subsets which are prefix-closed and satisfy $\mathbf{s} \cdot n \in A \rightarrow \mathbf{s} \cdot m \in A$ for all $m < n$) and for any $\mathbf{s} \in A$ we define

$$\begin{array}{l} A @ \mathbf{s} = \text{the subtree of } A \text{ rooted at } \mathbf{s} \\ A^{>m} = \text{the tree obtained from } A \text{ by deleting the first } m \text{ branches.} \end{array}$$

For example, for the maximal single-tree arena U , we have $U @ \mathbf{s} = U = U^{>n}$ for all sequences \mathbf{s} and numbers n . Next fix an innocent strategy in economical form f and let $\mathbf{v} \in \text{dom}(f)$. We shall use the following shorthand:

$$\begin{array}{l} \mathbf{m}_{\mathbf{o}}^f(\mathbf{v}) = \text{the last move of the P-view encoded by } \mathbf{v} \\ \mathbf{m}_{\mathbf{p}}^f(\mathbf{v}) = \text{the response of } \sigma_f \text{ at the P-view.} \end{array}$$

Note that the former is by definition an O-move and the latter a P-move. We omit the superscript f wherever it is clear which strategy is intended. For example, for any innocent strategy f the O-move $\mathbf{m}_{\mathbf{o}}(\varepsilon)$ is the initial move ε and $\mathbf{m}_{\mathbf{p}}(\varepsilon)$ is the first P-move made by σ_f in response. Now we can define a new property of strategies:

Definition 2. Consider an innocent strategy in economical form $f : \mathbb{N}^* \rightarrow \mathbb{N} \times \mathbb{N}_0$, over some single-tree arena A . We say that f is everywhere copycat (EC) at $\mathbf{v} \in \mathbb{N}^*$ if f is undefined at \mathbf{v} or the following hold:

- (i) The arenas $A @ \mathbf{m}_o(\mathbf{v})$ and $A @ \mathbf{m}_p(\mathbf{v})$ are order-isomorphic (with respect to the prefix ordering).
- (ii) Whenever $\mathbf{w} \geq \mathbf{v}$ we have that for all $i \in \mathbb{N}$ $f(\mathbf{w} \cdot i) = (i, 1)$.
- (iii) If $f(\mathbf{v}) = (i, p)$ then $p > 0$.


We say that f is almost-everywhere copycat (AC) at \mathbf{v} if f is undefined at \mathbf{v} or there exist numbers $t_v \in \mathbb{N}_0$ and $o_v \in \mathbb{Z}$ with $o_v \leq t_v$ called the copycat threshold and offset respectively, such that

- (i) The arenas $(A @ \mathbf{m}_o(\mathbf{v}))^{>(t_v - o_v)}$ and $(A @ \mathbf{m}_p(\mathbf{v}))^{>t_v}$ are isomorphic.
- (ii) For all $i > t_v$, $f(\mathbf{v} \cdot i) = (i - o_v, 1)$ and f is everywhere copycat at $\mathbf{v} \cdot i$.
- (iii) For all $\mathbf{w} \geq (\mathbf{v} \cdot k)$ with $k \leq t_v$, if $f(\mathbf{w}) = (i, |\mathbf{w}| - |\mathbf{v}|)$ then $i \leq t_v - o_v$.
- (iv) If $f(\mathbf{v}) = (i, 0)$ then $i \leq t_v - o_v$.

(Note that f is EC at \mathbf{v} if and only if f is AC at \mathbf{v} with $t_v = o_v = 0$.)

Finally, we say that f is effectively almost-everywhere copycat (EAC) if f is computable, almost-everywhere copycat at every sequence on which it is defined and the functions $\mathbf{v} \mapsto t_v$ and $\mathbf{v} \mapsto o_v$ are computable. A strategy σ over an arena A is EAC if its innocent function is EAC, and we can generalise to multiple-tree arenas in the usual way.

To illustrate the definition of everywhere copycat strategies, suppose f is defined at \mathbf{v} . Intuitively we say that f is everywhere copycat at \mathbf{v} if, from $\mathbf{m}_p(\mathbf{v})$ onwards, f 's behaviour is simply to play copycat for as long as the arena will allow it. So if O's move is mi , the i^{th} child of the justifying move m , then P responds with the i^{th} child of the move immediately preceding m in the P-view. Condition (i) in the definition guarantees that P's copycat move will always be available. As before we will primarily be interested in strategies on U . Since $U @ \mathbf{s} = U = U^{>n}$ for all sequences \mathbf{s} and numbers n , Condition (i) will always hold. Condition (ii) is best understood with reference to the Exact Correspondence Theorem which relates innocent strategies to Nakajima trees. It specifies that the subtree of the Nakajima tree corresponding to f , rooted at \mathbf{v} , has the following shape:



Condition (iii) of the definition is a technicality, which ensures that the variable y is not one of the x_i .

Definition 3. The category of arenas and EAC strategies, \mathbb{A}_{EAC} , has recursive arenas as objects and EAC strategies on $A \Rightarrow B$ as morphisms from A to B .

A main result in [4] is that the category \mathbb{A}_{EAC} is well-defined; the proof that EAC strategies compose is highly technical. In fact,

Theorem 2. \mathbb{A}_{EAC} is cartesian closed.

The arena U is still an object of \mathbb{A}_{EAC} and still equal to its function space. Thus we can define a $\lambda\eta$ -algebra $\langle \mathbb{A}_{\text{EAC}}, U, \text{id}_U, \text{id}_U \rangle$ which we shall denote by \mathcal{D}_{EAC} . Properties of \mathcal{D}_{EAC} will be presented later.

4 Effectively and Explicitly Almost-Everywhere Copycat Strategies

We wish to find a new game model which invalidates the rule of η -conversion. To do so, we would require the terms I and 1 to be denoted differently. They have the same variable-free form of Nakajima tree, so it is not apparent how this might be achieved. The key is to make use of the fact that the copycat thresholds are not unique — any number greater than a given valid copycat threshold is also a valid copycat threshold. Different thresholds (at some P-view) may be used to distinguish I and 1 .

This idea is prompted by the observation that when one compares a term with its denotation, the part of the EAC strategy which is specified by the rules of copycat corresponds precisely to the part of the Nakajima tree which has been generated by η -expansion (i.e. the part of the tree with the fresh variables as the head variables). Recall the Nakajima trees of I and 1 — the former has fresh variables appearing at every node except the root, whereas the latter is similar except that there is not a fresh variable at the first child of the root. Therefore we aim to find a model where I and 1 are represented by the strategy with the same moves, but the copycat threshold of $\llbracket I \rrbracket$ at the first P-view is 0, whereas that of $\llbracket 1 \rrbracket$ is 1.

However, the definition of an EAC strategy is stated in terms of the existence of some computable function which associates a pair of numbers to each P-view of the strategy and this function is *not* specified along with the strategy. (A consequence of this is that there is no computable procedure for finding valid thresholds for an EAC strategy.) It is really the thresholds (rather than the offsets) that are important because, for a certain P-view v of an EAC strategy σ , the copycat threshold t gives enough information to compute the offset o directly. This motivates the following definition:

Definition 4. *An effectively and explicitly almost-everywhere copycat strategy (EXAC strategy) is given by a pair $\langle \sigma, t_\sigma \rangle$, where σ is an EAC strategy and t_σ is an effective function mapping the P-views where σ is defined to valid copycat thresholds. We sometimes write the EXAC strategy $\langle \sigma, t_\sigma \rangle$ just as σ .*

We will usually refer to the first and second part of an EXAC strategy as the “(underlying) EAC strategy (part)” and the “threshold function (part)”, respectively. In view of our comments above, however, we will sometimes speak of the offsets as if they too are specified by the threshold function.

This definition allows us to make the intended finer distinction between strategies: two strategies with the same moves must be equal as EAC strategies, but may have different copycat thresholds and so can be distinguished as EXAC strategies. There is an obvious forgetful map from EXAC strategies to

EAC strategies, which takes only the strategy part (i.e. erasing the threshold information).

In a similar vein to the economical form of innocent strategies, using the same encoding of a P-view as a sequence of natural numbers, we can give an economical form of EXAC strategies over single-tree arenas. We can also take advantage of the fact that parts of the strategy are completely dictated by its copycat nature. Let us say that a P-view is *entirely explicit* if none of the O-moves in it exceed the copycat threshold of the P-view at which they are made. Thus if a P-view is not entirely explicit the ensuing move can be deduced from the threshold and offset of the P-view preceding the first O-move in it which did exceed the copycat threshold.

Definition 5. *The economical form of an EXAC strategy is a map from \mathbb{N}^* to $\mathbb{N} \times \mathbb{N}_0 \times \mathbb{N}_0 \times \mathbb{Z}$. The domain is the encoding of P-views in the usual way. The map is defined at a sequence \mathbf{v} only if the P-view encoded by \mathbf{v} is entirely explicit, in which case*

$$\mathbf{v} \mapsto (i, r, t, o)$$

where the resulting P-move is encoded as before — it is the i^{th} child of the move $2r$ from last of the P-view — and the copycat threshold and offset at this P-view are t and o respectively.

Example 3. We take the EXAC strategies η_0 and η_1 to be $\langle \llbracket I \rrbracket, t_0 \rangle$ and $\langle \llbracket 1 \rrbracket, t_1 \rangle$, where t_0 maps every P-view to the threshold 0 and t_1 does likewise except that the minimal P-view is mapped to the threshold 1. Since $\llbracket I \rrbracket = \llbracket 1 \rrbracket$, they have the same EAC strategy part, but different threshold functions. These are the suggestions we made for the denotations of I and 1 in a model not supporting η -conversion. Nearly every P-view of either is not entirely explicit, and the respective economical forms are given by:

$$\begin{aligned} \varepsilon &\mapsto (1, 0, 0, -1) \quad \text{and} \quad \varepsilon &\mapsto (1, 0, 1, -1) \\ \langle 1 \rangle &\mapsto (2, 1, 0, 0) \end{aligned}$$

We now need a method to compose EXAC strategies. Of course the EAC strategy part will just be the standard composition of innocent strategies, and we give below an algorithm for computing the composition of the threshold functions.

Algorithm (The Composition Algorithm). Let $\langle \sigma, t_\sigma \rangle$ be an EXAC strategy over $A \Rightarrow B$, and $\langle \tau, t_\tau \rangle$ be an EXAC strategy over $B \Rightarrow C$. Take a P-view \mathbf{v} on which the strategy $\sigma; \tau$ (which is given by the usual composition of innocent strategies) is defined and suppose that the last move of the P-view is \underline{m} and the resulting move is m .

We write $\mathbf{u} = \mathbf{u}(\mathbf{v}, \sigma, \tau)$ for the *uncovering* of the composition up to the move m . A formal definition can be found in [3] or [4], but we may describe it as the sequence of moves of the composition which result after the P-view \mathbf{v} , including any relevant B-moves which would be hidden by the composition. It will be of the form $\langle \varepsilon, \dots, \underline{m}, m_1, m_2, \dots, m_{p-1}, m_p, m \rangle$.

The moves m_i are the intermediate interactions which might have taken place between σ and τ before the move m became the visible outcome, and are all in the arena B . Possibly there are no such intermediate moves, in which case $p = 0$. We do not care about justification pointers, and for tidiness set $m_0 = \underline{m}$ and $m_{p+1} = m$.

For $1 \leq i \leq p + 1$ we consider the P-view \mathbf{u}_i that the strategies σ or τ are faced with when the move m_i was made. (For details on how one may define such a P-view precisely see [4] or [5]). Define t_i and o_i to be the copycat threshold and offset of σ , or τ as appropriate, at the P-view \mathbf{u}_i . These are specified by t_σ or t_τ . Then set:

$$\begin{array}{ll} t'_i = t_i + |A|, & \text{if } m_i \text{ is a root of the arena } B \\ t_i, & \text{otherwise} \\ o'_i = o_i + |A|, & \text{if } m_i \text{ is a root of the arena } B \\ o_i, & \text{otherwise} \end{array} \quad \begin{array}{l} T_1 = t'_1 \\ O_1 = o'_1 \\ T_{i+1} = \max(T_i + o'_{i+1}, t'_{i+1}) \\ O_{i+1} = O_i + o'_{i+1} \end{array}$$

$$\begin{array}{l} \mathbf{t} = T_{p+1} \\ \mathbf{o} = O_{p+1} - |A| + |B|, & \text{if } \underline{m} \text{ is a root of the arena } C \\ O_{p+1}, & \text{otherwise} \end{array}$$

(By $|A|$ we mean the number of trees in the arena A). Then \mathbf{t} and \mathbf{o} are the copycat threshold and offset of the composition $\langle \sigma, t_\sigma \rangle; \langle \tau, t_\tau \rangle$ at the P-view \mathbf{v} .

Now we must show that this method does indeed produce an EXAC strategy, i.e. that the composite threshold function specifies valid thresholds and offsets for the composite strategy. In fact it does so only under some restrictions, for which we need an additional definition.

Definition 6. *Let σ be an EAC strategy over a single-tree arena. If σ has a first move, then it has a copycat threshold and offset, say t and o , at the P-view consisting only of the root O-move (we call this P-view the minimal P-view). The l-number of σ is the value $t - o$, and we write it $\mathbf{l}(\sigma)$.*

If $\sigma = \langle \sigma_1, \dots, \sigma_n \rangle$ is an EAC strategy over an arena with n trees, and defined on at least one of the minimal P-views, then we say $\mathbf{l}(\sigma) = \min_{i=1, \sigma_i \neq \perp}^n \{\mathbf{l}(\sigma_i)\}$.

This is termed the l-number of σ because, as will eventually be shown, it corresponds to the number of λ -abstractions at the root of the Böhm tree of the term whose denotation is σ . For example $\mathbf{l}(\eta_0) = 1$ and $\mathbf{l}(\eta_1) = 2$, and we will be able to show that η_0 is the denotation of $\lambda x.x$ and η_1 the denotation of $\lambda xy.xy$.

Theorem 3. *If $\sigma : A \Rightarrow B$ and $\tau : B \Rightarrow C$ are EAC strategies satisfying $\mathbf{l}(\sigma) \geq |A|$ (or σ is everywhere undefined) and $\mathbf{l}(\tau) \geq |B|$ (or τ is everywhere undefined) then Algorithm 4 produces valid copycat thresholds and offsets for $\sigma; \tau$.*

There is an “obvious” category, which derives directly from the conditions required for the composition algorithm to work correctly.

Definition 7. The category of arenas and EXAC strategies, written \mathbb{A}_{EXAC} , has recursive arenas as objects, and the morphisms from A to B are the EXAC strategies on the arena $A \Rightarrow B$ which have l-number greater than or equal to $|A|$, or are everywhere undefined. The identity morphism on A is the EXAC strategy $\langle \text{id}_A, 0 \rangle$, i.e. the copycat threshold is zero everywhere.

One can show that this does indeed specify a category. Also, \mathbb{A}_{EXAC} has the obvious terminal object — the empty arena E — and products given in the usual way. However, \mathbb{A}_{EXAC} does not form a CCC with the usual constructions, as the following example shows:

Example 4. Suppose that $\sigma : A \times B \rightarrow C$. Then we know that $\mathbf{l}(\sigma) \geq |A| + |B|$. We need a morphism $\Lambda(\sigma) : A \rightarrow B \Rightarrow C$, which must have l-number at least $|A|$, so we could take $\Lambda(\sigma)$ to be the same EXAC strategy as σ . However this choice may not be unique. For example, consider η_0 and η_1 as defined earlier in this section. One can verify that both η_0 and η_1 can be considered as morphisms $U \rightarrow U \Rightarrow U$ and that in this case $\eta_0 \times \text{id}_U; \text{eval}_{U,U} = \eta_1 \times \text{id}_U; \text{eval}_{U,U} : U \times U \rightarrow U$, and that this is the same as the morphism $U \times U \rightarrow U$ described by η_1 . Hence there are two candidates for $\Lambda(\eta_1)$.

It is not clear that \mathbb{A}_{EXAC} forms a CCC with any unusual function space constructions either.

If we try to fix the definition of \mathbb{A}_{EXAC} , by cutting down the homsets some more, it becomes clear that one must also specify minimum copycat thresholds at the minimal P-views, along with minimum l-numbers. The obvious solution still does not work, and we can repeat the fixing-up process to obtain a sequence of failures — each is either not a category at all because identities fail to work properly, or has a non-uniqueness of curried morphisms as above.

We now present a new category based on EXAC strategies, which does form a CCC. Although it does appear to be much more complicated than the “almost-CCC” \mathbb{A}_{EXAC} , it seems to be the natural limit of the fixing-up process.

Firstly let us write $\text{br}(A)$ for the number of branches of a tree at the root (assuming that this is finite). Then we can write $\text{br}(A @ m)$, for any move m of a finitely-branching forest A , to mean the number of direct children of m in A . Then we make the following definition:

Definition 8. Let A be an arena and X a finitely-branching subarena³ of A . We say that an EXAC strategy σ over A is X -explicit if the following holds:

Let $\sigma : \mathbf{v} \mapsto (i, r, t, o)$ be the economical form of any clause of the innocent function. Suppose that the sequence \mathbf{v} codes a P-view ending in the O-move \underline{m} , and that the consequent P-move encoded by this clause is m . Then:

- (i) if \underline{m} is in the subarena X then $t - o \geq \text{br}(X @ \underline{m})$,
- (ii) if m is in the subarena X then $t \geq \text{br}(X @ m)$.

³ We say that the arena $A = \langle A_1, \dots, A_m \rangle$ is a subarena of $B = \langle B_1, \dots, B_n \rangle$ if $m = n$, and for each i , A_i is a subset of B_i . We say that an arena is finitely-branching if every tree in it is finitely-branching

An intuitive description of this definition is the following: The subarena X determines a part of the arena A where the strategy is known to be explicitly defined, i.e. moves in X are neither in the “domain” nor the “range” of automatic copycat forced by the threshold information of σ . This means that given a strategy σ over A which is X -explicit, any P-view of σ with moves only in X is entirely explicit.

Definition 9. *The category $\mathbb{X}\mathbb{A}_{\text{EXAC}}$, or simply $\mathbb{X}\mathbb{A}$, is given by the following: objects are pairs (A, X) consisting of a recursive arena A and a finitely-branching recursive subarena X ; a morphism $\sigma : (A, X) \rightarrow (B, Y)$ is an EXAC strategy on $A \Rightarrow B$ which is $(X \Rightarrow Y)$ -explicit. Composition of morphisms is composition of EXAC strategies, and the identity strategy on (A, X) , $\text{id}_{(A, X)}$, is the EXAC strategy $\langle \text{id}_A, t \rangle$, where id_A is the EAC identity strategy on A , and t is the function that takes the least value on every P-view which still leaves the EXAC strategy $\langle \text{id}_A, t \rangle$ as $(X \Rightarrow X)$ -explicit.*

The fact that the composition algorithm gives valid thresholds and offsets for EXAC strategies satisfying these conditions comes from Theorem 3 — a morphism $\sigma : (A, X) \rightarrow (B, Y)$ is $(X \Rightarrow Y)$ -explicit, and X has the same number of trees as A , so in particular $\mathbf{l}(\sigma) \geq |A|$.

Theorem 4. *$\mathbb{X}\mathbb{A}$ forms a CCC with the following constructions: the terminal object $\mathbf{1}$ is (E, E) , where E is the empty arena; the product $(A, X) \times (B, Y)$ is $(A \times B, X \times Y)$, with the threshold functions for projections specified in the same style as identities; the exponential object $(A, X) \Rightarrow (B, Y)$ is $(A \Rightarrow B, X \Rightarrow Y)$, and the evaluation map $\text{eval}_{(B, Y), (C, Z)}$ is the same EXAC strategy as $\text{id}_{(B, Y) \Rightarrow (C, Z)}$.*

Now that we have found an ambient CCC for the EXAC strategies, we can construct another λ -algebra based on it. In this category, however, the reflexive object is not isomorphic to its function space — exactly as we would hope for a model invalidating η -conversion.

Let us write U^0 for the object $\langle U, M \rangle$ of $\mathbb{X}\mathbb{A}$, and U^1 for the object $U^0 \Rightarrow U^0$. Here U and M are the maximal and minimal single-tree arenas described in Sect. 2. We define morphisms $F : U^0 \rightarrow U^1$ and $G : U^1 \rightarrow U^0$ to both be given by the EXAC strategy η_1 (the definition of which can be found in Sect. 4). It is straightforward to check that this does give proper morphisms, and that they satisfy $G; F = \text{id}_{U^1}$ and $F; G \neq \text{id}_{U^0}$.

Hence we can identify a new λ -algebra $\langle \mathbb{X}\mathbb{A}, U^0, F, G \rangle$ which invalidates the rule of η -conversion; we denote this λ -algebra \mathcal{M} . By erasing all threshold information, we can reduce \mathcal{M} to (a subset of) \mathcal{D} and deduce that \mathcal{M} is also sensible.

In the same way that the denotation of a term in the model \mathcal{D}_{EAC} had a strong connection with its Nakajima tree, the denotation in \mathcal{M} corresponds closely to (a variable-free version of) the Böhm tree. The variable-free form of the Böhm tree of a term is similar to the construction VFF mentioned earlier and defined in [4], but it includes extra information describing how many abstractions there are at each node, and how many children.

Definition 10. For a $(\mathbb{N} \times \mathbb{N}_0 \times \mathbb{N}_0 \times \mathbb{Z})$ -labelled tree p the tree p^* is the same tree labelled identically, except that nodes at depth d labelled $(i, d + 1, t, o)$ are relabelled $(i, d + 2, t, o)$.

Similarly the tree $\{p\}^n$, for $n \in \mathbb{N}_0$, is labelled identically except that firstly the node at the root (i, r, t, o) is first relabelled to $(i, r, t, o - n)$, and then nodes of depth d are relabelled as follows:

- (i) those labelled (i, d, t, o) are relabelled $(i + n, d, t, o)$;
- (ii) those labelled $(i, d + 1, t, o)$ for $i \leq n$ are relabelled $(n - i + 1, d, t, o)$;
- (iii) those labelled $(i, d + 1, t, o)$ for $i > n$ are relabelled $(i - n, d + 1, t, o)$.

For a term s with free variables within Δ the variable-free form of the Böhm tree of s , $\text{VFBT}_\Delta(s)$, is the following $(\mathbb{N} \times \mathbb{N}_0 \times \mathbb{N}_0 \times \mathbb{Z})$ -labelled tree:

$$\begin{aligned} \text{VFBT}_\Delta(s) &= \perp, && \text{the empty tree, for unsolvable } s. \\ \text{VFBT}_\Delta(\lambda x_1 \dots x_n \bullet s) &= \{ \text{VFBT}_{\Delta \cdot \langle x_1, \dots, x_n \rangle}(s) \}^n, \\ &&& \text{if } s \text{ is of the form } v_j s_1 \dots s_m. \\ \text{VFBT}_\Delta(v_j s_1 \dots s_m) &= && \begin{array}{c} (j, 1, m, m) \\ \swarrow \quad \searrow \\ \text{VFBT}_\Delta(s_1)^* \quad \text{VFBT}_\Delta(s_m)^* \end{array} \\ &&& \text{where } \Delta = \langle v_k, \dots, v_1 \rangle \text{ (note the reverse order)}. \end{aligned}$$

Exactly as before we can show that, at each node of the Böhm tree of s , the first two elements of the tuple at the corresponding node of $\text{VFBT}(s)$ encode the head variable by counting how many levels one goes up the tree, and how many abstractions along, to find where the variable is abstracted. The third component just counts the number of children at the node, and the fourth is the number of children minus the number of abstractions. We choose to encode the number of abstractions in this rather elliptic fashion in order to make the following theorem easier to state:

Theorem 5 (Exact Correspondence for \mathcal{M}). If $s \in \Lambda$ with free variables in $\Delta = \langle v_k, \dots, v_1 \rangle$ then $\llbracket s \rrbracket_\Delta^{\mathcal{M}} = \{ \text{VFBT}_\Delta(s) \}^k$ when the former is considered as an EXAC strategy in economical form and the latter as a labelling function.

In particular for closed terms s , $\llbracket s \rrbracket_\varepsilon = \text{VFBT}_\varepsilon(s)$

Example 5. Although it is hard to see directly, the given definition of VFBT does work as intended. One may check that $\text{VFBT}(I)$ and $\text{VFBT}(1)$ are:

$$\begin{array}{ccc} (1, 0, 0, -1) & \text{and} & (1, 0, 1, -1) \\ & & \downarrow \\ & & (2, 1, 0, 0) \end{array}$$

The node $\lambda xy.x$, in the Böhm tree of 1, corresponds to the node of $\text{VFBT}(1)$ labelled $(1, 0, 1, -1)$, which is so labelled because the head variable is the first abstracted variable zero levels up the tree (namely x), the node has one child, and the number of abstractions at this level is $1 - (-1) = 2$. The Exact Correspondence Theorem gives us the economical forms of $\llbracket I \rrbracket$ and $\llbracket 1 \rrbracket$ which, as we hoped, are the EXAC strategies η_0 and η_1 described earlier in this section.

The Exact Correspondence results are of key importance in examining the local structure of the game models. In [4] we use them, and a powerful result which we call the Separation Lemma, to obtain proofs that \mathcal{D}_{EAC} is a *universal* and *order-extensional* λ -model, with equational theory given by \mathcal{H}^* (the maximal consistent sensible theory). The models \mathcal{D} and \mathcal{D}_{REC} are neither universal nor extensional, but do generate the same equational theory on terms.

In a similar way, we can use the Exact Correspondence Theorem for \mathcal{M} to prove the following:

Theorem 6. (i) \mathcal{M} is universal i.e. every element is definable as the denotation of some term of the λ -calculus. (ii) \mathcal{M} is weakly extensional, so it is a λ -model (for a discussion of weak extensionality see [1, §5]). (iii) \mathcal{M} equates two terms of the λ -calculus precisely when they have the same Böhm tree. Thus the local structure of the model is the Böhm tree theory \mathcal{B} .

5 Further Work

Although we succeeded in our aim to find a universal game model of \mathcal{B} , there are other questions which the work prompts. Firstly, one might ask if there is a less syntactic way to arrive at \mathcal{D}_{EAC} from \mathcal{D}_{REC} , perhaps by some sort of *extensional collapse*. In fact extensional collapse appears to be insufficient, and further investigation would be of interest. In another direction, we can use \mathcal{D} , \mathcal{D}_{REC} , \mathcal{D}_{EAC} and \mathcal{M} as very natural combinatory algebras which are in some sense *sequential*. Therefore one might wish to study realizability models over them. Finally, and more practically, we could examine the Böhm tree composition algorithm given by the game model: it is quite different from the usual method in that it is “demand-driven” – for every node of the answer only the relevant nodes of the composed trees are examined.

References

- [1] Barendregt, H.: *The Lambda Calculus, Its Syntax and Semantics*. North-Holland, revised edition, 1984.
- [2] Di Gianantonio, P., Franco, G., Honsell, F.: Games Semantics for Untyped λ -Calculus, Preprint, announced on the Types Mailing List. April 1998.
- [3] Hyland, J. M. E, Ong, C.-H. L.: On Full Abstraction for PCF: I, II and III. To appear in *Information and Computation*, 133 pages, ftp-able. 1994.
- [4] Ker, A., Nickau, H., Ong, C.-H. L.: Game Models of Untyped λ -Calculus. Submitted for publication, 49 pages, ftp-able. 1998.
- [5] McCusker, G.: *Games and Full Abstraction for a Functional Metalanguage with Recursive Types*. Cambridge University Press, 1998.
- [6] Nakajima, R.: Infinite Normal Forms for the λ -Calculus. In Goos, G. and Hartmanis, J., editors, *Symposium on λ -Calculus and Computer Science, LNCS 37*, pages 62–82. Springer-Verlag, 1975.
- [7] Nickau, H.: *Hereditarily Sequential Functionals: A Game-Theoretic Approach to Sequentiality*. Dissertation, Universität GH Siegen. Shaker-Verlag, 1996.