On the Formal Semantics of IF-Like Logics

Santiago Figueira[1,3], Daniel Gorín[1], and Rafael Grimson[2]

[1] Departamento de Computación, FCEyN, Universidad de Buenos Aires, Argentina
[2] Departamento de Matemática, FCEyN, Universidad de Buenos Aires, Argentina, Hasselt University and Transnational University of Limburg
[3] CONICET, Argentina

Abstract. In classical logics, the meaning of a formula is invariant with respect to the renaming of bound variables. This property, normally taken for granted, has been shown not to hold in the case of Information Friendly (IF) logics. In this work we propose an alternative formalization under which invariance with respect the renaming of bound variables is restored. We show that, when one restricts to formulas where each variable is bound only once, our semantics coincide with those previously used in the literature. We also prove basic metatheoretical results of the resulting logic, such as compositionality and truth preserving operations on valuations. We work on Hodges' *slash logic* (from which results can be easily transferred to other IF-like logics) and we also consider his flattening operator, for which we give a game-theoretical semantics.

1 Introduction

Independence Friendly logic (IF, for short) was introduced and promoted as a new foundation for mathematics by Jaako Hintikka over a decade ago [9,10]. Closely related to Henkin's logic of *branching quantifiers* [8,16,7,2], IF is an extension of first-order logic where disjunctions and existential quantifiers may be decorated with denotations of universally-quantified variables. The intended meaning of a formula $\forall x \exists y/_{\forall x} \varphi$ is that the value for y may not *depend* on x (in other words, it may not be function of x). This notion is nicely formalized using a two player game between Abélard and Eloïse, which, because of the independence restrictions, is of *imperfect information.*

It was conjectured by Hintikka that one could not formulate IF semantics in a *composable* way [9]. This was promptly rebutted by Hodges in [11], where he achieves compositionality by taking as the interpretation of a formula $\varphi(x_1, \ldots, x_n)$ over the domain A, the set of sets of n-tuples (called *trumps*) for which Eloïse has a uniform winning strategy.

Two things are worth observing. First, in [11] Hodges introduced two slight modifications in syntax and semantics, namely: conjunctions and universal quantifiers may also be decorated with restrictions, and restrictions on any of the player's choices may range also over any of his previous choices[1]. Hodges later

[1] In Hintikka's presentation [9], Eloïse is not allowed to take into account her previous choices. For implications of this fact see, e.g. [14].

W. Hodges and R. de Queiroz (Eds.): WoLLIC 2008, LNAI 5110, pp. 164–178, 2008.

coined the name *slash logic* for his formulation and noticed that many writers have transferred the name 'IF logic' to slash logic, often without realising the difference [12]. We will use the term IF-*like logics* to encompass this variety of related logics. In [13], Hodges shows that even if one restricts to Hintikka and Sandu's original formulation of IF, compositionality can be obtained. The second thing to note is that in both papers Hodges considers only the syntactic fragment where each variable may be bound only once. The underlying assumption is that, given an arbitrary formula, one can appropriately rename its variables, so no generality is loss. In the light of later findings, it is not obvious whether this is was a reasonable assumption.

Caicedo and Krynicki [4] proved a prenex normal form theorem for slash logic. To account for arbitrary formulas, where variables occur in any order, and may get rebound, they used compositional semantics in the line of Hodges, but with n-tuples replaced with valuations. This extension seemed so natural that in later papers it was taken as the standard semantics of slash logic.

Based on this formulation, in [14], Janssen pointed out several strange properties of these logics. At the root of them lies the idea of *signaling*, i.e., "the phenomenon that the value of a variable one is supposed not to know, is available through the value of another variable" [15]. He observes that if variables are reused, signaling may be blocked and, thus, the truth-value of formulas that only differ on bound-variables may differ. This can even be the case of formulas of IF-logic without restrictions, which would challenge Hintikka's claim of IF being a *conservative extension* of classical logic [9].

A systematic analysis of signaling in IF-like logics was later performed in [15], where several claims of "equivalence of formulas under syntactic transformations" made in [4] are questioned due to signalings that may get unexpectedly blocked. These results were fixed in [3] by restating them in a much weaker sense.

Summing up, on the one hand, we have a family of logics, aiming to be a conservative extension of first-order logic, for which several results have been proved, but that hold only for the regular fragment. On the other hand, we have that the attempts to formulate general results for the whole fragment failed. In the face of this, Dechesne advocated for the restriction of IF-like logics to the *regular fragment*, where no rebinding of variables occur (cf. Section 7.5 of [6]).

In this paper, we argue that there is no real need to restrict IF-like logics to regular formulas and that, in fact, most, if not all, of previous results can be generalized to the irregular case in a safe and natural way. In a nutshell, we claim that classical valuations are simply not adequate to formalize independence restrictions in a context where variables can get rebound.

In Section 2 we discuss briefly the interaction between irregular formulas and classical valuations with respect to signaling. This motivates Section 3, where we avoid these problems using alternative semantics, that are equivalent for the regular fragment. In Section 4 we consider also the *flattening* operator $\downarrow$, introduced by Hodges in [11] and illustrate that irregular formulas can be handled uniformly also in this setting; while doing this, we provide a new (to the best of our knowledge) game semantics for this logic. All the proofs are in Appendix A.

2 Preliminaries

2.1 Syntax

From here on, we restrict ourselves to Hodges' *slash logic* (but without indexed disjunctions) [11,12], in which Hintikka's IF logic can be trivially embedded. Formulas are built out of an infinite supply of constant symbols, function symbols and relation symbols just like in first-order logic, using the following set of connectives: $\sim$, $\vee/_{y_1,\ldots,y_k}$ and $\exists x/_{y_1,\ldots,y_k}$, where $y_1,\ldots,y_k$ stands for a set of variables. The derived connectives $\wedge/_{y_1,\ldots,y_k}$ and $\forall x/_{y_1,\ldots,y_k}$ are defined in the usual way. We will also write $\wedge$, $\vee$, $\exists x$ and $\forall x$ for $\wedge/_\emptyset$, $\vee/_\emptyset$, $\exists x/_\emptyset$ and $\forall x/_\emptyset$. Following [4] we don't impose any restriction on the variables occurring under the slashes.

The sets of free and bound variables of φ, $\mathrm{Fv}(\varphi)$ and $\mathrm{Bv}(\varphi)$ respectively, are defined in the usual way. Of course, variables that occur under slashes must be taken into consideration; observe, for example, that if $\theta := \exists x/_{x,y}[x = z]$ then $\mathrm{Fv}(\theta) = \{x, y, z\}$ and $\mathrm{Bv}(\theta) = \{x\}$.

Following Dechesne [6], we will say that a formula φ is *regular* whenever $\mathrm{Fv}(\varphi) \cap \mathrm{Bv}(\varphi) = \emptyset$ and there is no nested quantification over the same variable. To follow Hodges' presentation, when referring to regular formulas we will sometimes make the *context* (i.e. the free variables in scope) a parameter of the formula by writing: $\varphi(x_1,\ldots,x_n)$, where $(x_1,\ldots,x_n)$ is an n-tuple of distinct variables such that $\mathrm{Fv}(\varphi) \subseteq \{x_1,\ldots,x_n\}$. Observe that this means that for a fixed φ, $\varphi(x, y)$ and $\varphi(x, y, z)$ will generally denote two non-equivalent formulas. See [11] for further details.

2.2 Semantics

We will consider two related semantics. On the one hand, there is Hodges' trump semantics, which we will call *T-semantics*. It is compositional and based on sets of tuples but its formalization requires regular formulas with the context as a parameter. On the other, we have Caicedo and Kynicki's extension of trump semantics to arbitrary formulas, which we will call *V-semantics*. It is based on sets of valuations and has a natural game-based formulation from which compositionality can be proved [4,3].

Let us begin with V-semantics. A formula φ is true in a model $\mathcal{M}$ under a set of valuations V, written $\mathcal{M} \models^+ \varphi[V]$, iff Eloïse has a valid strategy that, when followed, wins every instance $\mathsf{G}(\mathcal{M}, \varphi, v)$ (for $v \in V$) of the classical satisfaction game between Abélard and Eloïse. Dually, a formula is false, written $\mathcal{M} \models^- \varphi[V]$, whenever Abélard has a valid strategy that is winning for every $\mathsf{G}(\mathcal{M}, \varphi, v)$, $v \in V$. For a strategy to be valid, it has to satisfy additional independence conditions. For a formal presentation refer to [4,3].

Hodges avoided valuations in the first place by restricting to regular formulas where the context is a parameter: a valuation for $\varphi(x_1,\ldots,x_n)$ is simply an n-tuple $(a_1,\ldots,a_n)$. Let us say that $v_{(a_1,\ldots,a_n)}$ is a valuation such that $v_{(a_1,\ldots,a_n)}(x_i) = a_i$ when $1 \le i \le n$ and $v(x) = c$, for some fixed c, otherwise; then, intuitively, a *trump* (resp. *cotrump*) T for $\varphi(x_1,\ldots,x_n)$ in $\mathcal{M}$, written

$\mathcal{M} \models^+ \varphi(x_1, \ldots, x_n)[T]$ (resp. $\mathcal{M} \models^- \varphi(x_1 \ldots x_n)[T]$), is just a set of n-tuples for which Eloïse (resp. Abélard) has a strategy that is winning for every instance of the game $\mathsf{G}(\mathcal{M}, \varphi, v_{(a_1 \ldots, a_n)})$ for $(a_1 \ldots a_n) \in T$. This can be alternatively defined in a composable way; we include for reference such a formulation in Appendix A and refer the reader to [11] for further details.

Notation. Thoughout this paper, "$\mathcal{M} \models^\pm X$ iff $\mathcal{M} \models^\pm Y$" will stand for "$\mathcal{M} \models^+ X$ iff $\mathcal{M} \models^+ Y$, and $\mathcal{M} \models^- X$ iff $\mathcal{M} \models^- Y$".

As usual, each of these semantics gives rise to a notion of formula equivalence.

V-equivalence: $\varphi_1 \equiv_V \varphi_2$ iff $\mathrm{Fv}(\varphi_1) = \mathrm{Fv}(\varphi_2)$ and for all $\mathcal{M}$ and every set of valuations V, $\mathcal{M} \models^\pm \varphi_1[V]$ iff $\mathcal{M} \models^\pm \varphi_2[V]$.

T-equivalence: Let $\overline{x} = x_1, \ldots, x_n$. $\varphi_1(\overline{x}) \equiv_T \varphi_2(\overline{x})$ iff $\mathrm{Fv}(\varphi_1) = \mathrm{Fv}(\varphi_2)$ and for all $\mathcal{M}$ and every $T \subseteq |\mathcal{M}|^n$, $\mathcal{M} \models^\pm \varphi_1(\overline{x})[T]$ iff $\mathcal{M} \models^\pm \varphi_2(\overline{x})[T]$.

2.3 Signaling Kicks in

It was first observed by Jannsen [14] that V-semantics and signaling don't interact well. Consider, for instance, the following example (from [14], section 7, formulas (32) and (33)): $\theta_1 := \forall x \forall y \forall z[x = y \vee \exists u \exists w/_x[w \neq x \wedge u = z]]$ and $\theta_2 := \forall x \forall y \forall z[x = y \vee \exists y \exists w/_x[w \neq x \wedge y = z]]$ Clearly, θ_1 is a regular formula while θ_2 is not. Moreover, they only differ in the symbol used for a bound variable: u vs. y. Since variable symbols are expected to be simple placeholders, both formulas should be equivalent. Now, Eloïse has a winning strategy for θ_1, regardless the structure: $f_\vee(v) = L$ if $v(x) = v(y)$ and $f_\vee(v) = R$ otherwise; $f_{\exists u}(v) = v(z)$; $f_{\exists w/_x}(v) = v(y)$. Observe that Eloïse's strategy for θ_1 relies heavily on *signaling*: she needs a value other than $v(x)$ but her strategy function may not depend on x; however, y is *signaling* such a value.

The problem is that this strategy is not winning for θ_2: whenever Abélard picks different initial values for x and y, Eloïse will be forced to reset the value of y to that of z, breaking the global invariant of her strategy (i.e., blocking the signal). In fact, it is not hard to show that for arbitrary structures, Eloïse has no winning strategy for θ_2 which implies that $\theta_1 \not\equiv_V \theta_2$.

Now, although this is an already known example, we feel its significance has been overlooked. Variables (and specially those that are bound) ought to be a mere syntactic device, a simple placeholder. They should bear no meaning in itself. The only thing we should care about two bound variables x and y is that they are distinct and, as such, stand for distinct placeholders. In that sense u, v or w should be as good as y. In fact, we should expect to be able to drop variables altogether and replace them with some equivalent syntactic device, such as de Bruijn indices [5].

This notion is so crucial that it even has a name: *α-equivalence*. (for formal definitions see any textbook on λ-calculus, e.g. [1]). In every sensible formalism, α-equivalence implies equivalence. We already saw this does not hold in slash logic under V-semantics in general and the following example shows that it neither does restricted to regular formulas. Consider these α-equivalent, regular formulas: $\theta_3 := \exists y \exists z/_{x,y}[z = x]$ and $\theta_4 := \exists u \exists z/_{x,u}[z = x]$. For $||\mathcal{M}|| \geq 2$ and $V = \{v \mid v(x) = v(u)\}$ it is easy to see that $\mathcal{M} \models^+ \theta_3[V]$ but $\mathcal{M} \not\models^+ \theta_4[V]$.

Invariance under α-equivalence is such a basic property that it is not surprising that neither Hodges nor Caicedo and Krynicki mention it in their papers. However the latter two assumed it to hold and this lead to some flawed results (see [15]). In the face of this, it is worth verifying that, fortunately, α-equivalence does hold under T-semantics (the proof is on Section A.2).

Proposition 1. *Let* $\overline{x} = (x_1, \ldots, x_n)$*; if* $\varphi_1(\overline{x}) \equiv_\alpha \varphi_2(\overline{x})$ *then* $\varphi_1(\overline{x}) \equiv_T \varphi_2(\overline{x})$.

The fact that invariance under α-equivalence holds on regular formulas under T-semantics but fails under V-semantics is, in our opinion, a clear indication that this is neither a feature of these logics nor they should be restricted to the regular fragment. V-semantics simply fail to generalize properly the meaning given to the slashed connectives by the T-semantics.

3 Uniform Semantics for Regular and Irregular Formulas

Classical valuations are an inadequate device to formalize the semantics of unrestricted IF-like formulas: under rebinding of variables, they simply fail to keep track of all the previous choices, which is crucial in a setting of independence restrictions. Our plan is, roughly, to replace valuations with tuples $\langle s, p \rangle$, where $s \in |\mathcal{M}|^\omega$ is an infinite sequence of choices, and p is a mapping of variables into positions of s. A variable x gets thus interpreted as $s(p(x))$. Observe one can think of the composition $s \circ p$ as denoting a classical valuation[2].

Using games, we will define what we call *S-semantics*, that is, the relations $\mathcal{M} \models^+ \varphi[S, p, h]$ and $\mathcal{M} \models^- \varphi[S, p, h]$ where S is a nonempty set of sequences taken from $|\mathcal{M}|^\omega$, and $h < \omega$ can be regarded as indicating how many "previous choices" are in scope. After checking that under this formalization some of the nice properties of classical logics hold, we will verify that, on regular formulas, S-semantics and T-semantics coincide.

The game $\mathsf{G}(\mathcal{M}, \varphi, S, p, h)$ we are about to define deviates from the customary semantic game for IF-like logics: it is a one-turn game where Abélard and Eloïse pick functions instead of elements. There are two reasons for this. On the one hand, we prefer this formulation since in this way the higher-order nature of the logic becomes arguably more apparent. On the other, this game will be generalized to an n-turn game in Section 4 to provide natural game-theoretical semantics for Hodges' flattening operator.

Before we go into the definitions, we need some notation for the manipulation of functions (and, in particular, infinite sequences). Let $f : X \to Y$, we denote with $f[x \mapsto y]$ the function such that $f[x \mapsto y](x) = y$ and $f[x \mapsto y](z) = f(z)$ for all $z \neq x$. As usual, if $X' \subseteq X$ then $f \upharpoonright X' : X' \to Y$ will be the restriction of f to X'.

[2] Almost all of our presentation can probably be done using sequences of finite length. Apart from an arguably more cumbersome presentation, a downside of this would be that $s \circ p$ would then represent a classical valuation but one with finite image.

The board. The game is played over the syntactic tree of a formula. Every node of the tree, except the $\sim$-nodes, belong to one of the players: those initially under an even number of $\sim$-nodes belong to Eloïse, the rest belongs to Abélard. The initial assignment of nodes to a player will be remembered along the game. Furthermore, some nodes may be decorated with functions during the game: $\exists$-nodes can be decorated with any function $f : |\mathcal{M}|^\omega \to |\mathcal{M}|$; $\vee$-nodes can be decorated with any function $f : |\mathcal{M}|^\omega \to \{L, R\}$. Initially, these nodes have no decoration. Plus, there is a triple $\langle S, p, h\rangle$ and a placeholder (initially empty) for a sequence in $|\mathcal{M}|^\omega$.

The turn. The turn is composed of two clearly distinguished phases. In the first phase, both players decorate all their nodes with proper functions. The order in which they tag their nodes is not important as long as they don't get to see their opponent's choices in advance. For simplicity, we will assume they both play simultaneously. For the second phase, we introduce a third agent, sometimes known as Nature, that can be seen as random choices. Nature first picks some sequence from S and puts it in the placeholder. Next, it proceeds to *evaluate* the result of the turn using the following recursive procedure:

R1. If the tree is of the form $\sim\psi$, Nature replaces it with ψ and evaluation continues.

R2. If the tree is of the form $\psi_1 \vee/_{y_1,\ldots,y_k} \psi_2$, then its root must have been decorated with some $f : |\mathcal{M}|^\omega \to \{L, R\}$. Nature then picks a sequence $r \in |\mathcal{M}|^\omega$ such that $r(i) = s(i)$ for every $i \notin \{p(y_1), \ldots, p(y_n)\} \cup \{k \mid k \geq h\}$, where s stands for the sequence on the placeholder, and evaluates $f(r)$. Observe that the values the player was not supposed to consider are replaced with arbitrary values prior to evaluating the function. The tree then is replaced with ψ_1 if the result is L or with ψ_2 otherwise, and evaluation proceeds.

R3. If the tree is of the form $\exists x/_{y_1,\ldots,y_k}\psi$, then it must be decorated with some $f : |\mathcal{M}|^\omega \to |\mathcal{M}|$. Nature here also picks a sequence $r \in |\mathcal{M}|^\omega$ such that $r(i) = s(i)$ for every $i \notin \{p(y_1), \ldots, p(y_n)\} \cup \{k \mid k \geq h\}$, where s stands for the sequence on the placeholder, and evaluates $f(r)$. Let us call this value b. Nature records this choice by replacing the sequence in the placeholder with $s[h \mapsto b]$; x is bound to b by replacing p with $p[x \mapsto h]$ and h is incremented by one. Finally, the tree is replaced with ψ and evaluation proceeds.

R4. Finally, if the root of the tree is of the form $R(t_1, \ldots, t_k)$, evaluation ends. Eloïse is declared the winner of the match whenever this node belongs to her and $\mathcal{M} \models R(t_1, \ldots, t_k)[s \circ p]$, or the node belongs to Abélard and $\mathcal{M} \not\models R(t_1, \ldots, t_k)[s \circ p]$. In any other case, the winner is Abélard.

Winning strategies. A *strategy* for a player of the game $\mathsf{G}(\mathcal{M}, \varphi, S, p, h)$ is just the collection of functions used to decorate the syntactic tree of φ. Furthermore, the strategy is *winning* if it guarantees that the player will win every match of the game, regardless the strategy of the opponent and the choices made by Nature. Observe this game is of *imperfect information*: Abélard and Eloïse must play *simultaneously* (i.e. ignoring the opponent move) and the initial valuation is “randomly” picked by Nature. Therefore, some games are probably undetermined, that is, none of the players have a winning strategy.

We are now ready to give our game-semantic notion of *truth* and *falsity*. Observe, though, that this will be restricted to only certain p and h. The rationale for this will become clear later (cf. Example 1 and Lemma 1).

Definition 1. *We say that* $p : \mathrm{Vars} \to \omega$ *and* $h < \omega$ *are a* proper context *for a formula* φ *if* $p \restriction \mathrm{Fv}(\varphi)$ *is injective and* $\{p(x) \mid x \in \mathrm{Fv}(\varphi)\} \subseteq \{0, \ldots, h-1\}$.

Definition 2 ($\models^+$ and $\models^-$ for S-semantics). *Given a formula* φ*, a suitable model* $\mathcal{M}$*, a nonempty set* $S \subseteq |\mathcal{M}|^\omega$ *and a proper context for* φ*,* $p : \mathrm{Vars} \to \omega$ *and* $h < \omega$*, we define:* $\mathcal{M} \models^+ \varphi[S, p, h]$ *iff Eloïse has a winning strategy for* $\mathsf{G}(\mathcal{M}, \varphi, S, p, h)$*;* $\mathcal{M} \models^- \varphi[S, p, h]$ *iff Abélard has a winning strategy for* $\mathsf{G}(\mathcal{M}, \varphi, S, p, h)$.

When S is the singleton set $\{s\}$ we may alternatively write $\mathcal{M} \models^+ \varphi[s, p, h]$ and $\mathcal{M} \models^- \varphi[s, p, h]$. Furthermore, we will write $\mathcal{M} \models^+ \varphi$ if $\mathcal{M} \models^+ \varphi[|\mathcal{M}|^\omega, p, h]$ whenever p and h are a proper context for φ (and analogously for $\mathcal{M} \models^- \varphi$).

Example 1. Consider $\theta := \exists x\ [x \neq y]$. One would expect that for any $\mathcal{M}$ with at least two elements, $\mathcal{M} \models^+ \theta$ show hold. However, Eloïse has no winning strategy on $\mathsf{G}(\mathcal{M}, \theta, S, p, h)$ when $p(y) = h$. The problem here is that the value selected by Eloïse's function for x, whatever it is, will be recorded in position h, thus *overwriting* the value of y. Observe, though, that if p and h are a proper context for θ, then it cannot be the case that $p(y) \geq h$.

Example 2. Let us revisit the irregular formula θ_2 from Section 2.3. We shall verify that for any model $\mathcal{M}$, $\mathcal{M} \models^+ \theta_2$. For this, consider the following strategy for Eloïse: $f_\vee(s) = L$ if $s(h) = s(h+1)$ and $f_\vee(s) = R$ otherwise; $f_{\exists y}(s) = s(h+2)$; $f_{\exists u/_x}(s) = s(h+1)$. The reader should verify that this is essentially the same strategy used for θ_1 in Section 2. Observe that, for example, $s(h+1)$ plays the same role that $v(y)$ played in the latter, except that by using an offset from h (i.e., from the position in s where the value for the outermost quantifier was recorded) instead of the variable name, we escape from the deathtraps created by the rebinding of variables. In fact, Eloïse's winning strategy in this example works for any renaming of variables of θ_2.

So far we have defined $\models^+$ and $\models^-$ with respect to sets of sequences using a game theoretical approach. We can also give a compositional characterization, in the line of [11] and [4] (the proof of Theorem 1 is on Section A.3).

Definition 3. *Let* $f : A^B \to C$ *and let* $Y \subseteq B$*. We say that* f *is* Y*-independent if for all* $g_1, g_2 \in A^B$ *such that* $g_1(x) = g_2(x)$ *whenever* $x \notin Y$*,* $f(g_1) = f(g_2)$.

Theorem 1 (Compositionality of S-semantics). *Let* $\mathcal{M}$ *be a suitable model, let* $S \subseteq |\mathcal{M}|^\omega$ *be nonempty and let* $p : \mathrm{Vars} \to \omega$ *and* $h < \omega$ *be a proper context.*

1. $\mathcal{M} \models^+ R(t_1, \ldots, t_k)[S, p, h]$ *iff* $\mathcal{M} \models R(t_1, \ldots, t_k)[s \circ p]$ *for all* $s \in S$
2. $\mathcal{M} \models^- R(t_1, \ldots, t_k)[S, p, h]$ *iff* $\mathcal{M} \not\models R(t_1, \ldots, t_k)[s \circ p]$ *for all* $s \in S$
3. $\mathcal{M} \models^+ {\sim}\psi[S, p, h]$ *iff* $\mathcal{M} \models^- \psi[S, p, h]$
4. $\mathcal{M} \models^- {\sim}\psi[S, p, h]$ *iff* $\mathcal{M} \models^+ \psi[S, p, h]$
5. $\mathcal{M} \models^+ \psi_1 \vee/_{y_1,\ldots,y_k} \psi_2[S, p, h]$ *iff there is an* $f : S \to \{L, R\}$ *such that*
   - f *is* $\{p(y_1), \ldots, p(y_k)\} \cup \{k \mid k \geq h\}$*-independent;*
   - $\mathcal{M} \models^+ \psi_1[S_L, p, h]$*, where* $S_L = \{s \mid s \in S, f(s) = L\}$*; and*
   - $\mathcal{M} \models^+ \psi_2[S_R, p, h]$*, where* $S_R = \{s \mid s \in S, f(s) = R\}$
6. $\mathcal{M} \models^- \psi_1 \vee/_{y_1,\ldots,y_k} \psi_2[S, p, h]$ *iff* $\mathcal{M} \models^- \psi_1[S, p, h]$ *and* $\mathcal{M} \models^- \psi_1[S, p, h]$
7. $\mathcal{M} \models^+ \exists x/_{y_1,\ldots,y_k}\psi[S, p, h]$ *iff there is a function* $f : S \to |\mathcal{M}|$ *such that*
   - f *is* $\{p(y_1), \ldots, p(y_k)\} \cup \{k \mid k \geq h\}$*-independent; and*
   - $\mathcal{M} \models^+ \psi[\widetilde{S}, p[x \mapsto h], h + 1]$*, where* $\widetilde{S} = \{s[h \mapsto f(s)] \mid s \in S\}$
8. $\mathcal{M} \models^- \exists x/_{y_1,\ldots,y_k}\psi[S, p, h]$ *iff* $\mathcal{M} \models^- \psi[\widetilde{S}, p[x \mapsto h], h + 1]$ *for* $\widetilde{S} = \{s[h \mapsto a] \mid s \in S, a \in |\mathcal{M}|\}$

Definition 4 ($\equiv_h$ and $\equiv$). *Given* $h < \omega$*, we write* $\varphi_1 \equiv_h \varphi_2$ *if* $\mathrm{Fv}(\varphi_1) = \mathrm{Fv}(\varphi_2)$ *and for every suitable model* $\mathcal{M}$*, every nonempty* $S \subseteq |\mathcal{M}|^\omega$ *and every* $p : \mathrm{Vars} \to \omega$ *such that* p*,* h *is a proper context for* φ_1*,* $\mathcal{M} \models^\pm \varphi_1[S, p, h]$ *iff* $\mathcal{M} \models^\pm \varphi_2[S, p, h]$*. Furthermore, if for every* $h < \omega$ *we have* $\varphi_1 \equiv_h \varphi_2$ *, then we say that the formulas are S-*equivalent*, notated* $\varphi_1 \equiv \varphi_2$*.*

Since strategies for the game $\mathsf{G}(\mathcal{M}, \varphi, S, p, h)$ must deal with sequences but not with variable values, it is straightforward to verify the following:

Proposition 2. *If* $\varphi_1 \equiv_\alpha \varphi_2$ *then* $\varphi_1 \equiv \varphi_2$*.*

In first-order logic, the truth of a formula depends only on the value of its free variables (i.e., if $\mathcal{M} \models \varphi[v]$ and v and v' differ only on variables that are not free in φ, then $\mathcal{M} \models \varphi[v']$). We will show next that in our setting, there are three operations on valuations that preserve satisfaction. In what follows, for $S \subseteq A^\omega$, we define $S \restriction n = \{(s(0), \ldots, s(n-1)) \mid s \in S\}$; we call h-permutation to any bijective function $\pi : \omega \to \omega$ such that $\pi(i) = i$ for all $i \geq h$; and $S \circ \pi = \{s \circ \pi \mid s \in S\}$.

Theorem 2. *For all suitable* $\mathcal{M}$*, nonempty* $S \subseteq |\mathcal{M}|^\omega$ *and proper contexts for* φ*,* p *and* h*:*

1. *If* $\widetilde{p} \restriction \mathrm{Fv}(\varphi) = p \restriction \mathrm{Fv}(\varphi)$*, then* $\mathcal{M} \models^\pm \varphi[S, p, h]$ *iff* $\mathcal{M} \models^\pm \varphi[S, \widetilde{p}, h]$*.*
2. *If* $\widetilde{S}$ *is such that* $\widetilde{S} \restriction h = S \restriction h$*, then* $\mathcal{M} \models^\pm \varphi[S, p, h]$ *iff* $\mathcal{M} \models^\pm \varphi[\widetilde{S}, p, h]$*.*
3. *If* π *is an* h*-permutation then* $\mathcal{M} \models^\pm \varphi[S, p, h]$ *iff* $\mathcal{M} \models^\pm \varphi[S \circ \pi, \pi \circ p, h]$*.*

We are now ready to show that, when restricted to regular formulas, the equivalence notions of S-semantics and T-semantics match. Of course, this implies that the set of valid (regular) formulas of both logics is the same and, because of Proposition 2, S-semantics is a proper generalization of T-semantics (the proof is on Section A.4).

Theorem 3. *Let* φ_1 *and* φ_2 *be regular formulas. Then* $\varphi_1(x_0, \ldots, x_{h-1}) \equiv_T \varphi_2(x_0, \ldots, x_{h-1})$ *iff* $\varphi_1 \equiv_h \varphi_2$*.*

4 Game Theoretical Semantics for IF with Flattening

We say that φ is *true* in $\mathcal{M}$ when $\mathcal{M} \models^+ \varphi$ and that is *false* if $\mathcal{M} \models^- \varphi$. Clearly, if $\mathcal{M} \models^+ \varphi$ then $\mathcal{M} \not\models^- \varphi$, and if $\mathcal{M} \models^- \varphi$ then $\mathcal{M} \not\models^+ \varphi$. However there are *sentences* which may be neither true nor false in a model. Hodges considers the problem of adding classical negation to slash logic. He wants, for instance, $\mathcal{M} \models^\pm \neg\varphi$ iff $\mathcal{M} \not\models^\pm \varphi$ to hold; restoring, for sentences, the identity between being not-true and being false. To this end, he introduces the *flattening operator* $\downarrow$, and stipulates $\neg\psi \equiv {\sim}{\downarrow}\psi$ [11].

Since in this section we move to slash logic enriched with the flattening operator, we assume from here on that $\downarrow$ may occur freely in a formula. First of all, we need to specify its semantics. Hodges used a compositional definition; therefore, we will take Theorem 1 to be a compositional definition of $\models^+$ and $\models^-$ for slash logic and extend it to handle $\downarrow$. Observe we are simply adapting his notation according to our presentation.

Definition 5 ($\models^+$ and $\models^-$ for S-semantics with $\downarrow$). *We define* $\models^+$ *and* $\models^-$ *as the relation induced by clauses 1–8 of Theorem 1, plus*

9. $\mathcal{M} \models^+ {\downarrow}\psi[S,p,h]$ *iff* $\mathcal{M} \models^+ \psi[s,p,h]$ *for every* $s \in S$
10. $\mathcal{M} \models^- {\downarrow}\psi[S,p,h]$ *iff* $\mathcal{M} \not\models^+ \psi[s,p,h]$ *for every* $s \in S$

Hodges seems to suggest that no natural game-theoretical semantics can be given for this logic[3]. In any case, this can indeed be done. We define next the game $\mathsf{G}_\downarrow(\mathcal{M},\varphi,S,p,h)$, which extends the rules of the game described in Section 3 to deal with formulas containing arbitrary occurrences of $\downarrow$.

The board. The board is essentially the same one used for $\mathsf{G}(\mathcal{M},\varphi,S,p,h)$. The syntactic tree of the formula now may contain $\downarrow$-nodes; these are assigned to players using the same criteria: those under an even number of $\sim$-nodes belong to Eloïse, the remaining ones to Abélard. Just like the leafs of the tree, $\downarrow$-nodes will not be decorated.

The turns. Unlike the one of Section 3, this game may last more than one turn. At any point of the game, the remaining number of turns will be bounded by the number of nested occurrences of $\downarrow$-nodes in the game-board. The opening turn is played exactly like in Section 3, although we still need to stipulate what happens, during the evaluation phase, if Nature arrives to a formula of the form ${\downarrow}\psi$. Observe that this means that if no $\downarrow$ occurs in φ, then $\mathsf{G}(\mathcal{M},\varphi,S,p,h)$ and $\mathsf{G}_\downarrow(\mathcal{M},\varphi,S,p,h)$ are essentially the same game.

So, summing up, when the game starts, both players decorate their nodes simultaneously; then Nature picks a sequence and puts it in the placeholder, and finally starts the evaluation phase (cf. rules **R1**–**R4** in Section 3). If evaluation reaches a leaf (i.e., an atom), then the game ends, and the winner is determined according to rule **R4**. For the extra case we add the following rule:

R5. If the tree is of the form ${\downarrow}\psi$, then the turn ends.

[3] The exact quote is: "In the presence of $\downarrow$, we can't define a game $G(\phi, A)$ for arbitrary A and ϕ." [11, p. 556].

The initial turn differs slightly from the subsequent ones, where the formula on the board will be always of the form $\downarrow\psi$. Now both players get to redecorate their nodes, except that in this case, they proceed one after the other. The player who owns the $\downarrow$-node at the root gets to do it first. After this, Nature replaces the tree with ψ and proceeds to the evaluation phase following rules **R1**–**R5**.

We won't go into a formal description of a winning strategy for this game. We simply take it to be some form of oracle that, when followed, guarantees that the game ends in a winning position.

Theorem 4 (Game semantics for $\downarrow$). *Given a formula φ, a suitable $\mathcal{M}$, a nonempty $S \subseteq |\mathcal{M}|^\omega$ and a proper context for φ, $\langle p, h\rangle$, the following holds: $\mathcal{M} \models^+ \varphi[S, p, h]$ iff Eloïse has a winning strategy for $\mathsf{G}_\downarrow(\mathcal{M}, \varphi, S, p, h)$; $\mathcal{M} \models^- \varphi[S, p, h]$ iff Abélard has a winning strategy for $\mathsf{G}_\downarrow(\mathcal{M}, \varphi, S, p, h)$.*

5 Conclusions

We think that invariance under α-equivalence is a property that no sane formalism can disregard. By decoupling *values* from *name for values* we have been able to successfully generalize Hodges' T-semantics from regular formulas to unrestricted ones. To achieve this we had to pay a small price: abandon the well-established use of classical valuations.

In [3], Caicedo, Dechesne and Janssen took a different path and investigated a weaker notion of equivalence for V-semantics. They say, for instance, that $\varphi \equiv_{xz} \psi$ if $\{x, y\} \not\subseteq \mathrm{Fv}(\varphi) \cup \mathrm{Fv}(\psi)$ and $\mathcal{M} \models^\pm \varphi[V]$ iff $\mathcal{M} \models^\pm \psi[V]$, *provided that x and z are excluded from the domain of the valuations in V* and go into great technical efforts to properly characterize the normal form equivalences initially presented in [4]. We believe this route leads ultimately to a dead-end: V-semantics are buying very little and are too hard to reason about.

We favor a simpler approach, akin to the usual practice in classical logics. For convenience, stick to regular formulas, use a lightweight formalism, like T-semantics, and finally resort to Theorem 3 and Propositions 1 and 2 to generalize the result. This way, for instance, the normal forms results of [4] can easily be shown to hold under S-semantics, in a much more general way than in [3].

In the last part of the paper we looked at the $\downarrow$-operator from a novel game-theoretical perspective. We believe this will ultimately help to gain more insight about this operator.

Acknowledgments. This work was partially supported by grant PICTR 0172 from the Agencia Nacional de Promoción Científica y Tecnológica (Argentina); by grant G.0344.05 from the Research Foundation Flanders (FWO-Vlaanderen) and by the European Union under the FP6-IST-FET programme, Project n. FP6-14915, GeoPKDD: Geographic Privacy-Aware Knowledge Discovery and Delivery. D. Gorín is partially supported by a grant from CONICET, Argentina.

References

1. Barendregt, H.: The Lambda Calculus: its syntax and semantics, 2nd edn. Studies in logic and the foundations of mathematics, vol. 103. North Holland, Amsterdam (1985)
2. Barwise, J.: On branching quantifiers in English. Journal of Philosophical Logic 8, 47–80 (1979)
3. Caicedo, X., Dechesne, F., Jansen, T.M.V.: Equivalence and quantifier rules for logic with imperfect information. Prepublication Series PP-2007-20. ILLC (2007)
4. Caicedo, X., Krynicki, M.: Quantifiers for reasoning with imperfect information and Σ_1^1-logic. Contemporary Mathematics 235, 17–31 (1999)
5. de Bruijn, N.G.: Lambda calculus notation with nameless dummies. A tool for automatic formula manipulation with application to the Church-Rosser theorem. Indagationes Mathematicae 34, 381–392 (1972)
6. Dechesne, F.: Game, Sets, Maths: formal investigations into logic with imperfect information. PhD thesis, Department of Philosophy, University of Tilburg, The Netherlands (2005)
7. Enderton, H.: Finite partially ordered quantifiers. Zeitschrift für Mathematische Logik und Grundlagen der Mathematik 16, 393–397 (1970)
8. Henkin, L.: Some remarks on infinitely long formulas. In: Infinitistic Methods: Proceedings of the Symposium on Foundations of Mathematics, pp. 167–183. Pergamon Press, Oxford (1961)
9. Hintikka, J.: The Principles of Mathematics Revisited. Cambridge University Press, Cambridge (1996)
10. Hintikka, J., Sandu, G.: Game-theoretical semantics. In: van Benthem, J., ter Meulen, A. (eds.) Handbook of logic and language, ch. 6. MIT Press, Cambridge (1997)
11. Hodges, W.: Compositional semantics for a language of imperfect information. Logic Journal of the IGPL 5(4) (1997)
12. Hodges, W.: Logics of imperfect information: why sets of assignments? In: Proceedings of the 7th Augustus de Morgan Workshop Interactive Logic: Games and Social Software (2005)
13. Hodges, W.: Some strange quantifiers. In: Mycielski, J., Rozenberg, G., Salomaa, A. (eds.) Structures in Logic and Computer Science. LNCS, vol. 1261, pp. 51–65. Springer, Heidelberg (1997)
14. Janssen, T.M.V.: Independent choices and the interpretation of IF logic. Journal of Logic, Language and Information 11(3), 367–387 (2002)
15. Janssen, T.M.V., Dechesne, F.: Signalling in IF games: a tricky business. In: The age of alternative logics, ch. 15, pp. 221–241. Springer, Heidelberg (2006)
16. Walkoe Jr., W.: Finite partially-ordered quantification. Journal of Symbolic Logic 35(4), 535–555 (1970)

A Technical Appendix

A.1 Compositionality for T-Semantics

For completeness, we introduce the compositional formulation of T-semantics. Our presentation is closer to the one due to Caicedo and Krynicki [4], but they can easily be shown to be equivalent.

In the following definition, if $t = (t_1, \ldots, t_n) \in |\mathcal{M}|^n$, then v_t : Vars $\to |\mathcal{M}|$ stands for any classical first order valuation such that $v(x_i) = t_i$ for $i = 1, \ldots, n$. For a definition of Y-independence, refer to Definition 3 on page 170.

Definition 6 (Compositionality of T-semantics). *Let $\overline{x} = x_1, \ldots, x_n$, let $\psi(\overline{x})$ be a regular formula, let $\mathcal{M}$ be a suitable model and let $T \subseteq |\mathcal{M}|^n$ be a set of deals of length n. We define $\models^+$ and $\models^-$ as follows:*

1. *if $\varphi(\overline{x})$ is atomic or negated atomic,*
   - $\mathcal{M} \models^+ \varphi(\overline{x})[T]$ *iff* $\mathcal{M} \models \varphi(\overline{x})[v_t]$ *for all* $t \in T$
   - $\mathcal{M} \models^- \varphi(\overline{x})[T]$ *iff* $\mathcal{M} \not\models \varphi(\overline{x})[v_t]$ *for no* $t \in T$
2. *if $\varphi(\overline{x}) = {\sim}\psi(\overline{x})$,*
   - $\mathcal{M} \models^+ \varphi(\overline{x})[T]$ *iff* $\mathcal{M} \models^- \psi(\overline{x})[T]$
   - $\mathcal{M} \models^- \varphi(\overline{x})[T]$ *iff* $\mathcal{M} \models^+ \psi(\overline{x})[T]$
3. *if $\varphi(\overline{x}) = \psi_1(\overline{x}) \vee/_{x_{n_1},\ldots,x_{n_k}} \psi_2(\overline{x})$ for some $\{n_1, \ldots, n_k\} \subseteq \{1, \ldots, n\}$,*
   - $\mathcal{M} \models^+ \varphi(\overline{x})[T]$ *iff there is a function* $g : T \to \{L, R\}$ *such that*
     - g *is* $\{n_1, \ldots, n_k\}$*-independent;*
     - $\mathcal{M} \models^+ \psi_1(\overline{x})[T_L]$, *where* $T_L = \{t \mid t \in T, g(t) = L\}$*; and*
     - $\mathcal{M} \models^+ \psi_2(\overline{x})[T_R]$, *where* $T_R = \{t \mid t \in T, g(t) = R\}$
   - $\mathcal{M} \models^- \varphi(\overline{x})[T]$ *iff* $\mathcal{M} \models^- \psi_1(\overline{x})[T]$ *and* $\mathcal{M} \models^- \psi_1(\overline{x})[T]$
4. *if $\varphi(\overline{x}) = \exists y/_{x_{n_1},\ldots,x_{n_k}} \psi(\overline{x}, y)$ and $y \notin \{x_1, \ldots, x_n\}$ for some $\{n_1, \ldots, n_k\} \subseteq \{1, \ldots, n\}$,*
   - $\mathcal{M} \models^+ \varphi(\overline{x})[T]$ *iff there is a function* $g : T \to |\mathcal{M}|$ *such that*
     - g *is* $\{n_1, \ldots, n_k\}$*-independent; and*
     - $\mathcal{M} \models^+ \psi(\overline{x}, y)[T']$, *where*

$$T' = \{(t_1, \ldots, t_n, g(t_1, \ldots, t_n)) \mid (t_1, \ldots, t_n) \in T\}$$

   - $\mathcal{M} \models^- \varphi(\overline{x})[T]$ *iff* $\mathcal{M} \models^- \psi(\overline{x}, y)[T']$ *for*

$$T' = \{(t_1, \ldots, t_n, a) \mid (t_1, \ldots, t_n) \in T, a \in |\mathcal{M}|\}$$

A.2 Proof of Proposition 1

Let $\overline{x} = x_1, \ldots, x_n$. It is enough to consider the case where $\varphi_1(\overline{x})$ is equal to $\varphi_2(\overline{x})$ except that the bound variable u of $\varphi_1(\overline{x})$ is replaced by v in $\varphi_2(\overline{x})$, within the scope of the same quantifier. By the Full Abstraction Theorem for T-semantics [11, Theorem 7.6] it suffices to prove that $\exists u/_{x_{n_1},\ldots,x_{n_k}} \psi_1(\overline{x}, u) \equiv_T \exists v/_{x_{n_1},\ldots,x_{n_k}} \psi_2(\overline{x}, v)$, where $\psi_2(\overline{x}, v)$ is obtained from $\psi_1(\overline{x}, v)$ when replacing the free variable u with v. One can prove by induction that $\psi_1(\overline{x}, v) \equiv_T \psi_2(\overline{x}, u)$. The key point here is that at item 1 of Definition 6 the name u or v is irrelevant, as long as they came in the same order in the lists $(\overline{x}, u)$ and $(\overline{x}, v)$.

A.3 Proof of Theorem 1

For the right-to-left implication, one proceeds by structural induction and shows that, for the $\exists$ and $\vee$ cases, the function f plus the strategy for the subformula(s) constitute a winning strategy. For the left-to-right implication, one only needs to see that if a player has a winning strategy on the game $\mathsf{G}(\mathcal{M}, \varphi, S, p, h)$, then he also has a winning strategy where all the functions that constitute it satisfy the independence restriction, and this is relatively straightforward (the full details can be seen, e.g., in [3, Theorems 4.7 and 4.8]). In every case, one also has to check that contexts are proper, but this is trivial.

A.4 Proof of Theorem 3

We first need to establish the following lemma.

Lemma 1. *Let $\varphi(x_0, \ldots, x_{h-1})$ be a regular formula such that in every branch of its syntactic tree, variables are bound in the same order. Furthermore, let $p : \mathrm{Vars} \to \omega$ be such that $p(x_i) = i$ for $0 \leq i < h$. Then $\mathcal{M} \models^{\pm} \varphi[S, p, h]$ iff $\mathcal{M} \models^{\pm} \varphi(x_0 \ldots x_{h-1})[S \restriction h]$.*

Proof. Let $\overline{x} = x_0, \ldots, x_{h-1}$. Suppose the list of occurrences of bound variables appearing in each branch of the syntactic tree of $\varphi(\overline{x})$ (from the root to the leaves) is a prefix of $x_h, x_{h+1}, x_{h+2}, \ldots$ The proof goes by induction in the complexity of φ. The atomic and negation are straightforward. Let us analyze the case $\varphi = \exists x_h/_{x_{n_1},\ldots,x_{n_k}} \psi(\overline{x}, x_h)$, for some $\{n_1, \ldots, n_k\} \subseteq \{0, \ldots, h-1\}$.

For the left to right implication, suppose $\mathcal{M} \models^{+} \varphi(\overline{x})[S, p, h]$. By Theorem 1 (item 7), there is a function $f : S \to |\mathcal{M}|$ such that f is $\{p(x_{n_1}), \ldots, p(x_{n_k})\} \cup \{k \mid k \geq h\}$-independent and $\mathcal{M} \models^{+} \psi(\overline{x}, x_h)[S', p[x_h \mapsto h], h+1]$, where $S' = \{s[h \mapsto f(s)] \mid s \in S\}$. Since $p = p[x_h \mapsto h]$, by inductive hypothesis we get $\mathcal{M} \models^{+} \psi(\overline{x}, x_h)[S' \restriction h+1]$. Fix $z \in S$ and define $g : S \restriction h \to |\mathcal{M}|$ as $g(s_0, \ldots, s_{h-1}) = f(s_0 \ldots s_{h-1} z(h) z(h+1) \ldots)$ for every $(s_0, \ldots, s_{h-1}) \in S \restriction h$. Since f is $\{n_1, \ldots, n_k\} \cup \{k \mid k \geq h\}$-independent then g is clearly well-defined and $\{n_1, \ldots, n_k\}$-independent. Furthermore,

$$S' \restriction h+1 = \{(s_0, \ldots, s_{h-1}, g(s)) \mid (s_0, \ldots, s_{h-1}) \in S \restriction h\}.$$

By Definition 6 (item 4), $\mathcal{M} \models^{+} \varphi(\overline{x})[S \restriction h]$.

For the other direction, suppose $\mathcal{M} \models^{+} \varphi(\overline{x})[S \restriction h]$. By Definition 6 (item 4), there exists some function $g : S \restriction h \to |\mathcal{M}|$ that is $\{n_1, \ldots, n_k\}$-independent and such that $\mathcal{M} \models^{+} \psi(\overline{x}, x_h)[T']$, where

$$T' = \{(t_1, \ldots, t_h, g(t_1, \ldots, t_h)) \mid (t_1, \ldots, t_h) \in S \restriction h\}.$$

Observe that $T' = S' \restriction h+1$, where

$$S' = \{s[h \mapsto g(s(0), \ldots, s(h-1))] \mid s \in S\}.$$

By inductive hypothesis and the fact that $p[x_h \mapsto h] = p$, we have $\mathcal{M} \models^+ \psi(\overline{x}, x_h)[S', p[x_h \mapsto h], h+1]$. Define $f : S \to |\mathcal{M}|$ as $f(s) = g(s(0), \ldots, s(h-1))$ for $s \in S$. By definition, f is clearly $\{k \mid k \geq h\}$-independent, and since g is $\{n_1, \ldots, n_k\}$-independent, f also is. By Theorem 1 (item 7) we conclude $\mathcal{M} \models^+ \varphi(\overline{x})[S, p, h]$.

The case for $\models^-$ and $\varphi = \exists x_h/_{x_{n_1},\ldots,x_{n_k}} \psi(\overline{x}, x_n)$ is straightforward. A similar argument can be used for the case $\varphi(\overline{x}) = \psi_1(\overline{x}) \vee/_{x_{n_1},\ldots,x_{n_k}} \psi_1(\overline{x})$.

We are now ready to prove the theorem. We will only show it for $\models^+$, the argument for $\models^-$ is similar. In what follows $\overline{x}$ will stand for $x_0, \ldots, x_{h-1}$. From left to right, by the counterpositive, suppose that $\mathcal{M} \models^+ \varphi_1(\overline{x})[S, p, h]$ and $\mathcal{M} \not\models^+ \varphi_2(\overline{x})[S, p, h]$, for some suitable model $\mathcal{M}$ and some $p : \text{Vars} \to \omega$ such that p, h is a proper context for φ_1 (and for φ_2, since $\text{Fv}(\varphi_1) = \text{Fv}(\varphi_2)$). One can build an h-permutation π such that $\pi(p(x_i)) = i$ for $0 \leq i < h$ and using Theorem 2 one gets $\mathcal{M} \models^+ \varphi_1(\overline{x})[S \circ \pi, \pi \circ p, h]$ but $\mathcal{M} \not\models^+ \varphi_2(\overline{x})[S \circ \pi, \pi \circ p, h]$. By Proposition 2, we can pick regular $\varphi_1' \equiv_\alpha \varphi_1$ and $\varphi_2' \equiv_\alpha \varphi_2$ where variables are bound in the same order on every branch of their syntactic trees and, using Lemma 1 we obtain $\mathcal{M} \models^+ \varphi_1'(\overline{x})[S \circ \pi \restriction h]$ and $\mathcal{M} \not\models^+ \varphi_2'(\overline{x})[S \circ \pi \restriction h]$, which implies $\varphi_1(\overline{x}) \not\equiv_T \varphi_2(\overline{x})$ using Proposition 1.

From right to left, suppose $\varphi_1(\overline{x}) \not\equiv_T \varphi_2(\overline{x})$, i.e., $\mathcal{M} \models^+ \varphi_1(\overline{x})[T]$ and $\mathcal{M} \not\models^+ \varphi_2(\overline{x})[T]$, for some suitable model $\mathcal{M}$ and some trump $T \subseteq |\mathcal{M}|^h$. Define

$$S = \{t_1 \cdots t_h s \mid (t_1, \ldots, t_h) \in T, s \in |\mathcal{M}|^\omega\}$$

and $p(x_i) = i$. Again, using invariance under α-equivalence and Lemma 1 we conclude $\mathcal{M} \models^+ \varphi_1[S, p, n]$ and $\mathcal{M} \not\models^+ \varphi_2[S, p, n]$.

A.5 Proof of Theorem 4

The proof goes by induction on φ and is, essentially equivalent to the one for Theorem 1 except that we also have to account for the case where φ is $\downarrow\psi$. Suppose first $\mathcal{M} \models^+ \downarrow\psi[S, p, h]$; this means that $\mathcal{M} \models^+ \psi[s, p, h]$ for all $s \in S$. We want to construct a winning strategy for Eloïse for the game $\mathsf{G}_\downarrow(\mathcal{M}, \downarrow\psi, S, p, h)$. The first turn is irrelevant; for the second one, Eloïse simply has to consider the valuation s in the placeholder and use the winning strategy for $\mathsf{G}_\downarrow(\mathcal{M}, \psi, \{s\}, p, h)$ that, by inductive hypothesis, she has. For the other direction, suppose Eloïse has a winning strategy for $\mathsf{G}_\downarrow(\mathcal{M}, \downarrow\psi, S, p, h)$. This implies she has a winning strategy for $\mathsf{G}_\downarrow(\mathcal{M}, \psi, \{s\}, p, h)$ for all $s \in S$: play whatever she would play as her second turn in $\mathsf{G}_\downarrow(\mathcal{M}, \downarrow\psi, S, p, h)$ if Nature happened to pick s. By inductive hypothesis, this means $\mathcal{M} \models^+ \psi[s, p, h]$ for all $s \in S$ and, thus, $\mathcal{M} \models^+ \downarrow\psi[S, p, h]$.

Suppose now $\mathcal{M} \models^- \downarrow\psi[S, p, h]$; then for every $s \in S$, $\mathcal{M} \not\models^+ \psi[s, p, h]$. From here we derive a winning strategy for Abélard on $\mathsf{G}_\downarrow(\mathcal{M}, \downarrow\psi, S, p, h)$ as follows. The first turn is irrelevant; for the second one, an $s \in S$ has been picked and Eloïse has played first following some strategy. Observe that this strategy is also a possible strategy for $\mathsf{G}_\downarrow(\mathcal{M}, \psi, \{s\}, p, h)$. But by inductive hypothesis, since $\mathcal{M} \not\models^+ \psi[s, p, h]$, it cannot be a winning strategy for this game, i.e. Abélard has some strategy that defeats hers. Abélard simply has to use this strategy from this point on and will win the game. Analogously, if Abélard has a winning strategy for $\mathsf{G}_{\downarrow}(\mathcal{M}, {\downarrow}\psi, S, p, h)$, then for every $s \in S$ picked by Nature and any strategy followed by Eloïse, there is a way in which Abélard can play and win the game. But this means that for no $s \in S$, Eloïse has a winning strategy for $\mathsf{G}_{\downarrow}(\mathcal{M}, \psi, \{s\}, p, h)$ and, thus, by inductive hypothesis, $\mathcal{M} \not\models^{+} \psi[s, p, h]$ and, finally, $\mathcal{M} \models^{-} {\downarrow}\psi[S, p, h]$.