On the formal semantics of IF-like logics

Santiago Figueira [a,b,*], Daniel Gorín [a], Rafael Grimson [c,d]

[a] *Departamento de Computación, FCEyN, Universidad de Buenos Aires, Argentina*
[b] *CONICET, Argentina*
[c] *Departamento de Matemática, FCEyN, Universidad de Buenos Aires, Argentina*
[d] *Theoretical Computer Science Group, Hasselt University and Transnational University of Limburg, Belgium*

A R T I C L E   I N F O

A B S T R A C T

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 Independence Friendly (IF) logics. In this paper we argue that this is not an inherent characteristic of these logics but a defect in the way in which the compositional semantics given by Hodges for the *regular* fragment was generalized to arbitrary formulas. We fix this by proposing an alternative formalization, based on a variation of the classical notion of valuation. Basic metatheoretical results are proven. We present these results for Hodges' *slash logic* (from which these can be easily transferred to other IF-like logics) and we also consider the flattening operator, for which we give novel game-theoretical semantics.

© 2009 Elsevier Inc. All rights reserved.

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 [1,2]. Closely related to Henkin's logic of *branching quantifiers* [3–6], IF is an extension of first-order logic where each disjunction and each existential quantifier may be decorated with denotations of universally-quantified variables:

$$\forall x \forall y \exists z /_{\forall y} \exists w /_{\forall y} [y = z \lor /_{\forall x, \forall y} w = y]. \tag{1}$$

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 a 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*. In terms of expressive power, IF corresponds to the Σ_1^1 fragment of second-order logic [7], although the set of valid IF sentences is recursively isomorphic to that of full second-order logic [8].

It was conjectured by Hintikka that one could not formulate IF semantics in a *composable* way [1]. This was promptly rebutted by Hodges in [9], 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. In doing this, Hodges also provided formal semantics for non-sentences.

Two things are worth observing. First, in [9] 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. In Hintikka's presentation [1], no decorations with exis-

---

* Corresponding author at: Departamento de Computación, FCEyN, Universidad de Buenos Aires, Argentina.
  E-mail addresses: santiago@dc.uba.ar (S. Figueira), dgorin@dc.uba.ar (D. Gorín), rgrimson@dm.uba.ar (R. Grimson).

tentially quantified variables occur. Hintikka adopts a convention, which remains implicit in the syntax, that Eloïse is not allowed to take into account any of her previous choices. For implications of this fact see, e.g. [10].

For instance, formula (1) would be written in Hodges formulation as:

$$\forall x \forall y \exists z /_y \exists w /_{y,z} [y = z \vee /_{x,y,w,z} w = y].$$

Hodges later coined the name *slash logic* and stated [11]:

> During recent years many writers in this area (but never Hintikka himself) have transferred the name 'IF logic' to slash logic, often without realizing the difference. Until the terminology settles down, we have to beware of examples and proofs that don't make clear which semantics they intend.

We will use the term IF-*like logics* to encompass this variety of related logics. In [12], 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 a syntactic fragment called *regular* in which there can be no nested binding of the same variable (i.e. along one branch of the syntactic tree) nor a variable that occurs both freely and bound in one formula.[1] This is a standard practice in first-order logic model theory. The underlying assumption is that, given an arbitrary formula φ, one can appropriately rename its variables, so no generality is loss. As it was later pointed-out, this assumption is not satisfied in Hintikka's original semantics, and in fact it is not clear whether it is reasonable in the whole context of IF-like logics [13,14].

Historically, Caicedo and Krynicki [15] were the first to attempt a generalization of Hodges' semantics to non-regular formulas. In their paper they give a prenex normal form theorem for slash logic (called $\mathcal{L}_{ii}$ in their paper). Instead of n-tuples, they used valuations to account for arbitrary formulas where variables occur in any order and may even get rebound. This extension seemed so natural that in later papers (e.g. [10,14,16]) it was taken as the standard semantics of slash logic.

Based on this semantics, in [10], Janssen points-out several strange properties of these logics. At the root of them lies the idea of *signaling*. Signaling is "the phenomenon that the value of a variable one is supposed not to know, is available through the value of another variable" [14]. He observes, for instance, that the quantifier-context where a subformula occurs is of vital importance. Already in [9], it was observed (with the formula $\forall x \exists z \exists y /_x [x = y]$) that a seemingly vacuous existential quantifier[2] may be, due to signaling, the difference between a formula being true or not. Janssen [10] points out that if variables are reused in a formula, signaling may be blocked and, thus, the truth-value of formulas that only differ on bound-variables may vary. 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 [1].

A systematic analysis of signaling in IF-like logics was later performed in [14], where several claims of "equivalence of formulas under syntactic transformations" made in [15] are questioned due to signalings that may get unexpectedly blocked. These results are fixed in [16], although the original semantics had to be subtly modified (valuations are required to have finite domain) and equivalences are shown to hold only on certain contexts, depending on the variables occurring in the formulas.

Summing up, on the one hand, we have a family of logics, aiming to be a conservative extension of first-order logic, for which many results have been proved, but 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* (cf. Section 7.5 of [13]).

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 valuations as functions from variables to values are simply not adequate to formalize independence restrictions in a context where variables can get rebound.

The paper runs as follows. In Section 2 we discuss in more detail the misbehavior of irregular formulas with respect to signaling. We will also see that even when restricted to the regular fragment, the valuation-based semantics fails to generalize Hodges' compositional semantics. This motivates Section 3, were we first observe that, according to some results of Section 2, there must be at least two different notions of *independence* for regular IF-like logic, one followed by [15,16] and the other by [9]. We then give both game-theoretical and compositional semantics for slash logic carefully avoiding these pitfalls. In Section 4 we move to slash logic with the *flattening* operator $\downarrow$ introduced by Hodges in [9] and illustrate that irregular formulas can be handled uniformly also in this setting; while doing this, we provide a novel game semantics for this logic.

---

[1] Hodges does not explicitly restrict to the regular fragment, however this is implicitly done in his work by means of his formalism. Indeed, writing $\varphi(x_1, \ldots, x_n)$ implies that the free variables of φ (according to our definition of Fv) occur in the list $x_1, \ldots, x_n$, and he prevents this list from having duplications [9, p. 559]. Then he defines the compositional semantics of $\varphi(x_1, \ldots, x_n) = \exists x /_{y_1, \ldots, y_k} \psi$ in terms of $\psi(x_1, \ldots, x_n, x)$, and when writing ψ in this way, he is implicitly saying that x cannot be one of $x_1, \ldots, x_n$.

[2] By vacuous we mean that the bound variable does not occur in the subformula.

2. Preliminaries

2.1. Syntax

From here on, we restrict ourselves to Hodges' *slash logic* (but without indexed disjunctions) [9,11], 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 [16] we don't impose any restriction on the variables occurring under the slashes. Observe that $\sim$ is *game negation*; in Section 4 we will consider also the flattening operator $\downarrow$ and, thus, classical negation $\neg$.

The sets of free and bound variables of φ, $Fv(\varphi)$ and $Bv(\varphi)$ respectively, are defined in the usual way:

- if φ is atomic, then $Fv(\varphi)$ is the set of all free variables occurring in φ and $Bv(\varphi) = \emptyset$;
- if $\varphi = {\sim}\psi$, then $Fv(\varphi) = Fv(\psi)$ and $Bv(\varphi) = Bv(\psi)$;
- if $\varphi = \psi \vee /_{y_1,\ldots,y_k}\rho$, then $Fv(\varphi) = Fv(\psi) \cup Fv(\rho) \cup \{y_1,\ldots,y_k\}$ and $Bv(\varphi) = Bv(\psi) \cup Bv(\rho)$;
- if $\varphi = \exists x/_{y_1,\ldots,y_k}\psi$, then $Fv(\varphi) = (Fv(\psi) \setminus \{x\}) \cup \{y_1,\ldots,y_k\}$ and $Bv(\varphi) = Bv(\psi) \cup \{x\}$.

Observe that variables that occur under slashes must be taken into consideration. For example, that if $\theta := \exists x/_{x,y}[x = z]$ then $Fv(\theta) = \{x, y, z\}$ and $Bv(\theta) = \{x\}$.

Following Dechesne [13], we will say that a formula φ is *regular* whenever $Fv(\varphi) \cap Bv(\varphi) = \emptyset$ 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 $Fv(\varphi) \subseteq \{x_1,\ldots,x_n\}$. Observe that this means that for a fixed φ, $\varphi(x, y)$ and $\varphi(x, y, z)$ will denote two non-equivalent formulas. See [9] 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 [15,16].

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 $G(\mathcal{M}, \varphi, v)$ (for $v \in V$) of the classical satisfaction game between Abélard and Eloïse (sometimes called *Falsifier* and *Verifier*). Dually, a formula is false, written $\mathcal{M} \models^- \varphi[V]$, whenever Abélard has a valid strategy that is winning for every $G(\mathcal{M}, \varphi, v)$, $v \in V$.

For a strategy to be valid, it has to satisfy additional independence conditions. For a formal presentation refer to [15,16]. We illustrate these requirements with an example. Let us say that at some point of a game Eloïse must play in a position given by the formula $\exists x/_{y,z}\varphi$; a valid strategy function for this point would be an f mapping valuations to elements of the domain, with the restriction that $f(v) = f(v')$ whenever v and v' differ, at most, on their values for y and z. We say that this function f is $\{y, z\}$-independent.

Definition 1. Let $f : A^B \to C$ and let $Y \subseteq B$. We say that f is *Y-independent* if for all $g_1, g_2 \in A^B$: if $g_1(x) = g_2(x)$ for all $x \notin Y$ then $f(g_1) = f(g_2)$.

In [15] there is no explicit definition of valuation for the V-semantics, but it seems it is intended that a valuation should be defined on the full (infinite) set of variables.[3] On the contrary, Hodges works with finite valuations, as the V-semantics defined in [16]. He achieves this 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 $1 \leqslant i \leqslant n$ 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 $G(\mathcal{M}, \varphi, v_{(a_1,\ldots,a_n)})$ for $(a_1,\ldots,a_n) \in T$. Using trumps, Hodges [9] was able to prove that $\models^+$ and $\models^-$ can be defined in a composable way.

Theorem 1 (Hodges, 1997). *Let $\psi(\bar{x})$ be a regular formula of signature σ (with $\bar{x} = x_1,\ldots,x_n$), let $\mathcal{M}$ be a σ-structure and $T \subseteq |\mathcal{M}|^n$. Then,*

---

[3] In Section 2.1 of [15] valuations should be defined on some $X \subseteq Var$. So it is not explicitly required to be finite.

1. if $\varphi(\bar{x})$ is atomic or negated atomic,
   - $\mathcal{M} \models^+ \varphi(\bar{x})[T]$ iff $\mathcal{M} \models \varphi(\bar{x})[v_t]$ for all $t \in T$;
   - $\mathcal{M} \models^- \varphi(\bar{x})[T]$ iff $\mathcal{M} \not\models \varphi(\bar{x})[v_t]$ for all $t \in T$;
2. if $\varphi(\bar{x}) = {\sim}\psi(\bar{x})$,
   - $\mathcal{M} \models^+ \varphi(\bar{x})[T]$ iff $\mathcal{M} \models^- \psi(\bar{x})[T]$;
   - $\mathcal{M} \models^- \varphi(\bar{x})[T]$ iff $\mathcal{M} \models^+ \psi(\bar{x})[T]$;
3. if $\varphi(\bar{x}) = \psi_1(\bar{x}) \vee /_{x_{n_1}, \ldots, x_{n_k}} \psi_2(\bar{x})$ for some $\{x_{n_1}, \ldots, x_{n_k}\} \subseteq \{\bar{x}\}$,
   - $\mathcal{M} \models^+ \varphi(\bar{x})[T]$ iff there is a function $g : T \to \{L, R\}$ such that
     - g is $\{n_1, \ldots, n_k\}$-independent;
     - $T_L \neq \emptyset$ implies $\mathcal{M} \models^+ \psi_1(\bar{x})[T_L]$, with $T_L = \{t \mid t \in T, g(t) = L\}$;
     - $T_R \neq \emptyset$ implies $\mathcal{M} \models^+ \psi_2(\bar{x})[T_R]$, with $T_R = \{t \mid t \in T, g(t) = R\}$;
   - $\mathcal{M} \models^- \varphi(\bar{x})[T]$ iff $\mathcal{M} \models^- \psi_1(\bar{x})[T]$ and $\mathcal{M} \models^- \psi_2(\bar{x})[T]$;
4. if $\varphi(\bar{x}) = \exists y /_{x_{n_1}, \ldots, x_{n_k}} \psi(\bar{x}, y)$, where $y \notin \{\bar{x}\}$ and $\{x_{n_1}, \ldots, x_{n_k}\} \subseteq \{\bar{x}\}$,
   - $\mathcal{M} \models^+ \varphi(\bar{x})[T]$ iff there is a function $g : T \to |\mathcal{M}|$ such that
     - g is $\{n_1, \ldots, n_k\}$-independent; and
     - $\mathcal{M} \models^+ \psi(\bar{x}, y)[T']$, where $T' = \{(\bar{t}, g(\bar{t})) \mid (\bar{t}) \in T\}$;
   - $\mathcal{M} \models^- \varphi(\bar{x})[T]$ iff $\mathcal{M} \models^- \psi(\bar{x}, y)[T']$ for $T' = T \times |\mathcal{M}|$.[4]

Notation 1. Throughout 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 $\mathcal{M} \models^- Y$".

In first-order logic, a formula φ of some signature σ is *true* in a σ-structure $\mathcal{M}$ with respect to a valuation v if $\mathcal{M} \models \varphi[v]$. Equivalently φ is *false* in $\mathcal{M}$ with respect to v whenever $\mathcal{M} \not\models \varphi[v]$. The *meaning* of a formula φ of some signature σ, in a σ-structure $\mathcal{M}$, is given by the set of all valuations for which φ is *true* and the set of all valuations for which φ is *false*. Of course, in first-order logic, being *true* is equivalent to being not *false*. Therefore, the meaning of φ in a structure $\mathcal{M}$ is simply given by the set of all v such that $\mathcal{M} \models \varphi[v]$.

Similarly, in our setting of IF-like logics, a formula is said to be *true* when Eloïse has a winning strategy, and it is said to be *false* when Abélard has one. Under imperfect information there are formulas for which none of the players has a winning strategy, therefore there are formulas which are neither *true* nor *false*. The *meaning* of a formula, in a structure $\mathcal{M}$, under V- and T-semantics is formally defined as

$$\text{Meaning}^V_{\mathcal{M}}(\varphi) = \langle \{V \mid \mathcal{M} \models^+ \varphi[V]\}, \{V \mid \mathcal{M} \models^- \varphi[V]\} \rangle;$$

$$\text{Meaning}^T_{\mathcal{M}}(\varphi(\bar{x})) = \langle \{T \mid \mathcal{M} \models^+ \varphi(\bar{x})[T]\}, \{T \mid \mathcal{M} \models^- \varphi(\bar{x})[T]\} \rangle$$

where $\bar{x} = x_1, \ldots, x_n$. Intuitively two formulas are *equivalent* when they have the same meaning over any structure. So for first-order logic φ and ψ are equivalent iff for every σ-structure $\mathcal{M}$ and every valuation v, $\mathcal{M} \models \varphi[v]$ iff $\mathcal{M} \models \psi[v]$. From the above notion of meaning for IF-like logics, we derive the following equivalences:

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

T-equivalence (for a fixed tuple $\bar{x} = x_1, \ldots, x_n$): $\varphi_1(\bar{x}) \equiv_T \varphi_2(\bar{x})$ iff for all $\mathcal{M}$ and every $\emptyset \subset T \subseteq |\mathcal{M}|^n$, $\mathcal{M} \models^\pm \varphi_1(\bar{x})[T]$ iff $\mathcal{M} \models^\pm \varphi_2(\bar{x})[T]$.

T-equivalence (over all tuples): $\varphi_1 \equiv_T \varphi_2$ iff for all

$$\{x_1, \ldots, x_n\} \supseteq \text{Fv}(\varphi_1) \cup \text{Fv}(\varphi_2),$$

we have $\varphi_1(x_1, \ldots, x_n) \equiv_T \varphi_2(x_1, \ldots, x_n)$.

Notice that in the last two cases, φ_1 and φ_2 must be regular formulas (see footnote 1).

2.3. Signaling kicks in

It was first observed by Janssen [10] that V-semantics and signaling don't interact well. Consider, for instance, the following example (from [10], Section 7, formulas (32) and (33)):

$$\theta_1 := \forall x \forall y \forall z [x = y \vee \exists u \exists w /_x [w \neq x \wedge u = z]],$$

$$\theta_2 := \forall x \forall y \forall z [x = y \vee \exists y \exists w /_x [w \neq x \wedge y = z]].$$

---

[4] Observe that we take $T \times |\mathcal{M}|$ to be $\{(\bar{t}, a) \mid (\bar{t}) \in T, a \in |\mathcal{M}|\}$.

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:

$$f_\vee(v) = \begin{cases} L & \text{if } v(x) = v(y), \\ R & \text{otherwise,} \end{cases} \qquad f_{\exists u}(v) = v(z), \qquad f_{\exists w/_x}(v) = v(y).$$

Let us check that this is indeed a valid and winning strategy for Eloïse in any suitable model. Whatever values Abélard chooses for x, y and z, when Eloïse gets to pick L or R, either $v(x) = v(y)$ or $v(x) \neq v(y)$. In the first case, choosing L is the only sensible thing to do. But this means that if Eloïse ever gets to use $f_{\exists u}$ and $f_{\exists w/_x}$, then $v(x) \neq v(y)$ must hold; this is a global invariant of her strategy. So, let us assume that indeed, $v(x) \neq v(y)$; now Eloïse has to pick a value for u and since Abélard can challenge it later, her only choice is to pick $v(z)$. Finally, she has to provide a value for w other than $v(x)$. However, since the quantifier has a restriction on x, not every strategy function is valid. As one can see, $f_{\exists w/_x}$ is a valid strategy function and, because of the global invariant, returns a proper result. 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 Eloïse has no analogous winning strategy 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 the strategy (i.e., blocking the signal). In fact, it is not hard to show that for arbitrary structures, Eloïse has no winning strategy at all 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. In our opinion, 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.[5] In that sense u, v or w should be as good as y. The idea of variables as a mere syntactic device is reinforced with the fact that, at least in classical first-order logic, one could drop variables altogether and replace them with de Bruijn indices [17]. Of course, using de Bruijn indices is to somehow 'regularizing the logic'.

This crucial notion has a name: α-*equivalence*. Two formulas φ_1 and φ_2 are said to be α-*equivalent* ($\varphi_1 \equiv_\alpha \varphi_2$) if their syntactic trees are isomorphic, they have the same occurrences of free variables and for each occurrence of a bound variable the quantifier that binds it matches. For example, if $\theta := \exists x \exists y[x = z]$, then $\theta \equiv_\alpha \exists y \exists w[y = z]$ but $\theta \not\equiv_\alpha \exists z \exists y[z = z]$ and $\theta \not\equiv_\alpha \exists x \exists x[x = z]$ (for formal definitions see any textbook on λ-calculus, e.g. [18]).

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:

$$\theta_3 := \exists y \exists z/_{x,y}[z = x], \qquad \theta_4 := \exists u \exists z/_{x,u}[z = x].$$

For $V = \{v \mid v(x) = v(u)\}$ and $\|\mathcal{M}\| \geqslant 2$ it is easy to see that $\mathcal{M} \models^+ \theta_3[V]$ but $\mathcal{M} \not\models^+ \theta_4[V]$. The above example shows the kind of problems that we are going to solve in this paper. Observe that in θ_1 and θ_2 of page 336 one can reconstruct the context of the subformulas from the sentence in which they are introduced, that is *syntactically*. In θ_3 and θ_4, the context is in the structure of V, but there is no traceable syntactic counterpart to that. In [16] this problem is solved by relativizing the equivalence by restricting to valuations that avoid the bound variables occurring in the formulas.

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 claimed in [15, Lemma 3.1(a)] that it holds and this led to some flawed results (see [14]). In the face of this, it is worth verifying that, fortunately, α-equivalence does hold under T-semantics.

Proposition 1. *$\varphi_1(\bar{x}) \equiv_\alpha \varphi_2(\bar{x})$ implies $\varphi_1(\bar{x}) \equiv_T \varphi_2(\bar{x})$, for $\bar{x} = x_1, \ldots, x_n$.*

Proof. Because $\equiv_T$ is a transitive relation, one only needs to prove the implication for the case where $\varphi_1(\bar{x})$ and $\varphi_2(\bar{x})$ differ in at most one bound variable. This is carried out by induction on $\varphi_1(\bar{x})$. The case that matters is when $\varphi_1(\bar{x})$ is of the form $\exists u/_{z_1,\ldots,z_k} \psi_1(\bar{x}, u)$ and $\varphi_2(\bar{x})$ is $\exists v/_{z_1,\ldots,z_k} \psi_2(\bar{x}, v)$, where $\psi_2(\bar{x}, v)$ is obtained from $\psi_1(\bar{x}, u)$ by replacing the free variable u by v. These two subformulas match on the bound variables but differ in one free variable (u vs. v). However, now we can prove by a straightforward induction that $\mathcal{M} \models^\pm \psi_1(\bar{x}, u)[T]$ iff $\mathcal{M} \models^\pm \varphi_2(\bar{x}, v)[T]$ for all suitable $\mathcal{M}$ and T. The key point here is that at item 1 of Theorem 1 the name u or v is irrelevant, as long as they come in the same order in the lists $(\bar{x}, u)$ and $(\bar{x}, v)$. □

Corollary 2. *Let φ_1 and φ_2 be regular. If $\varphi_1 \equiv_\alpha \varphi_2$ then $\varphi_1 \equiv_T \varphi_2$.*

Proof. Immediate from the definition of T-equivalence over all tuples. □

---

[5] That is, in the formula $\forall x \exists y[x \neq y]$, x and y are *distinct* and that is why we expect it to be equivalent to $\forall y \exists z[y \neq z]$ but not to $\forall y \exists y[y = y]$.

Now, observe that formulas θ_3 and θ_4 above are regular and α-equivalent. Hence, by Corollary 2, $\theta_3 \equiv_T \theta_4$. But we have already seen that $\theta_3 \not\equiv_V \theta_4$. That is, θ_3 and θ_4 have the same meaning under T-semantics but different meaning under V-semantics. This clearly shows that V-semantics simply fails to be a *generalization* of T-semantics.

Proposition 2. *There are regular formulas which are T-equivalent but not V-equivalent. Therefore, V-semantics is not an extension of T-semantics.*

3. Uniform semantics for regular and irregular formulas

Failure of α-equivalence has been regarded in the literature as an unexpected, unusual yet, characteristic feature of IF-like logics. This begs a simple question: what is the precise relation between independence restrictions and α-equivalence? Or, in other words, in what sense is this an unavoidable feature of these logics?

We believe that in order to answer these questions, one first needs to untangle the exact notion of *independence* being used. That is, we must realize what is the intended meaning[6] of the slashed quantifier in the formula $\exists x /_y \varphi$. At first sight, it may seem like the answer is clear and simple enough: the independence restriction in this case is meant to be informally read "there exists a value for x, *independent of the value of y*, such that it makes φ true".

Arguably, it is this notion of "independence with respect to values" that lies behind the formal semantics in the tradition of [15,16], which we have called V-semantics. Since independence restrictions range over the values of variables, in these game-theoretical formalizations one only needs to hold the *current* value of each variable. But this means that each time a variable is rebound there is a loss of information, typically observed as a signaling that gets blocked. This prevents α-equivalence from holding in the irregular case.

This informal argument is not enough to conclude that the failure of α-equivalence is an intrinsic characteristic of IF-like logics, though. In particular, Proposition 2 already shows that there are at least two distinct semantics for *regular* IF-like formulas, one where α-equivalence holds and one where it does not.

In a way, this paper is built on the observation that $\exists x /_y \varphi$ can alternatively be read as "there exists a value for x, *independent of the last value chosen for y*, such that it makes φ true". The difference is subtle: by talking about "the last value for y" it is implied that there might have been previous values and that we don't want to forbid the use of these values in picking a new value for x. Of course, a precise definition of this notion is required (in particular, what it means for variables that are free), and this is what will be done in this section. It will be also shown that the obtained logic coincides with T-semantics on regular formulas and that it has α-equivalence for the whole fragment. This clearly shows that failure of α-equivalence is not inherent to IF-like logics.

There are many equivalent ways in which this alternative notion of independence can be formalized. We have opted to simply decouple chosen values from valuations; others may prefer, for example, to map each variable to the list of all its previous values. As long as the finite *history* of moves is part of the game board, equivalent game-theoretical semantics can be given.

It is worth making two closing observations. First, game-theoretical semantics for classical first-order logic is based on a game of perfect information; in this context the history of moves is irrelevant. Second, in T-semantics, a tuple $(a_1, \ldots, a_n)$ represents a precise finite *history* of moves for $\varphi(x_1, \ldots, x_n)$. In this setting, the *history* grows up one element at a time, as a consequence of each player's move for a $\exists$ or $\forall$. That is, when some player picks a value a_{n+1} for x_{n+1} in $\varphi(x_1, \ldots, x_n) = \exists x_{n+1} / y_1, \ldots, y_k \psi$, the new history becomes $(a_1, \ldots, a_n, a_{n+1})$ and the new formula becomes $\psi(x_1, \ldots, x_n, x_{n+1})$. But since T-semantics is restricted to the regular fragment, the value chosen for x_{n+1} will be a_{n+1} *all along the play*, and therefore both notions of independence coincide after all.

The fact that both notions of independence coincide in game-theoretical semantics for classical first-order and in T-semantics, might explain why the necessity of "history preserving" valuations for correctly dealing with independence restrictions has been so far overlooked.

Let's now move to a more technical field and define in a precise way what will be our *history*, how this history will interact with rebinding and finally what will be our semantics of unrestricted IF-like formulas. Recall that we will base our semantics upon the observation that $\exists x /_y \varphi$ should be read as "there exists a value for x, *independent of the last value chosen for y*, such that it makes φ true". It is clear that in this setting, classical valuations are an inadequate device to formalize this semantics: under rebinding of variables, they simply fail to keep track of all the previous choices. 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 that one can think of the composition $s \circ p$ as denoting a classical valuation. Almost all of our presentation can be done using sequences of finite length (see item 2 of Theorem 4 below). In this case, $s \circ p$ would represent a valuation with finite domain, as the presentation of [16]. We opted for using infinite sequences for the sake of clarity.

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 of $|\mathcal{M}|^\omega$, and $h < \omega$ can be regarded as indicating how many "previous choices" are in

---

[6] Here we are using the word *meaning* in its informal, non-technical sense.

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 $G(\mathcal{M}, \varphi, S, p, h)$ 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. Although in this formulation the game-theoretic nature of the semantics becomes arguably less apparent, we prefer it since the higher-order nature of the logic becomes more evident.[7] In Section 4, this game will be generalized to an n-turn game 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 $f : A \to B$ be any function; we denote with $f[a \mapsto b]$ the function

$$f[a \mapsto b](x) = \begin{cases} b & \text{if } x = a, \\ f(x) & \text{otherwise.} \end{cases}$$

As usual, if $A' \subseteq A$ then $f \upharpoonright A' : A' \to B$ will be the restriction of f to A'.

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:

1. $\exists$-nodes can be decorated with any function $f : |\mathcal{M}|^\omega \to |\mathcal{M}|$.
2. $\vee$-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 generating 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 $i \in \{0, \ldots, h-1\} \setminus \{p(y_1), \ldots, p(y_n)\}$, 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 $i \in \{0, \ldots, h-1\} \setminus \{p(y_1), \ldots, p(y_n)\}$, 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 $R(t_1, \ldots, t_k)$, the 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.

One could argue that the in the play described above there is no such thing as *turns*, which usually implies alternation of some kind. The reason why we call it *the turn* will become more clear in Section 4.

Winning strategies. A *strategy* for a player of the game $G(\mathcal{M}, \varphi, S, p, h)$ 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 that since Abélard and Eloïse must play *simultaneously* (i.e. ignoring the opponent move) this is a game of *imperfect information*. Furthermore, the initial valuation is "randomly" picked by Nature.[8] Therefore, some games are probably undetermined, that is, none of the players have a winning strategy (Example 3 in Section 4 exhibits one such game).

---

[7] In this sense, our formulation is very close to the one due to Väänänen [19] in which Eloïse has to pick functions. In our case, both players pick functions for there are also restrictions on Abélard's moves.

[8] The elements picked by Nature during the evaluation phase are irrelevant in this respect: the game can be formulated without this although in a less convenient way: during evaluation, the player who did not play the function being evaluated picks the elements to obfuscate the current valuation. Caicedo

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 6).

Definition 2 (*Proper context*). We say that $p : \text{Vars} \to \omega$ and $h < \omega$ are a *proper context* for a formula φ if $p \upharpoonright \text{Fv}(\varphi)$ is injective and $\{p(x) \mid x \in \text{Fv}(\varphi)\} \subseteq \{0, \dots, h-1\}$.

Definition 3 ($\models^+$ *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, $p : \text{Vars} \to \omega$ 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, $\mathcal{M} \models^+ \theta[S, p, h]$ should hold for any $S \subseteq |\mathcal{M}|^\omega$, and for *any* (not necessarily proper contexts for θ) $p : \text{Vars} \to \omega$ and $h < \omega$. 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 $p(y) \geqslant h$.

The restriction to *proper contexts* is essential, and can be seen as the counterpart to the restrictions imposed in [16, Definitions 6.1 and 6.10] going from G-equivalence to the Z-equivalence, in order to avoid clashes. The role of h is to keep the domain of the 'valuations' essentially finite, which they are as well in [16]. What is different in this formalization is that the S carries potentially more information than values assigned to values, as becomes clear in the next example.

Example 2. Let us revisit the irregular formula

$$\theta_2 := \forall x \forall y \forall z \big[x = y \lor \exists y \exists w /_x [w \neq x \land y = z] \big]$$

from Section 2.3 (page 336). We shall verify that for any model $\mathcal{M}$, $\mathcal{M} \models^+ \theta_2$. For this, consider the following strategy for Eloïse for the game $\mathsf{G}(\mathcal{M}, \theta_2, S, p, h)$:

$$f_\lor(s) = \begin{cases} L & \text{if } s(h) = s(h+1), \\ R & \text{otherwise}, \end{cases} \qquad f_{\exists y}(s) = s(h+2), \qquad f_{\exists w/x}(s) = s(h+1).$$

The reader should verify that this is essentially the same strategy used for θ_1 in Section 2.3. 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. The only information Eloïse needs to know is that the value picked by the first three universal quantifiers will be stored in $s(h)$, $s(h+1)$ and $s(h+2)$, respectively.

Even though the inner variable y (i.e. the one quantified existentially) now refers to the value stored at location $h+3$, it is within the strategy still possible to refer to the value stored at location $h+1$, which is the value for the outer y (i.e. the one quantified universally). What is different here from [16] is that the value is preserved even though no free variable within the formula refers to this position anymore, assuming p is a proper context. (It might actually be the case that some other, non-occurring variable still refers to location $h+1$, but this is irrelevant.) While the main problems addressed in [16] involve situations where the choice of a value for a bound variable in the formula overwrites a crucial value it previously had *even* if it did not occur free in the formula, this overwriting no longer happens in this formalism. Therefore, it is possible to define all notions on the basis of the (visible) free variables of the formula, rather than also having to take domains of valuations into account (which is what [16] does).

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 [9] and [15].

Theorem 3 (*S-semantics compositionality*). *Let $\mathcal{M}$ be a suitable model, let $S \subseteq |\mathcal{M}|^\omega$ and let $p : \text{Vars} \to \omega$ and $h < \omega$ be a proper context. Then*

---

and Krynicki suggest that Nature can be completely removed, by making the opponent of the player who owns the node closest to the root be the one who picks the initial valuation [15, p. 20, Remark ii.]. However, this selection cannot be part of the player's strategy, as this example shows: let v and v' be a constant and an injective valuation respectively; now, if Abélard is allowed to use the strategy "pick v' as initial valuation", then $\mathbb{N} \models^- x = y[\{v, v'\}]$ which violates the standard compositional definition of $\models^-$.

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 $\{p(y_1), \ldots, p(y_k)\} \cup \{h, h+1, \ldots\}$-independent;
   - $S_L \neq \emptyset$ implies $\mathcal{M} \models^+ \psi_1[S_L, p, h]$, where $S_L = \{s \mid s \in S, f(s) = L\}$;
   - $S_R \neq \emptyset$ implies $\mathcal{M} \models^+ \psi_2[S_R, p, h]$, 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 $\mathcal{M} \models^- \psi_2[S, p, h]$;

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 $\{p(y_1), \ldots, p(y_k)\} \cup \{h, h+1, \ldots\}$-independent; and
   - $\mathcal{M} \models^+ \psi[S', p[x \mapsto h], h+1]$, where $S' = \{s[h \mapsto f(s)] \mid s \in S\}$;

8. $\mathcal{M} \models^- \exists x /_{y_1, \ldots, y_k} \psi[S, p, h]$ iff $\mathcal{M} \models^- \psi[S', p[x \mapsto h], h+1]$ for

$$S' = \bigcup_{a \in |\mathcal{M}|} \{s[h \mapsto a] \mid s \in S\}.$$

Proof. 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 $G(\mathcal{M}, \varphi, S, p, h)$, then he also has a winning strategy in which all the functions that constitute it satisfy the independence restriction, and this is relatively straightforward (the full details can be seen, e.g. in [16, Theorems 4.7 and 4.8]). In every case, one also has to check that contexts are proper, but this is trivial. $\quad\square$

In first-order logic, satisfaction is invariant under modification of the values assigned to variables that are not free in the formula (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 $S \upharpoonright n = \{(s(0), \ldots, s(n-1)) \mid s \in S\}$; we use the term h-permutation for any bijective function $\pi : \omega \to \omega$ such that $\pi(i) = i$ for all $i \geqslant h$; and $S \circ \pi = \{s \circ \pi \mid s \in S\}$.

Theorem 4. *For all suitable $\mathcal{M}$, nonempty $S \subseteq |\mathcal{M}|^\omega$ and proper contexts for φ, p and h:*

1. $\mathcal{M} \models^\pm \varphi[S, p, h]$ iff $\mathcal{M} \models^\pm \varphi[S, \tilde{p}, h]$ *whenever* $\tilde{p} \upharpoonright \mathrm{Fv}(\varphi) = p \upharpoonright \mathrm{Fv}(\varphi)$.

2. $\mathcal{M} \models^\pm \varphi[S, p, h]$ iff $\mathcal{M} \models^\pm \varphi[\tilde{S}, p, h]$ *whenever* $\tilde{S} \upharpoonright h = S \upharpoonright h$.

3. $\mathcal{M} \models^\pm \varphi[S, p, h]$ iff $\mathcal{M} \models^\pm \varphi[S \circ \pi, \pi \circ p, h]$ *whenever* π *is an h-permutation.*

Proof. We only prove the last equivalence; the other are similar. We proceed by induction on φ and rely on Theorem 3. For the base case, it is enough to observe that since π is a permutation, then $(s \circ \pi) \circ (\pi \circ p) = s \circ p$ for every $s \in |\mathcal{M}|^\omega$. Negation follows trivially by inductive hypothesis. Suppose now that φ is of the form $\exists x /_{y_1, \ldots, y_k} \psi$. We know that $\mathcal{M} \models^+ \exists x /_{y_1, \ldots, y_k} \psi[S, p, h]$ iff $\mathcal{M} \models^+ \psi[S', p[x \mapsto h], h+1]$, where $S' = \{s[h \mapsto f(s)] \mid s \in S\}$ and f is $\{p(y_1), \ldots, p(y_k)\} \cup \{h, h+1, \ldots\}$-independent. Let f' be such that $f'(s) = f(s \circ \pi)$, which implies $f'(s \circ \pi) = f(s)$. Furthermore, since, by hypothesis, $\pi(i) = i$ for all $i \geqslant h$ we have that f' is $\{\pi \circ p(y_1), \ldots, \pi \circ p(y_k)\} \cup \{h, h+1, \ldots\}$-independent. We only need to verify that $\mathcal{M} \models^+ \psi[S'', (\pi \circ p)[x \mapsto h], h]$ where $S'' = \{s[h \mapsto f'(s)] \mid s \in S \circ \pi\}$. But because $\pi(h) = h$, we have:

$$
\begin{aligned}
S'' = {} & \{s[h \mapsto f'(s)] \mid s \in S \circ \pi\} = \\
& \{(s \circ \pi)[h \mapsto f'(s \circ \pi)] \mid s \in S\} = \\
& \{(s \circ \pi)[h \mapsto f(s)] \mid s \in S\} = \\
& \{s[h \mapsto f(s)] \circ \pi \mid s \in S\} = S' \circ \pi
\end{aligned}
$$

and the rest follows by inductive hypothesis using the fact that $(\pi \circ p)[x \mapsto h] = \pi \circ (p[x \mapsto h])$. The remaining cases are analogous. $\quad\square$

Observe that item 2 of Theorem 4 illustrates that h makes our given alternative to valuations essentially finite.

In [9, Theorem 7.6], Hodges shows that the regular fragment of slash logic has the full abstraction property. This basically says that two formulas with the same free variables have the same semantic value (i.e. the same *meaning*) if and only if

replacing one of them by the other in a sentence never changes the truth value of the sentence. We agree with Hodges that "any sensible compositional semantics must be fully abstract" [9, p. 541]. Before verifying that S-semantics preserves full abstraction, we need to give a precise definition of the *meaning* of a formula in a structure.

Definition 4 *(Meaning for S-semantics).* The *meaning* of a formula φ of signature σ over a σ-structure $\mathcal{M}$ is defined as $\text{Meaning}_{\mathcal{M}}^{\text{S}}(\varphi) = \langle a, b \rangle$ where

$$a = \left\{ (S, p, h) \mid p, h \text{ is a proper context for } \varphi \text{ and } \mathcal{M} \models^+ \varphi[S, p, h] \right\};$$
$$b = \left\{ (S, p, h) \mid p, h \text{ is a proper context for } \varphi \text{ and } \mathcal{M} \models^- \varphi[S, p, h] \right\}.$$

As before, we say that two formulas are equivalent whenever their meanings coincide over any appropriate structure.

Definition 5 *($\equiv_h$ and $\equiv$).* We write $\varphi_1 \equiv_h \varphi_2$, if for every suitable $\mathcal{M}$, S and p such that p, h is a proper context for both φ_1 and φ_2, $\mathcal{M} \models^{\pm} \varphi_1[S, p, h]$ iff $\mathcal{M} \models^{\pm} \varphi_2[S, p, h]$. We use $\varphi_1 \equiv \varphi_2$ (*S-equivalence*) whenever $\varphi_1 \equiv_h \varphi_2$ for every $h \geqslant \|\text{Fv}(\varphi_1) \cup \text{Fv}(\varphi_2)\|$.

Theorem 5 *(Full abstraction). Let φ_1 and φ_2 be two formulas of signature σ such that $\text{Fv}(\varphi_1) = \text{Fv}(\varphi_2)$. The following are equivalent*:

1. *For any σ-structure $\mathcal{M}$, $\text{Meaning}_{\mathcal{M}}^{\text{S}}(\varphi_1) = \text{Meaning}_{\mathcal{M}}^{\text{S}}(\varphi_2)$.*
2. *For any sentence ψ and for any appropriate structure $\mathcal{M}$, if ψ' comes from ψ by replacing an occurrence of φ_1 as a constituent of ψ by an occurrence of φ_2, $\mathcal{M} \models^{\pm} \psi$ iff $\mathcal{M} \models^{\pm} \psi'$.*

Proof. The implication $(1) \Rightarrow (2)$ is straightforward from Theorem 3. For $(2) \Rightarrow (1)$, suppose that $\text{Meaning}_{\mathcal{M}}^{\text{S}}(\varphi_1) \neq \text{Meaning}_{\mathcal{M}}^{\text{S}}(\varphi_2)$ for some σ-structure $\mathcal{M}$, where $\text{Fv}(\varphi_1) = \{x_1, \ldots, x_n\}$. Without loss of generality, assume there is a nonempty set $S \subseteq |\mathcal{M}|^{\omega}$, a natural number $k \geqslant n$ and a function $p : \text{Vars} \to \omega$ such that

$$\mathcal{M} \models^+ \varphi_1[S, p, k] \quad \text{and} \quad \mathcal{M} \not\models^+ \varphi_2[S, p, k]. \tag{2}$$

(The case $\mathcal{M} \models^- \varphi_1[S, p, k]$ and $\mathcal{M} \not\models^- \varphi_2[S, p, k]$ is similar.) By item 2 of Theorem 4 above, one may assume that if $s \in S$ and $s' \upharpoonright k = s \upharpoonright k$ then $s' \in S$. Furthermore, by item 3 one may also assume that $p(x_i) = i - 1$ for $i \in \{1, \ldots, k\}$. Observe that $x_{n+1}, \ldots, x_k$ are not free variables of φ_1.

Let us extend σ with a new k-relation symbol R and let $\mathcal{M}'$ be the model extending $\mathcal{M}$ with R interpreted as $S \upharpoonright k$. By the properties of S described above, $s \in S$ iff $s \upharpoonright k \in R^{\mathcal{M}'}$. For $i = 1, 2$ let

$$\chi_i := \forall x_1 \cdots \forall x_k \big(\sim R(x_1, \ldots, x_k) \vee \varphi_i \big).$$

It is clear that χ_1 and χ_2 are both sentences and that χ_2 comes from χ_1 by replacing the occurrence of φ_1 by φ_2. To conclude the proof, we now show that

$$\mathcal{M}' \models^+ \chi_1 \big[|\mathcal{M}|^{\omega}, p, 0 \big] \quad \text{and} \quad \mathcal{M}' \not\models^+ \chi_2 \big[|\mathcal{M}|^{\omega}, p, 0 \big].$$

The winning strategy for Eloïse on the game $\text{G}(\mathcal{M}', \chi_1, |\mathcal{M}|^{\omega}, p, 0)$ consists in playing, for the outermost $\vee$ of χ_1, an $f_{\vee}$ such that $f_{\vee}(s) = L$ iff $s \notin S$, in combination with the strategy she already has – according to (2) – for $\text{G}(\mathcal{M}, \varphi_1, S, p, k)$. Indeed, after evaluating the first k functions played by Abélard, Nature ends up with some sequence $s \in |\mathcal{M}|^{\omega}$ in the placeholder, $h = k$ and the pointer function p. Then Nature evaluates $f_{\vee}(s)$; in case $s \notin S$ then Nature picks $\sim R(x_1, \ldots, x_k)$ and evaluates it in $s \circ p$. Since $s \upharpoonright k \notin R^{\mathcal{M}'}$, then $\mathcal{M}' \not\models R(x_1, \ldots, x_k)[s \circ p]$ and therefore Eloïse wins. Alternatively, if $s \in S$ then Nature continues with φ_1 and Eloïse wins since her strategy is winning for $\text{G}(\mathcal{M}, \varphi_1, S, p, k)$.

On the other hand, Eloïse does not have a winning strategy for the game $\text{G}(\mathcal{M}', \chi_2, |\mathcal{M}|^{\omega}, p, 0)$, as having such strategy would imply having one for $\text{G}(\mathcal{M}, \varphi_2, S, p, k)$ – just remove the function picked for the outermost $\vee$-node – and this would contradict (2). $\square$

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

Proposition 3. *If $\varphi_1 \equiv_{\alpha} \varphi_2$ then $\varphi_1 \equiv \varphi_2$.*

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 3, S-semantics is a proper generalization of T-semantics. To do this, we will use the following technical lemma.

Lemma 6. *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 $0 \leqslant i < h$. Then $\mathcal{M} \models^{\pm} \varphi[S, p, h]$ iff $\mathcal{M} \models^{\pm} \varphi(x_0 \ldots x_{h-1})[S \upharpoonright h]$.*

Proof. Let $\bar{x} = x_0, \ldots, x_{h-1}$. Suppose the list of occurrences of bound variables appearing in each branch of the syntactic tree of $\varphi(\bar{x})$ (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. Both the base case and outermost negation are straightforward. Let us analyze the case $\varphi = \exists x_h /_{x_{n_1}, \ldots, x_{n_k}} \psi(\bar{x}, x_h)$, for some $\{n_1, \ldots, n_k\} \subseteq \{0, \ldots, h-1\}$.

For the left-to-right implication, suppose $\mathcal{M} \models^+ \varphi(\bar{x})[S, p, h]$. Just like we did in the proof of Theorem 5, we may assume, without loss of generality, that if $s \in S$ and $s' \upharpoonright h - 1 = s \upharpoonright h - 1$, then $s' \in S$. By Theorem 3 (item 7), $\mathcal{M} \models^+ \psi(\bar{x}, x_h)[S', p[x_h \mapsto h], h+1]$, where $S' = \{s[h \mapsto f(s)] \mid s \in S\}$, for some $f : S \to |\mathcal{M}|$ such that f is $\{p(x_{n_1}), \ldots, p(x_{n_k})\} \cup \{h, h+1, \ldots\}$-independent. Since $p = p[x_h \mapsto h]$, by inductive hypothesis we get $\mathcal{M} \models^+ \psi(\bar{x}, x_h)[S' \upharpoonright h + 1]$. Fix $a \in |\mathcal{M}|$ and define $g : S \upharpoonright h \to |\mathcal{M}|$ as $g(s_0, \ldots, s_{h-1}) = f(s_0 \ldots s_{h-1} aaa \ldots)$ for every $(s_0, \ldots, s_{h-1}) \in S \upharpoonright h$. Since $s_0 \ldots s_{h-1} aaa \ldots \in S$ (because of the above assumption), g is clearly well defined; moreover, because f is $\{p(x_{n_1}), \ldots, p(x_{n_k})\} \cup \{h, h+1, \ldots\}$-independent, g is $\{n_1, \ldots, n_k\}$-independent. But this means that $S' \upharpoonright h + 1 = \{(s_0, \ldots, s_{h-1}, g(s_0, \ldots, s_{h-1})) \mid (s_0, \ldots, s_{h-1}) \in S \upharpoonright h\}$ and, therefore, by Theorem 1 (item 4), $\mathcal{M} \models^+ \varphi(\bar{x})[S \upharpoonright h]$.

For the other direction, suppose $\mathcal{M} \models^+ \varphi(\bar{x})[S \upharpoonright h]$. By Theorem 1 (item 4), there exists some function $g : S \upharpoonright h \to |\mathcal{M}|$ that is $\{n_1, \ldots, n_k\}$-independent and such that $\mathcal{M} \models^+ \psi(\bar{x}, x_h)[T']$, with

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

Observe that $T' = S' \upharpoonright h + 1$, 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 conclude that $\mathcal{M} \models^+ \psi(\bar{x}, x_h)[S', p[x_h \mapsto h], h+1]$.

Define $f : S \to |\mathcal{M}|$ as $f(s) = g(s(0), \ldots, s(h-1))$ for $s \in S$. By definition, f does not depend on positions $\{h, h+1, \ldots\}$, and since g is $\{n_1, \ldots, n_k\}$-independent, then f is $\{p(n_1), \ldots, p(n_k)\} \cup \{h, h+1, \ldots\}$-independent. By Theorem 3 (item 7) we conclude $\mathcal{M} \models^+ \varphi(\bar{x})[S, p, h]$.

The case for $\models^-$ and $\varphi = \exists x_h /_{x_{n_1}, \ldots, x_{n_k}} \psi(\bar{x}, x_n)$ is straightforward. A similar argument can be used for the case $\varphi(\bar{x}) = \psi_1(\bar{x}) \vee /_{x_{n_1}, \ldots, x_{n_k}} \psi_2(\bar{x})$. $\quad\square$

Theorem 7. *If φ_1 and φ_2 are regular then the following are equivalent*:

1. $\varphi_1(x_0, \ldots, x_{h-1}) \equiv_T \varphi_2(x_0, \ldots, x_{h-1})$.
2. $\varphi_1 \equiv_h \varphi_2$.

Proof. We will only show it for $\models^+$, the argument for $\models^-$ is similar. In what follows $\bar{x}$ will stand for $x_0, \ldots, x_{h-1}$. To see that (1) implies (2), by the counterpositive, suppose that $\mathcal{M} \models^+ \varphi_1(\bar{x})[S, p, h]$ and $\mathcal{M} \not\models^+ \varphi_2(\bar{x})[S, p, h]$, for some suitable model $\mathcal{M}$ and some p such that p, h is a proper context for φ_1 and φ_2. One can build an h-permutation π such that $\pi(p(x_i)) = i$ for $0 \leqslant i < h$ and using Theorem 4 one gets $\mathcal{M} \models^+ \varphi_1(\bar{x})[S \circ \pi, \pi \circ p, h]$ but $\mathcal{M} \not\models^+ \varphi_2(\bar{x})[S \circ \pi, \pi \circ p, h]$. By Proposition 3, 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 6 we obtain $\mathcal{M} \models^+ \varphi_1'(\bar{x})[S \circ \pi \upharpoonright h]$ and $\mathcal{M} \not\models^+ \varphi_2'(\bar{x})[S \circ \pi \upharpoonright h]$, which implies $\varphi_1(\bar{x}) \not\equiv_T \varphi_2(\bar{x})$ using Proposition 1.

For the other direction, suppose $\varphi_1(\bar{x}) \not\equiv_T \varphi_2(\bar{x})$, i.e., $\mathcal{M} \models^+ \varphi_1(\bar{x})[T]$ and $\mathcal{M} \not\models^+ \varphi_2(\bar{x})[T]$, 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 $p(x_i) = i$. Since $\mathrm{Fv}(\varphi_1) \subseteq \{\bar{x}\}$ and $\mathrm{Fv}(\varphi_2) \subseteq \{\bar{x}\}$, it follows that p, h is a proper context for both φ_1 and φ_2. Again, using invariance under α-equivalence and Lemma 6 we conclude $\mathcal{M} \models^+ \varphi_1[S, p, n]$ and $\mathcal{M} \not\models^+ \varphi_2[S, p, n]$. $\quad\square$

Corollary 8. *If φ_1 and φ_2 are regular, then $\varphi_1 \equiv_T \varphi_2$ iff $\varphi_1 \equiv \varphi_2$.*

Proof. From left to right, suppose that for all $\{y_1, \ldots, y_k\} \supseteq \mathrm{Fv}(\varphi_1) \cup \mathrm{Fv}(\varphi_2)$, $\varphi(y_1, \ldots, y_k) \equiv_T \psi(y_1, \ldots, y_k)$. Given $h \geqslant \|\mathrm{Fv}(\varphi_1) \cup \mathrm{Fv}(\varphi_2)\|$, let

$$\{x_0, \ldots, x_{h-1}\} \supseteq \mathrm{Fv}(\varphi_1) \cup \mathrm{Fv}(\varphi_2).$$

By Theorem 7 it follows that $\varphi_1 \equiv_h \varphi_2$.

From right to left, suppose $\varphi_1 \equiv_h \varphi_2$ for every $h \geqslant \|\mathrm{Fv}(\varphi_1) \cup \mathrm{Fv}(\varphi_2)\|$. Given any $\{x_1, \ldots, x_n\} \supseteq \mathrm{Fv}(\varphi_1) \cup \mathrm{Fv}(\varphi_2)$ it is clear that $\varphi_1 \equiv_n \varphi_2$, since $n \geqslant \|\mathrm{Fv}(\varphi_1) \cup \mathrm{Fv}(\varphi_2)\|$. By Theorem 7 it follows that $\varphi_1(x_1, \ldots, x_n) \equiv_T \varphi_2(x_1, \ldots, x_n)$. $\quad\square$

From this corollary we may conclude that S-semantics is in fact an extension of T-semantics to arbitrary formulas (cf. Proposition 2).

4. Game theoretical semantics for IF with flattening

It is clear that 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; Example 3 below exhibits one such sentence. When working with imperfect information, being false is in general stronger than being not true.

Example 3. Let $\theta := \forall x \exists y /_x [y \neq x]$ and let $\mathcal{M}$ be an arbitrary model with at least two elements. For $h = 0$, Eloïse has certainly no winning strategy, but whatever Abélard chooses for x, it may happen that Eloïse is able to "accidentally guess" it. Thus $\mathcal{M} \not\models^+ \theta$ and $\mathcal{M} \not\models^- \theta$.

Hodges [9] considers the problem of adding classical negation ($\neg$) 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 with respect to classical negation. To this end, he introduces the *flattening operator* $\downarrow$, and stipulates $\neg\psi \equiv \sim \downarrow \psi$.

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 3 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 6 ($\models^+$ *and* $\models^-$ *for S-semantics with* $\downarrow$). We define $\models^+$ and $\models^-$ as the relation induced by clauses 1–8 of Theorem 3, 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 $s \in S$.

Example 4. Observe from Example 3 that $\mathcal{M} \not\models^+ \theta$ and $\mathcal{M} \not\models^+ \sim\theta$, since $\sim$ is game-negation. Let us see what happens if we use $\neg$ instead (recall $\neg\psi$ is defined as $\sim\downarrow \psi$). Using Theorem 3 (item 3) and Definition 6 (item 10), one gets $\mathcal{M} \models^+ \neg\theta$ iff $\mathcal{M} \models^- \downarrow \theta$ iff $\mathcal{M} \not\models^+ \theta[s, p, h]$ for every $s \in |\mathcal{M}|^\omega$ and proper context p, h. Since for no $s \in \mathcal{M}^\omega$ Eloïse has a winning strategy for $\mathsf{G}(\mathcal{M}, \theta, S, p, h)$, we conclude $\mathcal{M} \models^+ \neg\theta$.

Hodges could not provide a game-theoretical counterpart for the $\downarrow$ operator.[9] Next, we will fill this gap by defining the game $\mathsf{G}_\downarrow(\mathcal{M}, \varphi, S, p, h)$, that extends the rules of the game described in Section 3 to handle formulas containing arbitrary occurrences of $\downarrow$, and proving its equivalence.

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 leaves 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.

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**.

Observe that in $\mathsf{G}_\downarrow(\mathcal{M}, \downarrow \psi, S, p, h)$ what Abélard and Eloïse do in their first turn is completely irrelevant: they will get to redecorate ψ in the second turn. Furthermore, notice that if S is a singleton set, then the first turn can be ignored

---

[9] "In the presence of $\downarrow$, we can't define a game $G(\phi, A)$ for arbitrary A and ϕ." [9, p. 556].

altogether (Nature has to pick a valuation from a singleton set) and thus the game becomes of perfect information (cf. footnote 8 on page 339).

We won't go into a formal description of a winning strategy for $G_\downarrow(\mathcal{M}, \varphi, S, p, h)$. We simply take it to be some form of oracle that, when followed, guarantees that the game will end in a winning position, regardless what the opponent or Nature may play.

Theorem 9 *(Game semantics for $\downarrow$). 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 $h < \omega$, the following holds*:

- $\mathcal{M} \models^+ \varphi[S, p, h]$ *iff Eloïse has a winning strategy for* $G_\downarrow(\mathcal{M}, \varphi, S, p, h)$;
- $\mathcal{M} \models^- \varphi[S, p, h]$ *iff Abélard has a winning strategy for* $G_\downarrow(\mathcal{M}, \varphi, S, p, h)$.

Proof. The proof goes by induction on φ and is, essentially equivalent to the one for Theorem 3 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 $G_\downarrow(\mathcal{M}, \downarrow \psi, S, p, h)$. 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 $G_\downarrow(\mathcal{M}, \psi, \{s\}, p, h)$ that, by inductive hypothesis, she has. For the other direction, suppose Eloïse has a winning strategy for $G_\downarrow(\mathcal{M}, \downarrow \psi, S, p, h)$. This implies she has a winning strategy for $G_\downarrow(\mathcal{M}, \psi, \{s\}, p, h)$ for all $s \in S$: play whatever she would play as her second turn in $G_\downarrow(\mathcal{M}, \downarrow \psi, S, p, h)$ 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, $\mathcal{M} \models^- \downarrow \psi[S, p, h]$; this implies that for every $s \in S$, $\mathcal{M} \not\models^+ \psi[s, p, h]$. From here we derive a winning strategy for Abélard on the game $G_\downarrow(\mathcal{M}, \downarrow \psi, S, p, h)$ 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 $G_\downarrow(\mathcal{M}, \psi, \{s\}, p, h)$. 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 $G_\downarrow(\mathcal{M}, \downarrow \psi, S, p, h)$, 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 $G_\downarrow(\mathcal{M}, \psi, \{s\}, p, h)$ and, thus, by inductive hypothesis, $\mathcal{M} \not\models^+ \psi[s, p, h]$ and, finally, $\mathcal{M} \models^- \downarrow \psi[S, p, h]$. □

We can now use game semantics as an alternative tool to reason about the properties of $\downarrow$, as illustrated by the following two results.

Theorem 10. *Let ψ' be the formula obtained from ψ by removing from the slashed connectives every variable that occurs free in ψ. Then, for every S and every proper context for ψ, p and h, $\mathcal{M} \models^\pm \downarrow \psi[S, p, h]$ iff $\mathcal{M} \models^\pm \downarrow \psi'[S, p, h]$.*

Proof. The left-to-right implication is trivial: any winning strategy for the game $G_\downarrow(\mathcal{M}, \downarrow \psi, S, p, h)$ is also a winning strategy for $G_\downarrow(\mathcal{M}, \downarrow \psi', S, p, h)$. For the other direction, we show how to proceed in the case of Eloïse, for Abélard a similar argument applies. Let us suppose, then, $\mathcal{M} \models^+ \downarrow \psi'[S, p, h]$; this means that for all $s \in S$, $\mathcal{M} \models^+ \psi'[s, p, h]$. Now, take any arbitrary $s \in S$; Eloïse has a winning strategy for $G_\downarrow(\mathcal{M}, \psi', \{s\}, p, h)$, let us call it T_s. Consider the next strategy: "whenever T_s mandates to play a function f for a slashed connective and $x_1, \ldots, x_n$ are free variables of ψ that occur among the restrictions of the connective, play $f'(r) = f(r[p(x_1) \mapsto s(p(x_1)), \ldots, p(x_n) \mapsto s(p(x_n))])$ instead". It is straightforward to see that this is a winning strategy for $G_\downarrow(\mathcal{M}, \psi, \{s\}, p, h)$ and, thus, $\mathcal{M} \models^+ \downarrow \psi'[S, p, h]$. □

Observe that in the above theorem, $\downarrow \psi'$ may have fewer free variables than $\downarrow \psi$. In that case, they will not be equivalent formulas.

Theorem 11. *Let ψ be a formula such that in its syntax tree there is a path from the root to a $\downarrow$-node and every node in the path is either an $\sim$-node or is under an even number of $\sim$-nodes. Furthermore, let ψ' be the formula obtained by replacing that $\downarrow$-node in ψ by its descendent. Then, for all S and every proper context for ψ, p and h, $\mathcal{M} \models^\pm \downarrow \psi[S, p, h]$ iff $\mathcal{M} \models^\pm \downarrow \psi'[S, p, h]$.*

Proof. It suffices to observe that in this case, whatever Abélard plays in the second turn of the game $G_\downarrow(\mathcal{M}, \downarrow \psi, S, p, h)$, it cannot modify any node in the path between the root of ψ and the $\downarrow$-node in question. Thus there is nothing Eloïse can exploit if she gets the chance to redecorate the subtree below this node. □

In particular, the above result implies that if a formula contains two consecutive occurrences of $\downarrow$ without any occurrence of $\sim$ between them, then the second occurrence can be eliminated, leaving the meaning of the formula unaltered.

5. Conclusions

It is well known that IF-like formulas are very sensitive to context; much more than formulas in classical logics. The failure of basic properties (e.g. no α-equivalence, $\varphi \not\equiv \varphi \vee \varphi$, etc. [10]) witnessed to this fact. However, one of the points we tried to make in this paper is that so far there has been a disagreement in what is the context of an IF-like formula. This clearly follows, we believe, from Proposition 2.

Hodges [9] opted to split the context in two: syntactically, a list of variables in scope is a constituent of every formula; semantically, the value assigned to each variable is denoted by a list. In [16], the context is expressed only at the semantic level: the *finite* domain of valuations tries to capture the idea that only a finite number of variables have been bound. Similarly, we express the context of a formula only at the semantic level. Like in [9], our contexts contain a *finite list of values* (the prefixes of length h of S, cf. item 2 of Theorem 4) but since we don't want to restrict to regular formulas, we need to be able to interpret the free variables.

We have shown that the logic we obtain is an extension to the whole fragment of Hodges' trumps semantics (cf. Corollary 8). We can thus say that there are two alternative IF-like logics. Given the choice, we believe α-equivalence is a key property, one we would rather not sacrifice. For instance, because α-equivalence now holds, it is straightforward to take the flawed normal form results of [15] and show them to work under S-semantics. For most of the part one can simply reproduce Caicedo and Krynicki's original argument. Observe that they have also been fixed for V-semantics in [16] but it required a lot of technical effort (and yet, the equivalence results are context-dependent).

The last part of the paper was devoted to investigate the $\downarrow$-operator from a novel game-theoretical perspective. This lead us to some insight about its properties, specially when occurring nested and deeply inside a formula (cf. Theorems 10 and 11). Furthermore, the alternating selection of functions in the game-theoretical semantics for $\downarrow$ is suggesting that one can express alternating higher-order quantifiers although the precise expressive power of the operator still needs to be pinpointed.

Acknowledgments

References

[1] J. Hintikka, The Principles of Mathematics Revisited, Cambridge University Press, 1996.
[2] J. Hintikka, G. Sandu, Game-theoretical semantics, in: J. van Benthem, A. ter Meulen (Eds.), Handbook of Logic and Language, The MIT Press, 1997, Ch. 6.
[3] L. Henkin, Some remarks on infinitely long formulas, in: Infinitistic Methods: Proceedings of the Symposium on Foundations of Mathematics, Pergamon Press, 1961, pp. 167–183.
[4] W. Walkoe Jr., Finite partially-ordered quantification, J. Symbolic Logic 35 (4) (1970) 535–555.
[5] H. Enderton, Finite partially ordered quantifiers, Zeitschrift für Mathematische Logik und Grundlagen der Mathematik 16 (1970) 393–397.
[6] J. Barwise, On branching quantifiers in English, J. Philos. Logic 8 (1979) 47–80.
[7] G. Sandu, On the logic of informational independence and its applications, J. Philos. Logic 22 (1) (1993) 29–60.
[8] J.A. Väänänen, Second-order logic and foundations of mathematics, Bull. Symbolic Logic 7 (4) (2001) 504–520.
[9] W. Hodges, Compositional semantics for a language of imperfect information, Log. J. IGPL 5 (4) (1997) 539–563.
[10] T.M.V. Janssen, Independent choices and the interpretation of IF logic, J. Logic Lang. Inform. 11 (3) (2002) 367–387.
[11] W. Hodges, Logics of imperfect information: Why sets of assignments? in: Proceedings of the 7th Augustus de Morgan Workshop "Interactive Logic: Games and Social Software", 2005.
[12] W. Hodges, Some strange quantifiers, in: J. Mycielski, G. Rozenberg, A. Salomaa (Eds.), Structures in Logic and Computer Science, A Selection of Essays in Honor of Andrzej Ehrenfeucht, in: Lecture Notes in Computer Science, vol. 1261, Springer, 1997, pp. 51–65.
[13] F. Dechesne, Game, sets, maths: Formal investigations into logic with imperfect information, Ph.D. thesis, Department of Philosophy, University of Tilburg, The Netherlands, 2005.
[14] T.M.V. Janssen, F. Dechesne, Signalling in IF games: a tricky business, in: The Age of Alternative Logics, Springer, 2006, pp. 221–241, Ch. 15.
[15] X. Caicedo, M. Krynicki, Quantifiers for reasoning with imperfect information and Σ_1^1-logic, Contemp. Math. 235 (1999) 17–31.
[16] X. Caicedo, F. Dechesne, T.M. Janssen, Equivalence and quantifier rules for logic with imperfect information, Log. J. IGPL 17 (1) (2009) 91–129, doi:10.1093/jigpal/jzn030.
[17] N.G. de Bruijn, Lambda calculus notation with nameless dummies. A tool for automatic formula manipulation with application to the Church–Rosser theorem, Indag. Math. 34 (1972) 381–392.
[18] H. Barendregt, The Lambda Calculus: Its Syntax and Semantics, 2nd ed., Studies in Logic and the Foundations of Mathematics, vol. 103, North-Holland, 1985.
[19] J.A. Väänänen, On the semantics of informational independence, Log. J. IGPL 10 (3) (2002) 339–352.