

CLASSIFYING TOPOI AND FINITE FORCING

Andreas BLASS*

The University of Michigan, Ann Arbor, MI 48109, USA

Andrej ŠČEDROV

University of Pennsylvania, Philadelphia, PA 19104, USA

Communicated by P. Freyd

Received 3 November 1982

We show that Robinson's finite forcing, for a theory \mathcal{T} , is a universal construction in the sense of categorical algebra: it is the satisfaction relation for the universal model in the classifying topos \mathcal{E} of a certain universal Horn theory defined from \mathcal{T} . Assuming, without loss of generality, that \mathcal{T} is axiomatized by universal sentences, we construct, as sheaf subtopoi of \mathcal{E} , the classifying topoi for (i.e., universal examples of) finitely generic models, existentially closed models, and arbitrary models of \mathcal{T} (with complemented primitive predicates).

AMS 1980 Subject Classifications: 03C25, 03G30, 18B25.

Keywords: Finite forcing, Generic model, Classifying topos, Sheaf semantics, Existentially closed model, Atomic model.

Our purpose in this paper is to investigate the connection between the concepts of finite forcing in model theory [1, 7] and classifying topoi in category theory [8, 13, 15, 14]. After a preliminary section on the classifying topoi of universal Horn theories, we establish in Section 2 that the forcing relation for a theory \mathcal{T} is essentially the same as the satisfaction relation for the universal model in the classifying topos of a related universal Horn theory \mathcal{T}_{dVH} . In Sections 3 and 4 we give explicit constructions of the classifying topoi for the finitely generic models and the existentially closed models of a universal theory \mathcal{T} . We also discuss the relationship of these topoi to each other, to the classifying topos of (a theory classically equivalent to) \mathcal{T} , and to their common subtopos of double-negation sheaves.

For notation and background information in model theory and category theory, we refer to [4] and [12, 8] respectively, but we briefly review some of the topos-theoretic concepts that we shall need. We use the word 'topos' to mean Grothendieck topos [6, 8, 13], i.e., the category of sheaves on a site (\mathcal{C}, J) where \mathcal{C} is a small category and J is a Grothendieck topology on it. A geometric morphism $f: \mathcal{E} \rightarrow \mathcal{F}$ consists of two functors, $f_*: \mathcal{E} \rightarrow \mathcal{F}$ and its left adjoint $f^*: \mathcal{F} \rightarrow \mathcal{E}$ such that the 'inverse' part f^* is left exact. A natural transformation $f \rightarrow g$ is defined to be a natural transformation $f^* \rightarrow g^*$ (or equivalently, by adjointness, $g_* \rightarrow f_*$). The internal logic

* Partially supported by NSF grant MCS-8101560.

of topoi [3, 8 §5.4] permits us to define the notion of a model, in a topos \mathcal{E} , of a first order theory \mathcal{T} (possibly involving infinite conjunctions and disjunctions); all the axioms and rules of intuitionistic predicate logic are sound for such models. Geometric morphisms do not generally preserve the semantics of models, but the inverse parts f^* do preserve the truth values of existential positive formulas, also called coherent or geometric formulas (in finitary logic; infinite disjunctions are also permissible but infinite conjunctions are not). It follows that, if \mathcal{T} is a geometric theory, that is, one axiomatized by sentences $\forall x(\phi(x) \rightarrow \psi(x))$ where ϕ and ψ are existential positive formulas, then f^* sends models of \mathcal{T} to models of \mathcal{T} .

A classifying topos for a geometric theory \mathcal{T} is a topos $\mathcal{E}(\mathcal{T})$ such that, for any topos \mathcal{E} , the category of models of \mathcal{T} in \mathcal{E} (and homomorphisms) is equivalent, naturally in \mathcal{E} , to the category of geometric morphisms $\mathcal{E} \rightarrow \mathcal{E}(\mathcal{T})$ (and natural transformations). Thus $\mathcal{E}(\mathcal{T})$ contains a model \mathcal{M} of \mathcal{T} with the universal property that any model of \mathcal{T} in any topos \mathcal{E} is (isomorphic to) $f^*\mathcal{M}$ for a unique (up to natural isomorphism) geometric morphism $f: \mathcal{E} \rightarrow \mathcal{E}(\mathcal{T})$. This \mathcal{M} is usually called the generic model of \mathcal{T} , but we shall call it the universal model of \mathcal{T} to avoid conflict with the terminology ‘generic’ in forcing theory. (Fortunately, we shall not need the model-theoretic concept of universal model.) Every geometric theory has a classifying topos [15, 8 §6.5, 7.4], and conversely, if we allow our geometric theories to be infinitary and multi-sorted, then every topos classifies some such theory. In particular, Diaconescu’s theorem [5, 8 §4.3] asserts that the topos of presheaves on a small category \mathcal{C} classifies flat functors on \mathcal{C} . (The definition of flatness is in Section 1.) More generally, the topos of sheaves on a site (\mathcal{C}, J) classifies flat functors that are continuous in the sense that the covering families of J are sent to epimorphic families.

In Section 1, we shall construct the classifying topos of a universal Horn theory as a topos of presheaves over a certain syntactically defined category \mathcal{C} whose dual is the category of finitely presented models of the theory. In Section 2, we describe the satisfaction relation for the universal model of such a theory in terms of a concept of pseudo-forcing that essentially agrees with Robinson’s concept of finite forcing [1, 7] except that no negations are allowed in the forcing conditions. We then construct for any universal theory \mathcal{T} , a universal Horn theory \mathcal{T}_{dVH} such that pseudo-forcing for \mathcal{T}_{dVH} is essentially the same as finite forcing for \mathcal{T} . Thus, finite forcing for \mathcal{T} is identified with satisfaction in the universal model of \mathcal{T}_{dVH} . In Section 3, we define (finitely) generic models of \mathcal{T} in arbitrary topoi, we prove the ‘forcing equals truth’ lemma for these models, and we construct the classifying topos for generic \mathcal{T} -models as the topos of sheaves on a certain (\mathcal{C}, J_R) , where \mathcal{C} is the syntactic category associated to \mathcal{T}_{dVH} . Finally, in Section 4, we relate the topology J_R to other topologies J_D and J_E on \mathcal{C} whose sheaf topoi classify the models and the existentially closed models of (a geometric theory classically equivalent to) \mathcal{T} . We also discuss the double-negation topology on \mathcal{C} .

Loyal and Reyes [9] have discussed, from the point of view of categorical logic, some of the same concepts treated here. In particular, they define existentially

closed models and generic models – but not forcing – in this context. They do not, however, relate these definitions to the traditional ones in model theory. In fact, their ‘generic’ corresponds to what we call ‘pseudo-generic’ rather than to ordinary ‘finitely generic’; thus; for example, their definition, applied to the theory of groups, yields that the only generic group is the trivial one (see the end of Section 3). Nevertheless, it seems likely that their definition of ‘generic’ was motivated by something resembling our Theorem 2, perhaps in the context of pretopoi (the theories themselves, in categorical logic) rather than classifying topoi.

As a final introductory point, we mention some matters of notation. We systematically use boldface letters, like x , to abbreviate finite sequences (also called lists or tuples) x_1, \dots, x_n whose length n is usually not specified; we tacitly assume, of course, that the lengths of various sequences agree whenever the context requires this, e.g., when one sequence is to be substituted for another. If x is as above and we refer to a formula $\phi(x)$ or a term $t(x)$, we intend that all free variables of ϕ or t are among x , but we not require all (or even any) of x to actually occur in ϕ or t . We write $\forall x \phi(x)$ for $\forall x_1 \dots \forall x_n \phi(x)$, and we write $x=y$ for $(x_1=y_1) \wedge \dots \wedge (x_n=y_n)$. We do not distinguish between formulas or terms that differ only by a renaming of bound variables, subject to the usual conventions for avoiding clashes.

1. The classifying topos of a universal Horn theory

A universal Horn theory is a theory \mathcal{H} axiomatized by sentences of the form $\forall x(\phi \rightarrow \psi)$ where ϕ is a conjunction of atomic formulas (possibly the empty conjunction, *true*) and ψ is either an atomic formula or *false*. These are precisely the theories whose classes of models are closed under substructures and under direct products of one or more factors [4 §6.2]; we shall need only the easy half of this result, namely that universal Horn theories have these preservation properties. If we had not permitted *false* to occur as ψ in an axiom, then the class of models would have been closed under arbitrary products, even the empty product (a one element structure in which all primitive predicates hold of the unique tuple). This change would simplify much of this section but would exclude the intended applications in the next section. The preservation properties of universal Horn theories imply the following lemma, in which the hypothesis that the language of \mathcal{L} has a constant symbol is used only to avoid the need to consider empty structures; this hypothesis is not essential here but it will be important in some of our later results.

Lemma 1. *Let \mathcal{H} be a universal Horn theory in a language with at least one constant symbol. Let $\phi(x), \alpha_1(x, y), \dots, \alpha_n(x, y)$, with $n \geq 1$, be conjunctions of atomic formulas. If*

$$\forall x \left(\phi(x) \rightarrow \bigvee_{i=1}^n \exists y \alpha_i(x, y) \right) \quad (1)$$

is provable in \mathcal{H} , then so is

$$\forall x(\phi(x) \rightarrow \alpha_i(x, t(x))) \quad (2)$$

for some i and some terms $t(x)$ with only x free.

Proof. Suppose not. For each i and each list of appropriate terms t , let $\mathcal{A}_{i,t}$ be a model of \mathcal{H} in which (2) is false. Choose witnesses $a_{i,t} \in \mathcal{A}_{i,t}$ attesting to the failure of (2), so

$$\mathcal{A}_{i,t} \models \phi(a_{i,t}) \wedge \neg \alpha_i(a_{i,t}, t(a_{i,t})). \quad (3)$$

Let \mathcal{A} be the direct product of all of the $\mathcal{A}_{i,t}$, and let a be the elements of \mathcal{A} whose (i,t) -components are the elements $a_{i,t}$ of $\mathcal{A}_{i,t}$. Let \mathcal{B} be the substructure of \mathcal{A} generated by the elements a . Then \mathcal{B} is a model of \mathcal{H} because \mathcal{H} is a universal Horn theory, and it satisfies $\phi(a)$ because ϕ is a conjunction of atomic formulas and (3) holds. Therefore, by (1), there exists an i such that $\mathcal{B} \models \exists y \alpha_i(a, y)$. By definition of \mathcal{B} , this means that there are terms $t(x)$ such that $\mathcal{B} \models \alpha_i(a, t(a))$. Since α_i is a conjunction of atomic formulas, it is preserved by the embedding $\mathcal{B} \hookrightarrow \mathcal{A}$ and the projection $\mathcal{A} \rightarrow \mathcal{A}_{i,t}$, so we have $\mathcal{A}_{i,t} \models \alpha_i(a_{i,t}, t(a_{i,t}))$, contradicting (3). \square

We are now ready to construct the classifying topos for a universal Horn theory. The construction, as a presheaf topos, is a straightforward generalization of the corresponding construction for algebraic theories outlined in [13 §9.4]; it also has connections, which we shall explore later, with the construction, presented in [8 §7.4], of the classifying topos for a geometric theory, i.e., one axiomatized by sentences $\forall x(\phi(x) \rightarrow \psi(x))$ where ϕ and ψ are existential positive formulas. Of course, since universal Horn theories are geometric, the latter construction could be applied directly, but our construction is considerably simpler; in particular, we need only presheaves, not sheaves.

Let \mathcal{H} be a universal Horn theory in a language with at least one constant symbol. We define a category \mathcal{C} , a simplified version of the syntactic category in [8 §7.4], as follows. An object of \mathcal{C} is formal class term $\{x \mid \phi(x)\}$, where x is a list of variables and $\phi(x)$ is a conjunction of atomic formulas, with free variables among x , such that $\phi(x)$ is consistent with \mathcal{H} . The variables x are, of course, considered bound in $\{x \mid \phi(x)\}$, so this class term is unchanged by any proper substitution of other variables for x . To define morphisms from $\{x \mid \phi(x)\}$ to $\{y \mid \psi(y)\}$, we assume, by making a substitution if necessary, that the lists x and y are disjoint. A morphism is then given by a system of equations, $y = t(x)$, expressing each of the variables in y as a term involving only free variables from x , such that

$$\mathcal{H} \vdash \forall x(\phi(x) \rightarrow \psi(t(x)));$$

another such system, $y = t'(x)$, defines the same morphism if

$$\mathcal{H} \vdash \forall x(\phi(x) \rightarrow t(x) = t'(x)).$$

Composition is defined by substitution:

$$[z = u(y)] \circ [y = t(x)] = [z = u(t(x))];$$

it is easily verified that \mathcal{C} is a category.

It will be useful to single out a class of morphisms of \mathcal{C} , which seem to be quite special but are in fact, as we shall see, fairly representative of morphisms in general. Suppose $\{x, y \mid \phi(x, y)\}$ and $\{y \mid \psi(y)\}$ are objects of \mathcal{C} such that

$$\mathcal{M} \vdash \forall x, y (\phi(x, y) \rightarrow \psi(y)).$$

Then there is a morphism between these objects given, after the bound variables are renamed so as to be distinct, by

$$[y' = y] : \{x, y \mid \phi(x, y)\} \rightarrow \{y' \mid \psi(y')\}.$$

We call this the *simple* morphism from $\{x, y \mid \phi(x, y)\}$ to $\{y \mid \psi(y)\}$; this involves a harmless abuse of language in that we are using the specific choice of bound variables (the y used in both objects) to indicate which simple morphism is meant. The following lemma shows that, in a certain sense, every morphism is equivalent to a simple one.

Lemma 2. *Every morphism $A \rightarrow B$ of \mathcal{C} can be factored as $A \rightarrow A' \rightarrow B$ where $A \rightarrow A'$ is an isomorphism and $A' \rightarrow B$ is simple.*

Proof. Let the given morphism be $[y = t(x)] : \{x \mid \phi(x)\} \rightarrow \{y \mid \psi(y)\}$. The factorization is given by

$$[x' = x, y = t(x)] : \{x \mid \phi(x)\} \rightarrow \{x', y \mid \phi(x') \wedge y = t(x')\}$$

and the simple morphism

$$\{x', y \mid \phi(x') \wedge y = t(x)\} \rightarrow \{y \mid \psi(y)\}.$$

The first of these is an isomorphism with (simple) inverse $[x = x']$. We leave the verification of the details to the reader. \square

The syntactic objects and morphisms of \mathcal{C} have a natural semantic interpretation as presentations of models of \mathcal{M} and homomorphisms between them. Specifically, an object $\{x \mid \phi(x)\}$ determines a model $\langle x \mid \phi(x) \rangle$ of \mathcal{M} , generated by the (formal symbols) x subject to the relations $\phi(x)$. The elements of $\langle x \mid \phi(x) \rangle$ are equivalence classes of terms $t(x)$ modulo the equivalence relation

$$\mathcal{M} \vdash \forall x (\phi(x) \rightarrow t(x) = t'(x)).$$

Function symbols are interpreted in $\langle x \mid \phi(x) \rangle$ in the obvious way, and relation symbols R are interpreted to hold of (the equivalence classes of) $t(x)$ if and only if

$$\mathcal{M} \vdash \forall x (\phi(x) \rightarrow R(t(x))).$$

It is easy to check that $\langle x \mid \phi(x) \rangle$ is well defined, that it is generated by (the equivalence classes of) x , and that a conjunction of atomic formulas holds of certain elements in $\langle x \mid \phi(x) \rangle$ if and only if its holding is deducible in \mathcal{H} from $\phi(x)$. It follows that $\langle x \mid \phi(x) \rangle$ is a model of \mathcal{H} and its homomorphisms f into arbitrary models \mathcal{A} of \mathcal{H} are in canonical one-to-one correspondence with lists a of elements satisfying $\phi(a)$ in \mathcal{A} , the correspondence being that a is the image of x under f . In other words, $\langle x \mid \phi(x) \rangle$ is freely generated by x subject to $\phi(x)$.

Any morphism

$$[y = t(x)] : \{x \mid \phi(x)\} \rightarrow \{y \mid \psi(y)\}$$

in \mathcal{C} defines a homomorphism

$$\langle y = t(x) \rangle : \langle y \mid \psi(y) \rangle \rightarrow \langle x \mid \phi(x) \rangle$$

of \mathcal{H} -models, namely the homomorphism sending (the equivalence class of) $u(y)$ to (that of) $u(t(x))$. It is an easy consequence of the definitions that $\langle \rangle$ is a fully faithful contravariant functor from \mathcal{C} to the category of models of \mathcal{H} ; its image consists of the *finitely presented* models of \mathcal{H} . Thus, \mathcal{C} is the dual of the category of finitely presented models of \mathcal{H} .

In showing that $\mathcal{C}^{(\text{op})}$ classifies models of \mathcal{H} , we shall want to know that \mathcal{C} has finite limits. Unfortunately, it might not. In the first place, \mathcal{C} would not have a terminal object, or any object at all, if \mathcal{H} were inconsistent. We leave this trivial case to the reader and assume \mathcal{H} is consistent. Then \mathcal{C} has a terminal object $\{ \mid true \}$. (We never assumed that the list of variables x in an object $\{x \mid \phi(x)\}$ is nonempty.) Pullbacks and even products do not exist in general. Indeed, if $\phi(x)$ and $\psi(y)$, with disjoint sets of variables x and y , are individually consistent but not jointly consistent in \mathcal{H} , then $\{x \mid \phi(x)\}$ and $\{y \mid \psi(y)\}$ are objects with no product, since no object admits morphisms to both of them. For pullbacks in general, the situation is this. A pair of morphisms

$$\begin{array}{ccc}
 & \{y \mid \psi(y)\} & \\
 & \downarrow [z = u(y)] & \\
 \{x \mid \phi(x)\} & \xrightarrow{[z = t(x)]} & \{z \mid \theta(z)\}
 \end{array} \tag{4}$$

has a pullback if and only if it can be completed to a commutative square, if and only if $\phi(x) \wedge \psi(y) \wedge t(x) = u(y)$ is consistent with \mathcal{H} when x and y are disjoint lists. In this case, the pullback is given by the object $\{x, y \mid \phi(x) \wedge \psi(y) \wedge t(x) = u(y)\}$ with simple morphisms to $\{x \mid \phi(x)\}$ and $\{y \mid \psi(y)\}$. Notice in particular that the pullback exists whenever the two morphisms in (4) are the same; thus the usual characterization of monomorphisms in terms of pullbacks works in \mathcal{C} .

Theorem 1. $\mathcal{C}^{(\text{op})}$ is the classifying topos for the universal Horn theory \mathcal{H} .

Proof. By Diaconescu's theorem [5, 8 §4.3], geometric morphisms from any Grothendieck topos \mathcal{F} to $\mathcal{P}^{\mathcal{C}(\text{op})}$ correspond to flat internal functors on \mathcal{C} in \mathcal{F} . We must prove that these functors correspond (functorially and naturally) to models of \mathcal{H} in \mathcal{F} . For simplicity, we shall prove this under the assumption that $\mathcal{F} = \mathcal{S}$, but our argument can be made to apply to arbitrary \mathcal{F} by interpreting suitable parts of it in the internal logic of \mathcal{F} .

Consider an arbitrary flat functor F on \mathcal{C} . Flatness means that

(a) given finitely many (possibly zero) objects A_i of \mathcal{C} and elements $a_i \in F(A_i)$, we have an object B , morphisms $\alpha_i : B \rightarrow A_i$, and an element $b \in F(B)$ such that $F(\alpha_i)(b) = a_i$ for all i , and

(b) given $a \in F(A)$ and finitely many morphisms $\alpha_i : A \rightarrow A'$ such that all of the $F(\alpha_i)(a)$ are equal, we have a morphism $\beta : B \rightarrow A$ such that all of the composites $\alpha_i \circ \beta : B \rightarrow A'$ are equal and a is in the image of $F(\beta)$.

An easy consequence of (a) and (b) is that

(c) given finitely many morphisms $\alpha_i : A_i \rightarrow A'$ and elements $a_i \in F(A_i)$ such that all of the $F(\alpha_i)(a_i)$ are equal, we have an object B , morphisms $\beta_i : B \rightarrow A_i$, and an element $b \in F(B)$ such that all of the composites $\alpha_i \circ \beta_i$ are equal and $F(\beta_i)(b) = a_i$ for all i .

Another well-known and fairly easy consequence of flatness is that F preserves finite limits. Using these facts, we can construct a model \mathcal{M} from a flat functor F on \mathcal{C} as follows:

The universe M of \mathcal{M} is $F\{x \mid \text{true}\}$. It follows, by preservation of finite limits, that $M^n \cong F\{y \mid \text{true}\}$, where y is a list of n variables. The \mathcal{M} -interpretation of an n -ary function symbol f is $F[x = f(y)] : F\{y \mid \text{true}\} \rightarrow F\{x \mid \text{true}\}$. The interpretation of an n -ary relation symbol R is the subobject $F\{y \mid R(y)\} \hookrightarrow F\{y \mid \text{true}\}$ obtained by applying F to the trivial morphism (easily seen to be a monomorphism), provided $R(y)$ is consistent with \mathcal{H} so that $\{y \mid R(y)\}$ is an object of \mathcal{C} . If $R(y)$ is inconsistent with \mathcal{H} then we interpret R in \mathcal{M} as the empty relation.

It is easy to verify, by induction on terms, that the interpretation in \mathcal{M} of a term $t(x)$ is

$$F[y = t(x)] : M^n = F\{x \mid \text{true}\} \rightarrow F\{y \mid \text{true}\} = M,$$

where n is the number of variables in the list x . We shall show that, if $\phi(x)$ is a conjunction of atomic formulas, then its extension in \mathcal{M} , $\{a \in M^n \mid \mathcal{M} \models \phi(a)\}$ is the subobject of $M^n = F\{x \mid \text{true}\}$ obtained by applying F to the simple morphism $\{x \mid \phi(x)\} \rightarrow \{x \mid \text{true}\}$ if this morphism exists, i.e., if $\phi(x)$ is consistent with \mathcal{H} , and this extension is empty if $\phi(x)$ is inconsistent with \mathcal{H} . To prove this, suppose first that $\phi(x)$ is consistent with \mathcal{H} , and let it be the conjunction of n atomic formulas $R_i(t_i(x))$. Then it follows, from our description of pullbacks in \mathcal{C} , that $\{x \mid \phi(x)\}$ is the limit of the diagram consisting of the n objects $\{x \mid R_i(x)\}$ together with their simple morphisms to $\{x \mid \text{true}\}$; it also follows that each of these simple morphisms is the pullback along $[y_i = t_i(x)] : \{x \mid \text{true}\} \rightarrow \{y_i \mid \text{true}\}$, of a simple morphism $\{y_i \mid R_i(y_i)\} \rightarrow \{y_i \mid \text{true}\}$. The latter simple morphisms define, via F , the interpreta-

tions in \mathcal{M} of the R_i , so, since F preserves finite limits, it sends the former simple morphisms to the extensions of the formulas $R_i(t_i(x))$, and therefore sends $\{x \mid \phi(x)\}$ to the extension of $\phi(x)$, as claimed. There remains the case that $\phi(x)$ is inconsistent with \mathcal{M} . In this case, some of the limits in the preceding argument fail to exist; suppose, for concreteness, that all the $\{x \mid R_i(t_i(x))\}$ exist but $\{x \mid \phi(x)\}$ does not, i.e., each $R_i(t_i(x))$ is consistent with \mathcal{M} but their conjunction is not. (The other cases are (i) that some $R_i(t_i(x))$ is inconsistent with \mathcal{M} but $R_i(y_i)$ is consistent and (ii) that some $R_i(y_i)$ is inconsistent with \mathcal{M} . Case (i) is similar to the case to be treated in detail, and case (ii) is trivial since R_i is interpreted as the empty relation in \mathcal{M} .) Suppose, toward a contradiction, that \mathbf{a} are elements satisfying ϕ in \mathcal{M} . Then, for each i , \mathbf{a} is in the extension of $R_i(t_i(x))$, which we already know to be the subobject

$$F(\alpha_i) : F\{x \mid R_i(t_i(x))\} \hookrightarrow F\{x \mid \text{true}\}.$$

where α_i is a simple morphism; let \mathbf{a}'_i be such that $F(\alpha_i)\mathbf{a}'_i = \mathbf{a}$. By the flatness of F , specifically by (c) above, we can find an object $\{y \mid \psi(y)\}$, morphisms

$$\beta_i = [x = u_i(y)] : \{y \mid \psi(y)\} \rightarrow \{x \mid R_i(t_i(x))\},$$

and an element b of $F\{y \mid \psi(y)\}$ such that all the composites $\alpha_i \circ \beta_i$ are equal and $F(\beta_i)(b) = \mathbf{a}'_i$ for every i . The equality of the composites means that, for all i and j ,

$$\mathcal{M} \vdash \forall y (\psi(y) \rightarrow u_i(y) = u_j(y)),$$

and the definition of morphism in \mathcal{C} yields

$$\mathcal{M} \vdash \forall y (\psi(y) \rightarrow R_i(t_i(u_i(y)))).$$

Combining these results, we have, for each i ,

$$\mathcal{M} \vdash \forall y (\psi(y) \rightarrow \phi(u_i(y))).$$

But this is absurd, since ψ is consistent and ϕ is not. This contradiction completes the proof of our description of $\{\mathbf{a} \in M^n \mid \mathcal{M} \models \phi(\mathbf{a})\}$.

Consider an arbitrary axiom of \mathcal{M} , say $\forall x (\phi(x) \rightarrow \psi(x))$. If $\psi(x)$ is *false*, then $\phi(x)$ is inconsistent with \mathcal{M} , and the result just obtained shows that the axiom is true in \mathcal{M} . On the other hand, if ψ is atomic, then we can apply the result just obtained to both ϕ and ψ . In view of the commutative diagram of simple morphisms

$$\begin{array}{ccc} \{x \mid \phi(x)\} & \xrightarrow{\quad} & \{x \mid \text{true}\}, \\ & \searrow & \nearrow \\ & \{x \mid \psi(x)\} & \end{array}$$

we have that the extension of ϕ is included in that of ψ , so again the axiom holds. Therefore \mathcal{M} is a model of \mathcal{M} .

Conversely, given any model \mathcal{M} of \mathcal{M} , we can define a flat functor F on \mathcal{C} by taking $F\{x \mid \phi(x)\}$ to be the extension in \mathcal{M} of $\phi(x)$ and taking

$$F[y = t(x)] : F\{x \mid \phi(x)\} \rightarrow F\{y \mid \psi(y)\}$$

to be the map that sends any a satisfying ϕ to $t(a)$ (which satisfies ψ by definition of morphisms). It is easy to check that F is a functor. It is flat because (a) if we are given $a_i \in F\{x_i \mid \phi_i(x_i)\}$ for $i = 1, 2, \dots, n$, then they are all images, under F of simple morphisms, of the concatenated sequence a_1, a_2, \dots, a_n in

$$F\left\{x_1, x_2, \dots, x_n \mid \bigwedge_{i=1}^n \phi_i(x_i)\right\},$$

where we have taken the sequences of bound variables x_i to be disjoint, and (b) if we are given $a \in F\{x \mid \phi(x)\}$ and finitely many morphisms $\alpha_i = [y = t_i(x)] : \{x \mid \phi(x)\} \rightarrow \{y \mid \psi(y)\}$ such that all the $F(\alpha_i)(a)$ are equal, then a satisfies $t_i = t_j$ for each pair i, j and therefore lies in the image of F of the simple morphism

$$\beta : \left\{x \mid \phi(x) \wedge \bigwedge_{i,j} t_i(x) = t_j(x)\right\} \rightarrow \{x \mid \phi(x)\}$$

whose composites with all the α_i are equal. (Note that, in each part of this proof of flatness, the properties of the given elements guarantee the consistency with \mathcal{H} of the formula that defines the required object.)

We leave to the reader the straightforward verifications that the constructions of \mathcal{H} from F and vice versa are functorial and that they are inverse to each other up to natural isomorphism. These verifications complete the proof of Theorem 1. \square

Since universal Horn theories are a special sort of geometric theories, their classifying topoi can also be obtained by the general construction of Joyal and Reyes presented in §7.4 of [8]. It seems worthwhile to compare this construction with $\mathcal{V}^{(\text{op})}$; the comparison will yield a second, less direct, proof of Theorem 1.

The site (\mathcal{C}_g, J_g) used in the Joyal–Reyes construction has as objects class terms $\{x \mid \phi(x)\}$ where ϕ is an existential positive formula (not necessarily consistent with \mathcal{H}). Its morphisms from $\{x \mid \phi(x)\}$ to $\{y \mid \psi(y)\}$ are given by $[x \mapsto y \mid \theta(x, y)]$ where θ is an existential positive formula such that the sentences saying “ θ defines a single-valued function from $\{x \mid \phi(x)\}$ to $\{y \mid \psi(y)\}$ ” are provable in \mathcal{H} . Two θ ’s define the same morphism if they are \mathcal{H} -provably equivalent. (For details, see [8].) A sieve covers $\{y \mid \psi(y)\}$ if and only if it contains finitely many morphisms

$$[x_i \mapsto y \mid \theta_i(x_i, y)] : \{x_i \mid \phi_i(x_i)\} \rightarrow \{y \mid \psi(y)\}$$

such that

$$\mathcal{H} \vdash \forall y (\psi(y) \rightarrow \bigvee_i \exists x_i \theta_i(x_i, y)).$$

(In particular, the empty sieve covers $\{y \mid \psi(y)\}$ if and only if $\psi(y)$ is inconsistent with \mathcal{H} .)

There is a functor from \mathcal{C} to \mathcal{C}_g , sending each object $\{x \mid \phi(x)\}$ to itself and sending a morphism $[y = t(x)]$ to $[x \mapsto y \mid \phi(x) \wedge y = t(x)]$, where $\{x \mid \phi(x)\}$ is the domain of the morphism. The definitions of equality of morphisms in \mathcal{C} and \mathcal{C}_g

easily imply that this functor is faithful. In fact, it is also full. To see this, suppose

$$[x \mapsto y \mid \theta(x, y)] : \{x \mid \phi(x)\} \rightarrow \{y \mid \psi(y)\} \quad (5)$$

is a morphism in \mathcal{C}_g between objects of \mathcal{C} . This implies that

$$\mathcal{A} \vdash \forall x(\phi(x) \rightarrow \exists y \theta(x, y)).$$

Since θ is existential and positive, it is logically equivalent to a formula of the form $\bigvee_{i=1}^n \exists z \theta_i(x, y, z)$ where each θ_i is a conjunction of atomic formulas. (To put θ in this form, first put it in prenex form, then put the matrix in disjunctive normal form, and finally distribute existential quantifiers over disjunctions.) Since $\{x \mid \phi(x)\}$ is an object of \mathcal{C} , ϕ is consistent with \mathcal{A} , so $n \geq 1$ and Lemma 1 is applicable to

$$\mathcal{A} \vdash \forall x \left(\phi(x) \rightarrow \bigvee_{i=1}^n \exists y \exists z \theta_i(x, y, z) \right).$$

This lemma provides an index i and terms $t(x), u(x)$ such that

$$\mathcal{A} \vdash \forall x(\phi(x) \rightarrow \theta_i(x, t(x), u(x)))$$

and therefore

$$\mathcal{A} \vdash \forall x(\phi(x) \rightarrow \theta(x, t(x))).$$

Since θ \mathcal{A} -provably defines a function, it easily follows that it defines the same function as $\phi(x) \wedge y = t(x)$. Thus, the morphism (5) is the image, under our functor, of $[y = t(x)]$. Therefore, this functor identifies \mathcal{C} with a full subcategory of \mathcal{C}_g .

Every object $\{x \mid \phi(x)\}$ of \mathcal{C}_g is J_g -covered by morphisms whose domains are in \mathcal{C} . Indeed, transforming ϕ to the form $\bigvee_i \exists y \phi_i(x, y)$ where the ϕ_i are conjunctions of atomic formulas (as we did with θ in the preceding paragraph), we see that $\{x \mid \phi(x)\}$ is covered by simple morphisms from the objects $\{x, y \mid \phi_i(x, y)\}$ of \mathcal{C} .

What we have shown about the connection between \mathcal{C} and \mathcal{C}_g implies, by the comparison lemma [6, III.4.1], that the topos of sheaves on (\mathcal{C}_g, J_g) , the Joyal-Reyes version of the classifying topos of \mathcal{A} , is equivalent to the topos of sheaves on \mathcal{C} with the topology induced by J_g . To complete the identification of this form with $\mathcal{C}^{(\mathcal{C}^{\text{op}})}$, we still need to see that this induced topology is trivial, i.e., that every covering sieve contains the identity.

Suppose, therefore, that we have an object $\{y \mid \psi(y)\}$ of \mathcal{C} covered by a sieve R in the induced topology. By definition, this means that R contains finitely many morphisms

$$[y = t_i(x_i)] : \{x_i \mid \phi_i(x_i)\} \rightarrow \{y \mid \psi(y)\}$$

such that

$$\mathcal{A} \vdash \forall y \left(\psi(y) \rightarrow \bigvee_i \exists x_i \left(\phi_i(x_i) \wedge y = t_i(x_i) \right) \right).$$

As ψ is consistent with \mathcal{A} , the disjunction here cannot be empty, so Lemma 1 gives

us an index i and terms u such that

$$\mathcal{H} \vdash \forall y (\psi(y) \rightarrow \phi_i(u(y)) \wedge y = t_i(u(y))).$$

But then the morphism $[y = t_i(x_i)]$ in our covering sieve R , composed with $[x_i = u(y)]$, yields the identity morphism of $\{y \mid \psi(y)\}$, which is therefore also in R , as required.

The proof of Theorem 1 provides an explicit description of the universal model \mathcal{G} of \mathcal{H} in $\mathcal{S}^{\mathcal{C}^{\text{op}}}$, i.e., the model that corresponds to the identity geometric morphism on $\mathcal{S}^{\mathcal{C}^{\text{op}}}$. Indeed, it is well known [8 §4.3] that the flat functor corresponding to the identity morphism is the Yoneda embedding $Y: \mathcal{C} \rightarrow \mathcal{S}^{\mathcal{C}^{\text{op}}}$. Applying the proof of Theorem 1 with Y in place of F , we find that the underlying object G of the universal model \mathcal{G} is the presheaf $G = Y\{x \mid \text{true}\}$ whose value at an arbitrary object $\{y \mid \phi(y)\}$ is

$$\text{Hom}_{\mathcal{C}}(\{y \mid \phi(y)\}, \{x \mid \text{true}\}).$$

But an element $[x = t(y)]$ of this Hom-set is determined by an arbitrary term $t(y)$ in the variables y , two terms yielding the same element if and only if $\mathcal{H} \vdash \forall y (\phi(y) \rightarrow t(y) = t'(y))$. Thus, $G(\{y \mid \phi(y)\})$ is (in canonical one-to-one correspondence with) the underlying set of the finitely presented \mathcal{H} -model $\langle y \mid \phi(y) \rangle$. It is easy to verify that this correspondence respects the \mathcal{H} -model structure. Therefore, if we identify \mathcal{C}^{op} with the category of finitely presented \mathcal{H} -models, then $\mathcal{G} \in \mathcal{S}^{\mathcal{C}^{\text{op}}}$ is simply the underlying set functor on \mathcal{C}^{op} , equipped with its natural \mathcal{H} -model structure.

2. Forcing

Our objective in this section is to relate Robinson's concept of finite forcing in model theory [1, 7, 10] to classifying topoi. Our first step is to introduce 'pseudo-forcing', a concept that has some of the flavor of Robinson's forcing but is significantly different from it; the usefulness of this concept lies in its direct connection with classifying topoi.

Until further notice, let \mathcal{H} be a universal Horn theory, with at least one constant symbol, and let $\mathcal{S}^{\mathcal{C}^{\text{op}}}$ be its classifying topos as constructed in Section 1. Pseudo-forcing is a relation between consistent (finite) conjunctions $\phi(x)$ of atomic formulas, usually called conditions in this context, and formulas $\alpha(x)$ built from atomic formulas by means of conjunction, disjunction, negation, and existential quantification. Before giving the definition, we point out that the crucial difference between pseudo-forcing and (honest) forcing is that negations are not allowed in our (pseudo) conditions. Another apparent difference, our use of free variables in ϕ and α where Robinson used new constants, is only a matter of notation and has no effect on the theory. Formulas $\alpha(x)$ of the sort described above will, in accord with the terminology of [9], be called Robinson formulas. Although every formula is

equivalent, in classical logic, to a Robinson formula, the restriction to such formulas is non-vacuous because pseudo-forcing (like forcing) does not respect classical logic.

The definition of the pseudo-forcing relation, $\phi(x) \Vdash \alpha(x)$, is by induction on $\alpha(x)$.

If $\alpha(x)$ is atomic, then $\phi(x) \Vdash \alpha(x)$ if and only if $\mathcal{M} \vDash \forall x(\phi(x) \rightarrow \alpha(x))$.

If α is $\beta \wedge \gamma$ (resp. $\beta \vee \gamma$) then $\phi \Vdash \alpha$ if and only if $\phi \Vdash \beta$ and (resp. or) $\phi \Vdash \gamma$.

If α is $\exists y \beta(y)$ then $\phi \Vdash \alpha$ if and only if, for some term t , $\phi \Vdash \beta(t)$.

If α is $\neg \beta$, then $\phi \Vdash \alpha$ if and only if there is no ψ such that $(\phi \wedge \psi)$ is consistent and $\phi \wedge \psi \Vdash \beta$.

The list of free variables x in $\phi(x)$ and $\alpha(x)$, which we have omitted for the sake of brevity in most of the clauses, is never of any importance, as long as it contains all the free variables of ϕ and α . Even in the atomic clause, additional (dummy) variables added to the list x would make no difference, since the assumption that \mathcal{M} has a constant symbol precludes any difficulties arising from empty structures. We emphasize that the t in the existential quantifier clause and the ψ in the negation clause may well contain free variables other than those in ϕ and α . In this respect, our definition agrees with the usual definition of forcing.

The atomic clause in our definition of \Vdash differs from the usual one in not requiring $\alpha(x)$ to occur explicitly as a conjunct in $\phi(x)$. This difference will be useful when we relate pseudo-forcing to classifying topoi. It could be incorporated into the usual definition of forcing without any substantial effect on the theory; generic models, weak forcing, and forcing companions are all unaffected. It would have the rather pleasant consequence that \mathcal{M} -provably equivalent conditions force the same statements.

We leave it to the reader to check that the usual properties of forcing hold for pseudo-forcing. In particular, if $\phi \Vdash \alpha$ and if $\phi \wedge \psi$ is a condition, then $\phi \wedge \psi \Vdash \alpha$. Also, if $\phi \Vdash \alpha$ and if α' results from α by substitution of a closed term for a variable not free in ϕ , then $\phi \Vdash \alpha'$.

The following theorem is based on the observation that the definition of pseudo-forcing closely resembles the Kripke–Joyal sheaf semantics for presheaf topoi. The resemblance is not perfect, however, so some work is needed in the proof. In particular, an example to be given after the proof shows that the assumption that \mathcal{M} has a constant symbol is necessary. In the theorem, \mathcal{G} is the universal model of \mathcal{M} in $\mathcal{C}^{(\mathcal{M})}$, G is its underlying presheaf, and \models is the sheaf satisfaction relation [9, 8 §5.4] (whose definition will be recalled in the course of the proof).

Theorem 2. *Let $c = \{x \mid \phi(x)\}$ be an object of \mathcal{C} , and let a be elements of $G(c) = \langle x \mid \phi(x) \rangle$ given by (equivalence classes of) terms $t(x)$. Then, for any Robinson formula $\alpha(z)$, $\mathcal{G} \models_c \alpha(a)$ if and only if $\phi(x) \Vdash \alpha(t(x))$.*

Proof. We proceed by induction on $\alpha(z)$. If $\alpha(z)$ is atomic, then $\mathcal{G} \models_c \alpha(a)$ means, by definition, that $\alpha(a)$ holds in $\mathcal{G}(c) = \langle x \mid \phi(x) \rangle$. This means, also by definition,

that

$$\mathcal{H} \vdash \forall x(\phi(x) \rightarrow \alpha(t(x))).$$

And this is just the definition of $\phi(x) \dashv\vdash \alpha(t(x))$.

The cases of conjunction and disjunction are trivial, since $\mathcal{G} \vDash_c \beta \wedge \gamma$ (resp. $\beta \vee \gamma$) if and only if $\mathcal{G} \vDash_c \beta$ and (resp. or) $\mathcal{G} \vDash_c \gamma$. (For disjunction, it is important that we are working in a presheaf topos so it is not necessary to pass to a covering of c .)

If $\alpha(z)$ is $\exists y \beta(z, y)$, then $\mathcal{G} \vDash_c \alpha(a)$ if and only if there exists $b \in G(c) = \langle x \mid \phi(x) \rangle$ such that $\mathcal{G} \vDash_c \beta(a, b)$. Such a b is given by a term $u(x)$, so, applying the induction hypothesis to β , we find that $\mathcal{G} \vDash_c \alpha(a)$ if and only if there is a term $u(x)$ with only x free, such that $\phi(x) \dashv\vdash \beta(t(x), u(x))$. The restriction that all the free variables of u be among x can be removed, since any other free variables could be replaced with a constant symbol, which we have assumed to exist. But without the restriction on the variables of u , the requirement for $\mathcal{G} \vDash_c \alpha(a)$ reduces to $\phi(x) \dashv\vdash \alpha(t(x))$.

Finally, suppose $\alpha(z)$ is $\neg \beta(z)$. Then $\mathcal{G} \vDash_c \alpha(a)$ if and only if there is no morphism $\lambda : c' \rightarrow c$ in \mathcal{C} such that $\mathcal{G} \vDash_{c'} \beta(\lambda^*(a))$. Here λ^* means $G(\lambda)$, so if $\lambda = [x = u(y)]$, then λ^* sends $a = [t(x)]$ to $[t(u(y))]$. By Lemma 2, there is no loss of generality in assuming that λ is a simple morphism from, say $\{x, y \mid \psi(x, y)\}$ to $\{x \mid \phi(x)\}$. Since $\psi(x, y)$ provably implies $\phi(x)$, we may, by composing with a (simple) isomorphism, replace the domain of λ with $\{x, y \mid \phi(x) \wedge \psi(x, y)\}$. Thus, the statement $\mathcal{G} \vDash_c \alpha(a)$ is equivalent to: there is no $\psi(x, y)$ such that $\mathcal{G} \vDash_{c'} \beta(a)$, where $c' = \{x, y \mid \phi(x) \wedge \psi(x, y)\}$, and where $\lambda^*(a)$ has been simplified to a since λ is simple. Now the induction hypothesis applied to β gives the further equivalent form: there is no $\psi(x, y)$ such that $\phi(x) \wedge \psi(x, y) \dashv\vdash \beta(t(x))$, i.e., $\phi(x) \dashv\vdash \alpha(t(x))$. This completes the proof of Theorem 2. \square

We give some examples to clarify Theorem 2 and the concept of pseudo-forcing.

Example 1. Let \mathcal{H} be pure equality theory; it has no non-logical symbols and no non-logical axioms. Although \mathcal{H} is a universal Horn theory, it fails to satisfy our requirement that there be at least one constant symbol. We shall see that this failure results in Theorem 2 being false for \mathcal{H} . To see this, simply observe that the empty conjunction, *true*, pseudo-forces $\exists x \text{ true}$, since arbitrary terms are allowed in the \exists -clause of the definition of pseudo-forcing. However, the sentence $\exists x \text{ true}$ is not satisfied (in the sense of sheaf semantics) by \mathcal{G} at stage $\{ \mid \text{true} \} = 1$ in the classifying topos of \mathcal{H} (the object classifier), for, if it were, then, being a geometric formula, it would also be satisfied by all objects in all topoi, whereas in fact it is not satisfied by the empty set.

The temptation to remedy this defect by requiring, in the definition of $\phi \dashv\vdash \exists x \beta(x)$, that t have no free variables other than those of ϕ and $\exists x \beta(x)$ must be resisted, since it would prevent us from connecting pseudo-forcing with Robinson forcing where no such requirement is imposed.

Example 2. Let us modify the \mathcal{H} of the preceding example by adding one constant symbol $*$ to the language, so that Theorem 2 becomes applicable. Thus, \mathcal{H} is theory of pointed sets. We use Theorem 2 to check that the universal pointed set \mathcal{U} is remarkably small:

$$\mathcal{U} \models_c \neg \exists x \neg (x = *).$$

It suffices to check that no condition ϕ pseudo-forces $\exists x \neg (x = *)$. This means that, given ϕ and an arbitrary term t , we must find an extension $\phi \wedge \psi$ of ϕ that forces $t = *$. But $\phi \wedge (t = *)$ is such a condition; it is consistent because any conjunction of atomic formulas is consistent with \mathcal{H} .

The property of \mathcal{H} just cited, that every conjunction of atomic formulas is a condition, holds for any universal Horn theory in whose axioms, $\forall x(\phi \rightarrow \psi)$, ψ is always atomic (i.e. never *false*.) In particular, it holds for any equational theory. Thus, the universal algebra of any variety satisfies $\neg \exists x \neg (x = *)$ if $*$ is a nullary operation of the variety (and it satisfies $\neg \exists x \exists y \neg (x = y)$ in any case).

Rather than continuing with the theory of pseudo-forcing, by defining pseudo-generic models and the classifying topos for such models, we turn to the connection between pseudo-forcing and finite Robinson forcing. A comparison of our definition of \Vdash with Robinson's reveals the following differences:

- (a) Robinson uses new constants where we use free variables;
- (b) Robinson's conditions are sets of formulas, whereas ours are the conjunctions of those sets;
- (c) Robinson's definition and ours have different clauses for the atomic case;
- (d) Robinson's definition permits negated atomic formulas to occur in conditions, and ours does not.

Differences (a) and (b) are purely notational. Difference (c) is non-trivial, but, as we remarked above, if Robinson's theory were changed to agree with ours in this respect, nothing would be lost. The essential difference between pseudo-forcing and forcing is (d). For example, in contrast to Example 2, for the theory considered there (or any non-trivial variety with a nullary operation $*$), no condition forces $\neg \exists x \neg (x = *)$, because every condition can be consistently extended by adding $\neg (z = *)$ where z is a new variable. Despite this crucial difference between forcing and pseudo-forcing, we shall show that the former can be viewed as a special case of the latter by considering suitable theories.

Since the forcing relation (in Robinson's sense) for an arbitrary first-order theory \mathcal{T} depends only on the universal part of \mathcal{T} , we consider only universal theories in the following discussion of forcing.

Henceforth, \mathcal{T} is a consistent universal theory in a first-order language L that has at least one constant symbol.

Let L_d be the language obtained from L by adding, for each relation symbol R of L (including the equality symbol), a new relation symbol \mathcal{R} with the same number of arguments as R . Let \mathcal{T}'_d be the theory, in the language L_d , obtained by adding

to \mathcal{T} the axioms

$$\forall x(R(x) \Leftrightarrow \neg R(x)) \quad (6)$$

for all relation symbols R of L . Then \mathcal{T}'_d is a definitional extension of \mathcal{T} , so the two theories have essentially the same models; every model of \mathcal{T} is the L -reduct of a unique model of \mathcal{T}'_d . The category of models of \mathcal{T}'_d is, however, quite different from the category of models of \mathcal{T} , since a morphism of \mathcal{T}'_d -models must preserve not only the relations R of L but also their negations \bar{R} .

The theory \mathcal{T}'_d is equivalent, in classical logic, to a geometric theory $\tilde{\mathcal{T}}_d$ to be described below. Whenever we deal with interpretations in non-Boolean topoi, it will be $\tilde{\mathcal{T}}_d$, not \mathcal{T}'_d , that is relevant. The axioms of $\tilde{\mathcal{T}}_d$ are obtained as follows. First, there are the sentences

$$\forall x(\text{true} \rightarrow R(x) \vee \bar{R}(x)), \quad (7)$$

$$\forall x(R(x) \wedge \bar{R}(x) \rightarrow \text{false}), \quad (8)$$

for all relation symbols R of L . These axioms are jointly equivalent to (6) in classical logic. (Only the implication from (6) to (7) requires classical logic.) Second, each axiom of \mathcal{T} is rewritten in prenex form with its matrix in conjunctive normal form, and the universal quantifiers are distributed over the conjunction to yield an equivalent axiom that is a conjunction of sentences of the form $\forall x \bigvee_{i=1}^n \phi_i(x)$ where each ϕ_i is an atomic or negated atomic formula. Then each of these conjuncts is rewritten as

$$\forall x \left(\bigwedge_{i=1}^n -\phi_i(x) \rightarrow \text{false} \right) \quad (9)$$

where $-\phi_i(x)$ is $\bar{R}(x)$ (resp. $R(x)$) if $\phi_i(x)$ is $R(x)$ (resp. $\neg R(x)$). We take all the resulting sentences (9) as axioms of $\tilde{\mathcal{T}}_d$. In the presence of (7) and (8), these axioms (9) are clearly equivalent to the original axioms of \mathcal{T} from which they were derived.

The notation $\tilde{\mathcal{T}}_d$ was chosen to indicate that (in arbitrary topoi, not necessarily Boolean) the models of $\tilde{\mathcal{T}}_d$ are the decidable models of \mathcal{T} , i.e. those models \mathcal{M} such that the interpretation of each n -ary relation symbol is a complemented subobject of M^n .

Let $\tilde{\mathcal{T}}_{d\vee H}$ be the universal Horn part of $\tilde{\mathcal{T}}_d$; that is, its axioms are all the universal Horn sentences provable in $\tilde{\mathcal{T}}_d$. Thus, (8) and (9) are among the axioms of $\tilde{\mathcal{T}}_{d\vee H}$, but (7) is not. Other axioms of $\tilde{\mathcal{T}}_{d\vee H}$ include sentences that are like (9) except that one of the $\phi_i(x)$ has been left on the right of the implication sign instead of being transposed to (a $-\phi_i(x)$ on) the left. These sentences, and similar ones obtained from the universal theorems of \mathcal{T} (instead of only its axioms) constitute an axiomatization of $\tilde{\mathcal{T}}_{d\vee H}$, but we shall have no use for this fact, so we omit its proof.

Example 3. Let \mathcal{T} be the theory of pointed sets, as in Example 2. Then $\tilde{\mathcal{T}}_d$ is the

theory, with one binary predicate symbol \neq , axiomatized by

$$\forall x \forall y (\text{true} \rightarrow x = y \vee x \neq y), \quad (10)$$

$$\forall x (x \neq x \rightarrow \text{false}). \quad (11)$$

The second of these, but not the first, is an axiom of $\mathcal{T}_{d\forall H}$. It is not hard to check that $\mathcal{T}_{d\forall H}$ is axiomatized by (11) and

$$\forall x \forall y (x \neq y \rightarrow y \neq x).$$

Thus, $\mathcal{T}_{d\forall H}$ is simply the theory of (undirected, simple) graphs if we read \neq as ‘‘adjacent to’’. From this point of view, the models of \mathcal{T}_d are just the graphs satisfying (10), the complete graphs.

Returning to a general theory \mathcal{T} as above, observe that, since $\mathcal{T}_{d\forall H}$ is, by definition, a universal Horn theory, our earlier work is applicable to it. In particular, we have a syntactic category \mathcal{C} , dual to the category of finitely presented models of $\mathcal{T}_{d\forall H}$, such that $\mathcal{C}^{(\text{op})}$ is a classifying topos for $\mathcal{T}_{d\forall H}$ with universal model \mathcal{G} . We also have pseudo-forcing related to truth in \mathcal{G} by Theorem 2. But pseudo-forcing for $\mathcal{T}_{d\forall H}$ is essentially the same as Robinson forcing for \mathcal{T} . More precisely, if we modify Robinson’s definition of forcing so as to eliminate the differences (a), (b), (c) listed above, and if we rewrite forcing conditions by putting \mathcal{R} in place of $\neg R$ whenever a negated atomic formula occurs, then the resulting definition of forcing for \mathcal{T} agrees with our definition of pseudo-forcing for $\mathcal{T}_{d\forall H}$. The proof of this is a straightforward induction, since the definitions are virtually identical. The only non-trivial points to notice are that the formula $\forall x(\phi(x) \rightarrow \alpha(x))$ occurring in the atomic clause is universal Horn and that $\mathcal{R}(x)$ pseudo-forces $\neg R(x)$ because of (8). The previously crucial difference (d) between forcing and pseudo-forcing has been eliminated by the introduction of the new relation symbols \mathcal{R} , despite the fact that \mathcal{R} is equivalent to $\neg R$ only in \mathcal{T}_d , not in $\mathcal{T}_{d\forall H}$.

Ignoring the inessential changes ((a), (b), (c) and the use of \mathcal{R} for $\neg R$) in Robinson’s definition, we may summarize the preceding discussion combined with Theorem 2, as:

Finite forcing for \mathcal{T} is truth in the universal model of $\mathcal{T}_{d\forall H}$.

3. Classifying topoi for generic models

We continue the convention that \mathcal{T} is a consistent universal theory with at least one constant, and we let $L, L_d, \mathcal{T}_d, \mathcal{T}_{d\forall H}, \mathcal{C}$, and \mathcal{G} be as before. The symbol \Vdash will refer to forcing for \mathcal{T} , i.e., pseudo-forcing for $\mathcal{T}_{d\forall H}$.

A model \mathcal{M} of \mathcal{T} (in the topos \mathcal{C} of sets) is *generic* [1, 7] if, for every Robinson formula $\alpha(x)$ of L and every a in M^n , where n is the length of x , there is a condition $\phi(x, y)$ and there are elements b in M such that

$$\Vdash \phi(a, b) \quad \text{and} \quad \phi(x, y) \Vdash \alpha(x) \vee \neg \alpha(x).$$

The conditions $\phi(x, y)$ that force $\alpha(x) \vee \neg \alpha(x)$ are, according to Theorem 2, just the ones such that the simple morphism

$$\{x, y \mid \phi(x, y)\} \rightarrow \{x \mid \text{true}\} \quad (12)$$

is in the sieve (on $\{x \mid \text{true}\}$ in \mathcal{C}) that is the truth value $\|\alpha(x) \vee \neg \alpha(x)\|$ in \mathcal{G} of $\alpha(x) \vee \neg \alpha(x)$. Since \mathcal{S} is Boolean, a model \mathcal{M} of \mathcal{T} may be viewed as a model of \mathcal{T}_d , hence also of $\mathcal{T}_{d\vee H}$. The associated flat functor $F: \mathcal{C} \rightarrow \mathcal{S}$ (in the proof of Theorem 1) sends the simple morphism (12) to the projection

$$\{a, b \in M^{n+k} \mid \mathcal{M} \models \phi(a, b)\} \rightarrow M^n,$$

and the requirement for genericity is that these projections be jointly epimorphic. Lemma 2 allows us to ignore any non-simple morphisms in the sieve $\|\alpha(x) \vee \neg \alpha(x)\|$, so we conclude that a generic model of \mathcal{T} is one whose associated flat functor sends $\|\alpha(x) \vee \neg \alpha(x)\|$ to an epimorphic family, for all Robinson formulas $\alpha(x)$. This observation suggests defining a Grothendieck topology on \mathcal{C} such that this criterion for genericity becomes simply the continuity of F .

Definition. The *Robinson topology* J_R is the smallest Grothendieck topology on \mathcal{C} such that, for every object $\{x \mid \phi(x)\}$ of \mathcal{C} and every Robinson formula $\alpha(x)$ with free variables among x , the sieve $\|\alpha(x) \vee \neg \alpha(x)\|$ on $\{x \mid \phi(x)\}$ belongs to J_R .

Several equivalent descriptions of J_R will be useful. First, notice that the definition would be unchanged if we replaced all three occurrences of ' $\alpha(x)$ ' with ' $\alpha(t(x))$ ' where t ranges over all terms. This is simply because $\alpha(t(x))$ is another Robinson formula $\beta(x)$. Since every element of $\langle x \mid \phi(x) \rangle = G(\{x \mid \phi(x)\})$ is of the form $t(x)$ for some term t , we see that J_R can also be described as the smallest topology which, for each object c of \mathcal{C} , each Robinson formula $\alpha(x)$, and each $a \in G(c)$, contains the sieve $\|\alpha(a) \vee \neg \alpha(a)\|$ on c . In this form, the definition makes it clear that the given generating family for J_R is closed under pullbacks (of sieves). It also shows that the corresponding Lawvere–Tierney topology j_R in $\mathcal{S}^{(\mathcal{C}^{\text{op}})}$ is the smallest one that makes the interpretation in \mathcal{G} of every Robinson formula in n variables a complemented subobject of G^n .

Another reformulation of the definition of J_R is obtained by replacing both occurrences of ' $\phi(x)$ ' with ' true '. The reason this change does not affect the topology is that each sieve that is required to be in J_R by the original definition is the pullback, along a simple morphism $\{x \mid \phi(x)\} \rightarrow \{x \mid \text{true}\}$, of one that is required to be in J_R by the new definition. The new definition gives a slightly simpler generating family for J_R , and it connects more directly with our previous discussion of genericity, but the new generating family turns out to be less useful because it is not closed under pullbacks.

To study continuous flat functors on \mathcal{C} , we shall need the following lemma from the topos-theoretic folklore, which appears not to be published. To simplify its statement, we list for reference the axioms for a Grothendieck topology J .

(a) For every object c , the maximal sieve, consisting of all morphisms into c , is in J .

(b) If J contains a sieve R on c and if $f: c' \rightarrow c$ is any morphism, then J also contains the pullback f^*R of R along f , which consists of all morphisms into c' whose composite with f belongs to R .

(c) R and S are sieves on c , if $R \in J$, and if $f^*S \in J$ for all $f \in R$, then $S \in J$.

Lemma 3. *Let K be a family of sieves in a category \mathcal{C} , and assume that K is closed under pullbacks. Then the Grothendieck topology J generated by K is the smallest family of sieves in \mathcal{C} that includes K and satisfies (a) and (c) above. A flat functor from \mathcal{C} to a topos \mathcal{E} is continuous for J if and only if it sends every sieve in K to an epimorphic family in \mathcal{E} .*

Proof. To establish the first assertion, let J' be the smallest family that includes K and satisfies (a) and (c). Since J is the smallest family with these properties plus (b), all we need to show is that (b) holds for J' . We let J'' consist of those sieves all of whose pullbacks are in J' , and we prove that $J' \subseteq J''$ by showing that J'' has all the properties in the definition of J' . J'' includes K because J' does and K is closed under pullback. J'' satisfies (a) because J' does and pullbacks of maximal sieves are maximal. To show that J'' satisfies (c), let R and S be sieves on an object c , and assume that $R \in J''$ and $f^*S \in J''$ for all $f \in R$. We must show that $S \in J''$, so we consider an arbitrary $g: c' \rightarrow c$ and show that $g^*S \in J'$. Both g^*R and g^*S are sieves on c' and the former is in J' because $R \in J''$. So to show $g^*S \in J'$ it suffices, since J' satisfies (c), to show $h^*g^*S \in J'$ for every $h \in g^*R$. But $h \in g^*R$ means that $gh \in R$. By our assumption about S , $(gh)^*S = h^*g^*S$ is in J'' and therefore in J' , as required.

The second assertion will follow from the first if we show that, for any flat functor $F: \mathcal{C} \rightarrow \mathcal{E}$, the family J_F of sieves sent to epimorphic families satisfies (a) and (c); then if $K \subseteq J_F$ we can infer $J \subseteq J_F$ as required. Part (a) is obvious since F preserves identity morphisms. For part (c), assume R, S are sieves on c , $R \in J_F$, $f^*S \in J_F$ for all $f \in R$. To show that the family $F(S)$ of morphisms into $F(c)$ is epimorphic, suppose p and q are two morphisms $F(c) \rightarrow X$ such that $p \circ F(h) = q \circ F(h)$ for every $h \in S$. We must show that $p = q$, and for this it suffices to show $p \circ F(f) = q \circ F(f)$ for all $f \in R$, since $F(R)$ is epimorphic. For a fixed $f \in R$, to show $p \circ F(f) = q \circ F(f)$, it suffices to show $p \circ F(f) \circ F(g) = q \circ F(f) \circ F(g)$ for all $g \in f^*S$, since $F(f^*S)$ is epimorphic. But $g \in f^*S$ means $f \circ g \in S$, so we have $p \circ F(f \circ g) = q \circ F(f \circ g)$ by hypothesis, and we are done because F is a functor. \square

We apply this lemma to describe the continuous functors from the site (\mathcal{C}, J_R) to a topos \mathcal{E} . A flat functor $F: \mathcal{C} \rightarrow \mathcal{E}$ is given, according to Theorem 1, by a model \mathcal{M} of \mathcal{L}_{dVH} in \mathcal{E} . For F to be continuous, it is, according to the lemma, necessary and sufficient that it send each sieve $\|\alpha(x) \vee \neg \alpha(x)\|$ on any object $\{x \mid \phi(x)\}$ to an epimorphic family. By Lemma 2, we may restrict our attention to the simple

morphisms in this sieve. These are, up to isomorphism,

$$\{x, y \mid \phi(x) \wedge \psi(x, y)\} \rightarrow \{x \mid \phi(x)\}$$

where $\phi(x) \wedge \psi(x, y) \Vdash \alpha(x) \vee \neg \alpha(x)$. F sends these morphisms to projections $\{a, b \in M^{n+k} \mid .\# \models \phi(a) \wedge \psi(a, b)\} \rightarrow \{a \in M^n \mid .\# \models \phi(a)\}$. Continuity of F requires that, for each ϕ and α , the projections so obtained from various ψ 's constitute an epimorphic family. In other words, the following statement must hold in the internal logic of \mathcal{E} :

$$(\forall a \in M^n) \left[(.\# \models \phi(a)) \rightarrow \bigvee_{\psi} (\exists b \in M^k). \# \models \psi(a, b) \right]$$

where the disjunction is over ψ 's such that $\phi(x) \wedge \psi(x, y) \Vdash \alpha(x) \vee \neg \alpha(x)$. It is easy to see that this statement will hold for arbitrary ϕ if it holds when ϕ is *true*, so we obtain the following simplification.

Theorem 3. *A model $.\#$ of \mathcal{T}_{dVH} in a topos \mathcal{E} corresponds to a continuous functor from (\mathcal{C}, J_R) to \mathcal{E} if and only if, in the internal logic of \mathcal{E} , it is true, for each Robinson formula $\alpha(x)$, that*

$$\forall a \in M^n \exists b \in M^k \bigvee_{\psi} .\# \models \psi(a, b),$$

where the disjunction is over all conditions $\psi(x, y)$ forcing $\alpha(x) \vee \neg \alpha(x)$.

Definition. A model $.\#$ with the property in Theorem 3 is called a *generic model* cf. \mathcal{T}_d .

This terminology presupposes that such an $.\#$ is in fact a model of \mathcal{T}_d ; we shall confirm this supposition below.

Corollary. *The topos of sheaves on (\mathcal{C}, J_R) is the classifying topos for generic models of \mathcal{T}_d .*

Thus, the universal generic model of \mathcal{T}_d is obtained from the universal model of \mathcal{T}_{dVH} by forcing (with J_R) every Robinson formula to become decidable.

To see that generic models of \mathcal{T}_d are models of \mathcal{T}_d , we need the following analog of the well-known “forcing equals truth” lemma in finite forcing theory.

Lemma 4. *Let $.\#$ be a generic model of \mathcal{T}_d in \mathcal{E} . Then, for any Robinson formula $\alpha(x)$, the following holds in the internal logic of \mathcal{E} :*

$$(\forall a \in M^n) \left[.\# \models \alpha(a) \leftrightarrow \bigvee_{\phi} (\exists b \in M^k). \# \models \phi(a, b) \right]$$

where the disjunction is over conditions $\phi(x, y)$ that force $\alpha(x)$.

Proof. We proceed by induction on α . Suppose first that α is atomic. Working in the internal logic of \mathcal{E} , we have that, if $\mathcal{M} \models \alpha(a)$ then the disjunction in the desired formula holds because there is a disjunct in which ϕ is α (and $k=0$). Conversely, if the disjunct corresponding to $\phi(x, y)$ holds, then, as $\phi(x, y)$ forces $\alpha(x)$, we have that $\forall x, y(\phi(x, y) \rightarrow \alpha(x))$ is provable in $\mathcal{T}_{d \vee H}$, hence true in \mathcal{M} , so from $\mathcal{M} \models \phi(a, b)$ we can infer $\mathcal{M} \models \alpha(a)$.

Next, suppose $\alpha(x)$ is $\neg\beta(x)$. Again, we work in the internal logic of \mathcal{E} . If $\mathcal{M} \models \alpha(x)$, then, by induction hypothesis, we do not have $\mathcal{M} \models \phi(a, b)$ for any $\phi(x, y)$ forcing $\beta(x)$. But, by genericity, we do have $\mathcal{M} \models \phi(a, b)$ for some $\phi(x, y)$ forcing $\beta(x) \vee \neg\beta(x)$, so this ϕ must force $\neg\beta$, which is α . Conversely, suppose $\mathcal{M} \models \phi(a, b)$ and $\phi(x, y)$ forces $\alpha(x)$. Then \mathcal{M} cannot satisfy $\beta(a)$ because to do so it would have to satisfy $\psi(a, b, c)$ for some $\psi(x, y, z)$ forcing $\beta(x)$, and then $\phi \wedge \psi$ would be consistent (since satisfied in \mathcal{M}) and would be an extension of ϕ forcing β , which is absurd as ϕ forces $\neg\beta$.

If $\alpha(x)$ is $\beta(x) \wedge \gamma(x)$, then, by induction hypothesis, we have in the internal logic that $\mathcal{M} \models \alpha(a)$ if and only if $\mathcal{M} \models \phi(a, b) \wedge \psi(a, b)$ for some ϕ forcing β and some ψ forcing γ . (We assumed that the b is the same in both parts, since we may add dummy variables to ϕ and ψ .) Then $\phi \wedge \psi$ is consistent (since satisfied in \mathcal{M}), so it is a condition forcing $\beta \wedge \gamma$; the converse direction is even easier.

The case that α is a disjunction is trivial.

Finally, let $\alpha(x)$ be $\exists z \beta(x, z)$. Working again in the internal logic, we have that $\mathcal{M} \models \alpha(x)$ if and only if, for some b and c , $\mathcal{M} \models \phi(a, b, c)$ where $\phi(x, y, z) \Vdash \beta(x, z)$. This certainly implies $\phi(x, y, z) \Vdash \alpha(x)$. Conversely, suppose $\phi(x, y) \Vdash \alpha(x)$ and $\mathcal{M} \models \phi(a, b)$. Then there is a term $t(x, y)$ such that $\phi(x, y) \Vdash \beta(x, t(x, y))$; we have arranged that t has no free variables other than x, y by replacing the others with a constant symbol. By induction hypothesis we have $\mathcal{M} \models \beta(a, t(a, b))$, so $\mathcal{M} \models \alpha(a)$, as required. \square

Using Lemma 4, we show that a generic model \mathcal{M} is a model of $\tilde{\mathcal{T}}_d$. It is, of course, a model of $\tilde{\mathcal{T}}_{d \vee H}$ by definition; the only axioms of $\tilde{\mathcal{T}}_d$ that are not in $\tilde{\mathcal{T}}_{d \vee H}$ are those of the form

$$\forall x(\text{true} \rightarrow R(x) \vee \bar{R}(x)).$$

To show these hold in \mathcal{M} , we work in the internal logic and find, for any $a \in M^n$, a condition $\phi(x, y)$, forcing $R(x) \vee \bar{R}(x)$, and satisfied in \mathcal{M} by a, b for some b . By genericity, we can find $\phi_1(x, y)$, satisfied in \mathcal{M} by a, b , and forcing $R(x) \vee \neg R(x)$, and we can find $\phi_2(x, y)$, also satisfied in \mathcal{M} by a, b (the same b , by adding dummy variables if necessary), and forcing $\bar{R}(x) \vee \neg \bar{R}(x)$. If ϕ_1 forces $R(x)$ or ϕ_2 forces $\bar{R}(x)$, then it serves as the desired ϕ . It remains to consider the case that ϕ_1 forces $\neg R(x)$ and ϕ_2 forces $\neg \bar{R}(x)$; we shall show that this case cannot arise. If it did, then, by definition of forcing of negations, neither $\phi_1 \wedge \bar{R}(x)$ nor $\phi_2 \wedge R(x)$ is a condition, for the former (resp. latter) would be an extension of ϕ_1 (resp. ϕ_2)

forcing $R(x)$ (resp. $\bar{R}(x)$), which is impossible. So both

$$\forall x, y(\phi_1(x, y) \wedge R(x) \rightarrow \text{false}) \quad (13)$$

and

$$\forall x, y(\phi_2(x, y) \wedge \bar{R}(x) \rightarrow \text{false}) \quad (14)$$

are theorems of $\mathcal{T}_{d\forall H}$, hence a fortiori of \mathcal{T}_d . But \mathcal{T}_d has the axiom $\forall x(\text{true} \rightarrow R(x) \vee \bar{R}(x))$ which together with (13) and (14) yields

$$\forall x, y(\phi_1(x, y) \wedge \phi_2(x, y) \rightarrow \text{false}).$$

This sentence, being a universal Horn theorem of \mathcal{T}_d , is an axiom of $\mathcal{T}_{d\forall H}$, hence is true in \mathcal{M} . But in \mathcal{M} , $\phi_1 \wedge \phi_2$ is satisfied by \mathbf{a}, \mathbf{b} . This contradiction completes the proof that generic models satisfy \mathcal{T}_d .

Remark. Except for the verification that generic models satisfy \mathcal{T}_d , the material in this section has really involved only $\mathcal{T}_{d\forall H}$, not \mathcal{T}_d or \mathcal{T} . We could, therefore, have started with an arbitrary universal Horn theory with a constant symbol, defined the Robinson topology J_R on its syntactic category \mathcal{C} , shown that the topos of sheaves on (\mathcal{C}, J_R) classifies pseudo-generic models (defined just like generic models but using pseudo-forcing), and proved a “pseudo-forcing equals truth” result like Lemma 4 for these pseudo-generic models. We chose not to present the results in this generality, although no additional work would have been required, for two reasons. First, we did not want to postpone for too long establishing contact with forcing, since it is our primary interest in this paper. Second, it turns out that, in some natural examples, pseudo-generic models are rather uninteresting. Specifically, it follows easily from Example 2 and the subsequent discussion that the only pseudo-generic algebras for any variety are the trivial algebras.

4. Other topologies

In this section, we discuss some naturally occurring topologies, other than J_R , on \mathcal{C} , the syntactic category of $\mathcal{T}_{d\forall H}$; we are interested particularly in their relationships to J_R and to each other and in the concepts classified by their sheaf topoi.

The theory, \mathcal{T}_d , being geometric, has a classifying topos; the methods of [14, 15] enable us to describe this topos as a sheaf subtopos of the classifying topos $\mathcal{Y}^{(\text{op})}$ of $\mathcal{T}_{d\forall H}$. Indeed, it is clear from the discussion of \mathcal{T}_d and $\mathcal{T}_{d\forall H}$ in Section 2 that the former is obtained from the latter by adding the axioms

$$\forall x(R(x) \vee \bar{R}(x)). \quad (7)$$

Therefore, by [14] or [15] the classifying topos for \mathcal{T}_d is the topos of sheaves on (\mathcal{C}, J_D) , where J_D is the smallest topology on \mathcal{C} such that, for each relation symbol R of L , the object $\{x \mid \text{true}\}$ is covered by the sieve $\|R(x) \vee \bar{R}(x)\|$. This sieve is generated by the simple morphisms to $\{x \mid \text{true}\}$ from $\{x \mid R(x)\}$ and $\{x \mid \bar{R}(x)\}$. Since

generic models satisfy $\tilde{\mathcal{T}}_d$, we have $J_D \subseteq J_R$, a fact which can also be verified directly (by a method closely resembling our proof that generic models satisfy \mathcal{T}_d) or deduced from more detailed information given below.

In ordinary (in \mathcal{L}) model theory, the existentially closed models of \mathcal{T} form an important class intermediate between the classes of all models of \mathcal{T} and of generic models of \mathcal{T} . We shall generalize the concept of existentially closed model to arbitrary topoi and construct a classifying topos for it as a sheaf subtopos of $\mathcal{L}^{(\text{op})}$. Joyal and Reyes [9] have defined the concept ‘existentially closed’ for geometric theories in the context of categorical logic and have related it to a certain Grothendieck topology on the pretopos that (from their point of view) is the geometric theory. We work only with universal theories, since in ordinary model theory the existentially closed models of an arbitrary theory are just those models of that theory which, considered as models of its universal part, are existentially closed. Also, we view theories and models in a less abstract way than Joyal and Reyes do. As a result, our definition of ‘existentially closed’ seems quite different from theirs. However, the similarity between their Grothendieck topology on the theory and the Grothendieck topology J_E on \mathcal{V} that we define below leads us to suspect that the two definitions are basically the same.

Definition. A model \mathcal{M} of $\tilde{\mathcal{T}}_{dVH}$ in a topos \mathcal{E} is *existentially closed* if and only if, for every conjunction $\psi(\mathbf{x}, \mathbf{y})$ of atomic formulas of L_d , the following statement holds in the internal logic of \mathcal{E} :

$$(\forall \mathbf{a} \in M^n) \left[(\exists \mathbf{b} \in M^k) (\mathcal{M} \models \psi(\mathbf{a}, \mathbf{b})) \vee \bigvee_{\lambda} (\exists \mathbf{c} \in M^l) (\mathcal{M} \models \lambda(\mathbf{a}, \mathbf{c})) \right],$$

where the disjunction is over all conjunctions $\lambda(\mathbf{x}, \mathbf{z})$ of atomic formulas of L_d , such that $\psi(\mathbf{x}, \mathbf{y}) \wedge \lambda(\mathbf{x}, \mathbf{z})$ (with \mathbf{y} and \mathbf{z} disjoint lists of variables) is inconsistent with $\tilde{\mathcal{T}}_{dVH}$.

We point out immediately that such a model \mathcal{M} necessarily satisfies $\tilde{\mathcal{T}}_d$. Indeed, applying the definition of existentially closed, first with $R(\mathbf{x})$ and then with $\bar{R}(\mathbf{x})$ as $\psi(\mathbf{x})$, we find (in the internal logic) that every $\mathbf{a} \in M^n$ either satisfies $R(\mathbf{a}) \vee \bar{R}(\mathbf{a})$ as desired or else satisfies $\lambda_1(\mathbf{a}, \mathbf{c}) \wedge \lambda_2(\mathbf{a}, \mathbf{c})$ for some \mathbf{c} , where $\lambda_1(\mathbf{x}, \mathbf{y})$ and $\lambda_2(\mathbf{x}, \mathbf{y})$ contradict $R(\mathbf{x})$ and $\bar{R}(\mathbf{x})$ in $\tilde{\mathcal{T}}_{dVH}$, hence in $\tilde{\mathcal{T}}_d$. (We can take the same \mathbf{c} for both λ 's by adding dummy variables if necessary.) It follows that $\lambda_1(\mathbf{x}, \mathbf{y}) \wedge \lambda_2(\mathbf{x}, \mathbf{y})$ contradicts $\tilde{\mathcal{T}}_d$, hence also contradicts $\tilde{\mathcal{T}}_{dVH}$ because $\forall \mathbf{x}, \mathbf{y} (\lambda_1 \wedge \lambda_2 \rightarrow \text{false})$ is a universal Horn sentence. Yet \mathbf{a}, \mathbf{c} satisfies $\lambda_1 \wedge \lambda_2$ in a model \mathcal{M} of $\tilde{\mathcal{T}}_{dVH}$. This contradiction shows that $R(\mathbf{a}) \vee \bar{R}(\mathbf{a})$ must hold, so $\mathcal{M} \models \tilde{\mathcal{T}}_d$.

The next lemma serves to justify our definition of existentially closed by showing that, in \mathcal{L} , it agrees with the usual definition. Recall that, in \mathcal{S} or in any Boolean topos, models of $\tilde{\mathcal{T}}_d$ are essentially the same as models of \mathcal{T} ; we shall therefore ignore the distinction between these theories.

Lemma 5. *A model \mathcal{M} of \mathcal{T} in \mathcal{S} is existentially closed if and only if, for all existential formulas $\eta(x)$, for all $a \in M^n$, and for all extensions $\mathcal{N} \supseteq \mathcal{M}$ that are models of \mathcal{T} , if $\mathcal{N} \models \eta(a)$ then $\mathcal{M} \models \eta(a)$.*

Proof. Suppose first that \mathcal{M} is existentially closed, and let $\eta(x)$, a , and \mathcal{N} be as in the statement of the lemma, with $\mathcal{N} \models \eta(a)$. By putting η into prenex form, putting its matrix into disjunctive form, and distributing existential quantifiers across disjunctions, we can assume $\eta(x)$ is a disjunction of formulas of the form $\exists y \psi(x, y)$ where ψ is a conjunction of atomic formulas of L_d . (We have gone from \mathcal{T} to \mathcal{T}_d and replaced any $\neg R$ in η with \bar{R} .) Let $\exists y \psi(x, y)$ be one of the disjuncts that is satisfied by a in η ; we shall show that it and hence also $\eta(x)$ are satisfied by a in \mathcal{M} . If it were not, then, as \mathcal{M} is existentially closed, we could find $c \in M^l$ satisfying $\lambda(a, c)$ in \mathcal{M} , where $\lambda(x, z)$ is a conjunction of atomic formulas of L_d and contradicts $\psi(x, y)$ in \mathcal{T}_d . But then $\lambda(a, c)$ still holds in \mathcal{N} because $\mathcal{N} \supseteq \mathcal{M}$, and this is absurd since $\exists y \psi(a, y)$ holds in \mathcal{N} as well.

Conversely, suppose \mathcal{M} satisfies the criterion in the lemma. Let $\psi(x, y)$ be a conjunction of atomic formulas of L_d , and let $a \in M^n$. We must show that either $\mathcal{M} \models \exists y \psi(a, y)$ or there is a conjunction $\lambda(x, z)$ of atomic formulas of L_d , such that $\psi(x, y) \wedge \lambda(x, z)$ is inconsistent with \mathcal{T}_d and such that $\mathcal{M} \models \exists z \lambda(a, z)$. Suppose, therefore, that no such λ exists. Then, in the language obtained by adding to L_d names for all elements of M and additional constants p , the set of sentences

$$\mathcal{T}_d \cup \text{Diagram of } \mathcal{M} \cup \{\psi(a, p)\}$$

is consistent, by a compactness argument. A model \mathcal{N} of it is, up to isomorphism, an extension of \mathcal{M} in which p witnesses that $\exists y \psi(a, y)$ holds. Therefore, since \mathcal{M} is assumed to satisfy the criterion in the lemma, $\exists y \psi(a, y)$ holds also in \mathcal{M} , as required. \square

We now begin the construction of a classifying topos for existentially closed models of \mathcal{T}_d .

Definition. J_E is the smallest Grothendieck topology on \mathcal{C} such that, for each morphism $f: c' \rightarrow c$ in \mathcal{C} , J_E contains the sieve $\pm f$ that consists of

- (a) all morphisms into c that factor through f , and
- (b) all morphisms g into c such that no morphism into c factors through both f and g .

Note that (a) describes the sieve generated by f , while (b) describes its negation, the largest sieve disjoint from the one generated by f . This explains the notation $\pm f$.

Note also that no special properties of \mathcal{C} are used in defining J_E . The same definition gives a topology on any category, and the preceding remark shows that this topology is always included in the double-negation topology.

Theorem 4. *The topos of sheaves on (\mathcal{C}, J_E) is the classifying topos for existentially closed models of $\bar{\mathcal{T}}_d$.*

Proof. If f and g are morphisms with a common codomain c , then either they are incompatible, in the sense that no morphism factors through both, or they have a pullback. In the first case, $g^*\pm f$ is the maximal sieve on the domain of g . In the second case, $g^*\pm f$ is $\pm h$, where h is the pullback of f along g . Therefore, the generating family given in the definition of J_E becomes closed under pullbacks if we adjoin to it the maximal sieves. Therefore, by Lemma 3, to check that a flat functor F of \mathcal{C} is continuous for J_E , we need only check that it sends each sieve $\pm f$ to an epimorphic family. By Lemma 2, we need only consider the sieves $\pm f$ where the morphism f is simple.

So let $F: \mathcal{C} \rightarrow \mathcal{E}$ be a flat functor, corresponding to a model \mathcal{M} of $\bar{\mathcal{T}}_{d \vee H}$ in a topos \mathcal{E} , and let

$$f: \{x, y \mid \psi(x, y)\} \rightarrow \{x \mid \phi(x)\}$$

be a simple morphism in \mathcal{C} . To say that $F(\pm f)$ is an epimorphic family means, in view of Lemma 2, that f and the simple morphisms

$$\{x, z \mid \lambda(x, z)\} \rightarrow \{x \mid \phi(x)\},$$

such that $\lambda(x, z)$ implies $\phi(x)$ but is inconsistent with $\psi(x, y)$ in $\bar{\mathcal{T}}_d$, are sent by F to an epimorphic family. This can be expressed in the internal language of \mathcal{E} by saying that, for every a satisfying ϕ in \mathcal{M} , either there is b such that $\mathcal{M} \models \psi(a, b)$ or there is c such that $\mathcal{M} \models \lambda(a, c)$ for some such λ . In the special case that $\phi(x)$ is *true*, this is precisely (the ψ instance of) the definition of being existentially closed. And the general case, with arbitrary $\phi(x)$, easily follows from the special case, for if $\lambda(x, z)$ works in the special case then $\lambda(x, z) \wedge \phi(x)$ works with $\phi(x)$. \square

Theorem 5. *The topologies J_R, J_D, J_E and the double-negation topology $J_{\neg\neg}$ on \mathcal{C} satisfy*

$$J_D \subseteq J_E \subseteq J_R \subseteq J_{\neg\neg}.$$

Proof. That $J_D \subseteq J_E$ is just a restatement of the already established fact that all existentially closed models satisfy $\bar{\mathcal{T}}_d$. We sketch a direct proof. To show that the sieve on $\{x \mid \text{true}\}$ generated by the simple morphisms f and g from $\{x \mid R(x)\}$ and $\{x \mid \bar{R}(x)\}$ is in J_E , it suffices to check that it is the intersection of the sieves $\pm f$ and $\pm g$, since Grothendieck topologies are closed under intersection [13, p. 15]. And checking this amounts to checking that no morphism can be incompatible with both f and g . By Lemma 2, we need only consider simple morphisms, and the necessary checking was done in the course of showing that existentially closed models satisfy $\bar{\mathcal{T}}_d$.

$J_f \subseteq J_R$ means that each sieve $\pm f$ is in J_R . It suffices, by Lemma 2, to prove this

for simple f , say from $\{x, y \mid \psi(x, y)\}$ to $\{x \mid \phi(x)\}$. To do this, we show that every morphism g in the sieve $\|\exists y \psi(x, y) \vee \neg \exists y \psi(x, y)\|$ on $\{x \mid \phi(x)\}$ also belongs to $\pm f$. Again, we may assume g is a simple morphism, say from $\{x, z \mid \theta(x, z)\}$ to $\{x \mid \phi(x)\}$. Suppose first that g belongs to $\|\exists y \psi(x, y)\|$, so by Theorem 2 and the definition of forcing,

$$\theta(x, z) \Vdash \psi(x, t(x, z))$$

for some terms t . Since ψ is a conjunction of atomic formulas, we have

$$\mathcal{F}_{d \vee H} \vdash \forall x, z (\theta(x, z) \rightarrow \psi(x, t(x, z))).$$

Therefore, $[x' = x, y = t(x, z)]$ is a morphism in \mathcal{C} from $\{x, z \mid \theta(x, z)\}$ to $\{x', y \mid \psi(x, y)\}$, and it clearly gives a factorization of g through f , so $g \in \pm f$. There remains the case that g belongs to $\|\neg \exists y \psi(x, y)\|$. In this case, by Theorem 2 and the definition of forcing, no extension of the condition $\theta(x, z)$ can force $\exists y \psi(x, y)$. In particular, $\theta(x, z) \wedge \psi(x, y)$, with the lists z and y disjoint, must be inconsistent with $\mathcal{F}_{d \vee H}$, for otherwise it would be a condition extending $\theta(x, z)$ and forcing $\exists y \psi(x, y)$. It follows immediately that g and f are incompatible, so again $g \in \pm f$.

Finally, that $J_R \subseteq J_{\neg\neg}$ is clear because the double negation of a generating sieve $\|\alpha(x) \vee \neg \alpha(x)\|$ of R is the truth value of the intuitionistically valid formula $\neg\neg(\alpha(x) \vee \neg \alpha(x))$. \square

We shall not discuss in detail the concept classified by the topos of sheaves on $(\mathcal{C}, J_{\neg\neg})$, since such a discussion would lead us away from model-theoretic forcing to set-theoretic forcing. An instructive example in this connection is the theory \mathcal{F} , with one constant symbol c and denumerably many unary predicate symbols P_n , axiomatized by $\forall x (x = c)$. Models of \mathcal{F}_d are essentially functions f from the set of natural numbers into $\{\text{true}, \text{false}\}$, giving the truth values of the sentences $P_n(c)$. Sheaves on $(\mathcal{C}, J_{\neg\neg})$ form the classifying topos for functions f that are Cohen-generic over \mathcal{F} . For further discussion of this and related examples, see [14].

We can, however, give a model-theoretic description of the points of the sheaf topos $\text{Sh}(\mathcal{C}, J_{\neg\neg})$. The description involves the following slight strengthening of the concept of atomic model.

Definition. A model \mathcal{M} of a theory \mathcal{F} is a *strongly atomic* model of \mathcal{F} if and only if, for each list of elements a in M , there is a formula $\phi(x)$, satisfied by a in \mathcal{M} , such that, for every formula $\psi(x)$ satisfied by a in \mathcal{M} ,

$$\mathcal{F} \vdash \forall x (\phi(x) \rightarrow \psi(x)).$$

This definition differs from Vaught's notion of atomic model only in that the last line has " $\mathcal{F} \vdash$ " instead of " $\mathcal{M} \models$ ". Thus, the two notions agree if \mathcal{F} is complete. Unlike atomicity, strong atomicity imposes a non-trivial requirement even when the list a is empty, for the definition then asserts that the complete theory of \mathcal{M} is

axiomatized by \mathcal{F} plus a single sentence. It is not hard to check that this requirement and atomicity together imply strong atomicity.

Recall that \mathcal{F}^f , the finite forcing companion of \mathcal{F} , consists of those sentences whose negation is not forced by any condition; see [1, 7]. Note that $(\mathcal{F}^f)_d = (\mathcal{F}_d)^f$, so we may unambiguously write \mathcal{F}_d^f .

Theorem 6. *A flat functor $F: \mathcal{C} \rightarrow \mathcal{D}$ is continuous for $J_{\neg\neg}$ if and only if the corresponding model of $\mathcal{F}_{d \vee H}$ is a strongly atomic model of \mathcal{F}_d^f .*

Proof. Let \mathcal{M} be the model corresponding to a $J_{\neg\neg}$ -continuous flat functor F on \mathcal{C} . Since $J_{\neg\neg} \supseteq J_R$, Theorem 3 tells us that \mathcal{M} is a generic model of \mathcal{F}_d , so it is a model of \mathcal{F}_d^f .

Consider an arbitrary list a of elements of \mathcal{M} , and let Φ be the set of all the conditions $\phi(x, y)$ such that $\mathcal{M} \models \neg \exists y \phi(a, y)$. The simple morphisms

$$\{x, y \mid \phi(x, y)\} \rightarrow \{x \mid \text{true}\} \quad (15)$$

for $\phi \in \Phi$ cannot form a $J_{\neg\neg}$ -covering of $\{x \mid \text{true}\}$ since none of the projections obtained by applying F to them has a in its range. So there exists a morphism, without loss of generality a simple one,

$$\{x, z \mid \psi(x, z)\} \rightarrow \{x \mid \text{true}\}$$

incompatible with (15) for all $\phi \in \Phi$. Thus, the existential formula $\exists z \psi(x, z)$ \mathcal{F}_d -provably implies all the formulas $\neg \exists y \phi(x, y)$ satisfied by a in \mathcal{M} , where ϕ ranges over conditions. It follows immediately that $\mathcal{M} \models \exists z \psi(a, z)$, for otherwise ψ would belong to Φ , hence imply its own negation, which contradicts the fact that ψ is consistent. To complete the proof that \mathcal{M} is a strongly atomic model of \mathcal{F}_d^f , we show that every formula $\gamma(x)$ satisfied in \mathcal{M} by a is deducible in \mathcal{F}_d^f from $\exists z \psi(x, z)$. Suppose, therefore, that $\gamma(x)$ were a counterexample. Then $\exists z \psi(x, z) \wedge \neg \gamma(x)$ is consistent with \mathcal{F}_d^f and is therefore forced by some condition $\phi(x, y)$. Then $\exists y \phi(x, y)$ is consistent with $\exists z \psi(x, z)$ in \mathcal{F}_d , so $\phi \notin \Phi$. This means that we can find b in \mathcal{M} such that $\mathcal{M} \models \phi(a, b)$. By the choice of ϕ and the genericity of \mathcal{M} , we have $\mathcal{M} \models \neg \gamma(a)$, a contradiction.

Conversely, suppose \mathcal{M} is a strongly atomic model of \mathcal{F}_d^f . To show that the associated flat functor $F: \mathcal{C} \rightarrow \mathcal{D}$ is $J_{\neg\neg}$ -continuous, consider an arbitrary double-negation dense sieve R on an object $\{x \mid \phi(x)\}$. To show that the family $F(R)$ is epimorphic, let a be an arbitrary element of $F\{x \mid \phi(x)\}$, so a satisfies $\phi(x)$ in \mathcal{M} . Since \mathcal{M} is strongly atomic, find a formula $\psi(x)$, also satisfied by a in \mathcal{M} , that \mathcal{F}_d^f -provably implies every formula satisfied by a in \mathcal{M} . Then $\psi(x)$ is an atom in the Lindenbaum algebra of \mathcal{F}_d^f -equivalence classes of formulas with only x free. But, since \mathcal{F}_d^f is a forcing companion, every consistent formula is above a consistent existential formula in this Lindenbaum algebra [7]; therefore the atom $\psi(x)$ can be taken to be an existential formula. As in several previous arguments, we can write $\psi(x)$ as a disjunction of formulas $\exists y \theta(x, y)$ where each θ is a conjunction of atomic

formulas of L_d . Since ψ is an atom, it is equivalent to one of these disjuncts $\exists y \theta(x, y)$. Since ψ \mathcal{T}_d^f -provably implies ϕ , and since \mathcal{T}_d^f and \mathcal{T}_d prove the same universal sentences, the universal Horn sentence $\forall x, y(\theta(x, y) \rightarrow \phi(x))$ is provable in $\mathcal{T}_{d\forall H}$, so there is a simple morphism

$$\{x, y \mid \theta(x, y)\} \rightarrow \{x \mid \phi(x)\}$$

in \mathcal{C} . Some composite morphism of the form

$$\{x, y, z \mid \lambda(x, y, z)\} \xrightarrow{f} \{x, y \mid \theta(x, y)\} \rightarrow \{x \mid \phi(x)\} \quad (16)$$

is in the double-negation dense sieve R ; we are using Lemma 2 to justify assuming that f is simple. The fact that f is a morphism means that the sentence

$$\forall x(\exists y, z \lambda(x, y, z) \rightarrow \exists y \theta(x, y)) \quad (17)$$

is provable in $\mathcal{T}_{d\forall H}$, hence also in \mathcal{T}_d^f ; the fact that the domain of f is an object means that $\exists y \exists z \lambda(x, y, z)$ is consistent with $\mathcal{T}_{d\forall H}$, hence also with \mathcal{T}_d^f . But $\exists y \theta(x, y)$ is an atom of \mathcal{T}_d^f , so the implication in (17) can be turned into an equivalence. Since a satisfies $\exists y \theta(x, y)$ in \mathcal{A} , it also satisfies $\exists y \exists z \lambda(x, y, z)$ and therefore lies in the range of the projection obtained by applying F to the morphism (16) in R . Thus, $F(R)$ is epimorphic and F is $J_{\neg\neg}$ -continuous. \square

This theorem implies that strictly atomic models of \mathcal{T}^f are automatically generic.

We conclude this paper with a discussion of the possibilities for equalities and strict inclusions in the chain

$$J_D \subseteq J_E \subseteq J_R \subseteq J_{\neg\neg}$$

of Theorem 5.

If \mathcal{T} , or equivalently \mathcal{T}_d , has a model companion \mathcal{T}^* , or \mathcal{T}_d^* , then the existentially closed models of \mathcal{T} , the finitely generic models of \mathcal{T} , and the models of \mathcal{T}^* are all the same in \mathcal{S} ; see [7]. In this situation, the topologies J_E and J_R are the same; in other words all existentially closed models are generic in arbitrary topoi, not just \mathcal{S} . To prove this, it suffices to check that each sieve $\|\alpha(x) \vee \neg\alpha(x)\|$ on $\{x \mid \text{true}\}$ is in J_E , for these sieves generate J_R . In every model (in \mathcal{S}) of \mathcal{T}_d^* , every a satisfies $\exists y \phi(a, y)$ for some condition $\phi(x, y)$ forcing $\alpha(x) \vee \neg\alpha(x)$, by genericity. A compactness argument shows that there are finitely many such ϕ 's, say ϕ_1, \dots, ϕ_n , such that $\mathcal{T}_d^* \vdash \forall x \bigvee_{i=1}^n \exists y_i \phi_i(x, y_i)$. It follows that no morphism into $\{x \mid \text{true}\}$ is incompatible with all n of the simple morphisms

$$f_i : \{x, y_i \mid \phi_i(x, y_i)\} \rightarrow \{x \mid \text{true}\}.$$

Thus, the intersection of the sieves $\pm f_i$, which is in J_E , consists entirely of morphisms that factor through an f_i and therefore lie in $\|\alpha(x) \vee \neg\alpha(x)\|$. Thus, $\|\alpha(x) \vee \neg\alpha(x)\|$ is also in J_E .

Of the eight possible patterns of equalities and strict inclusions in $J_D \subseteq J_E \subseteq J_R \subseteq J_{\neg\neg}$, the four that have $J_E = J_R$ all occur, with theories that have model companions. The simplest case, $J_D = J_E = J_R = J_{\neg\neg}$ occurs only in trivial situations, for this chain of equalities means that the classifying topos of the universal theory \mathcal{T}_d is Boolean, and it was shown in [2] that \mathcal{T}_d must then be the theory of a finite collection of finite models no one of which admits a homomorphism to another. We leave it to the reader to derive this description of \mathcal{T}_d directly from the assumption that every model of \mathcal{T}_d is a strongly atomic model of \mathcal{T}_d^f .

Pure equality theory and the theory of linear order, with a constant symbol added to conform to our convention, provide examples of $J_D \subsetneq J_E = J_R = J_{\neg\neg}$. More generally, so does any universal theory that is not model complete but has an \aleph_0 -categorical model companion; see [2].

If we add infinitely many constant symbols to pure equality theory, and if we add axioms saying that the constants are all distinct, the resulting theory is model complete and has $J_D = J_E = J_R \subsetneq J_{\neg\neg}$. Here the only (strongly) atomic model is the one where every element is denoted by a constant symbol. The same situation occurs, in a finite language, for the theory of $(\mathbb{Z}, 0, S, P)$ where S and P are the successor and predecessor functions.

The situation $J_D \subsetneq J_E = J_R \subsetneq J_{\neg\neg}$ occurs for most interesting companionable theories. For example, let \mathcal{T} be the theory of fields; then \mathcal{T}^* is the theory of algebraically closed fields, and the points of $(\mathcal{C}, J_{\neg\neg})$ are the algebraic closures of prime fields. Similarly, if \mathcal{T} is the theory of abelian groups, then \mathcal{T}^* is the theory of divisible groups with infinitely many elements of each finite order, and the points of $(\mathcal{C}, J_{\neg\neg})$ are those models of \mathcal{T}^* that are torsion groups. Perhaps the simplest example of this sort is the theory of one unary function and one constant, with no nonlogical axioms. \mathcal{T}^* is then given by axioms asserting about the function that every element has infinitely many pre-images and there are infinitely many cycles of every finite size. The points of $(\mathcal{C}, J_{\neg\neg})$ are the models of \mathcal{T}^* in which every element satisfies $f^m(x) = f^n(x)$ for some distinct m and n .

To get examples with $J_E \subsetneq J_R$, it is necessary to have $J_D \subsetneq J_E$ as well, because if all models of \mathcal{T}_d were existentially closed then \mathcal{T}_d would be model complete, i.e., it would serve as its own model companion, and we would have $J_E = J_R$. The standard examples, as in [7], of theories without model companions, such as the theory of groups, all have $J_D \subsetneq J_E \subsetneq J_R \subsetneq J_{\neg\neg}$. The last of these proper inclusions can, in the examples we have in mind, be deduced as follows from our results together with a result of Macintyre [11] asserting the existence of 2^{\aleph_0} non-isomorphic countable generic models under fairly general conditions. If $J_R = J_{\neg\neg}$, then, by Theorems 3 and 6, every generic model is a strongly atomic model of \mathcal{T}_d^f . It follows that the complete theory of any generic model is finitely axiomatized over \mathcal{T}_d^f , so, if \mathcal{T} is countable, there are only countably many such complete theories. It also follows from atomicity that any two countable generic models having the same complete theory are isomorphic. Thus, the existence of uncountably many non-isomorphic generic models implies $J_R \neq J_{\neg\neg}$. We leave it as an open problem

to find a universal theory (preferably a naturally occurring one) with $J_D \subsetneq J_E \subsetneq J_R = J_{\neg\neg}$.

More generally, it seems reasonable to look for a useful model-theoretic description of the theories for which $J_R = J_{\neg\neg}$, i.e., for which the finitely generic models are classified by a Boolean topos. (The equivalence of these two conditions follows from Lemma 1.2 of [2].) If, as before, we use the notation \mathcal{G} for the universal model of \mathcal{T}_{dVH} in $\mathcal{S}^{(\text{op})}$ and G for its underlying object, then J_R is the topology forcing, in the sense of [15], all subobjects of G^n that are definable by Robinson formulas to be complemented, while $J_{\neg\neg}$ forces *all* subobjects in $\mathcal{S}^{(\text{op})}$ to be complemented. Thus, we want to describe situations where every subobject is so closely related to definable subobjects of G^n that the property of complementation can be transferred from the latter to the former. We have already remarked that such a situation occurs when \mathcal{T} has an \aleph_0 -categorical model-companion. It can also be made to occur by so enriching the language that everything in $\mathcal{S}^{(\text{op})}$ becomes definable; the forcing languages (or Boolean-valued universes) used in set-theoretic forcing are an example of this phenomenon. For less rich languages, in particular for most countable theories, the topos of $J_{\neg\neg}$ -sheaves will have no points (set-theoretically generic objects, over the universe, do not really exist except in degenerate situations) while the topos of J_R -sheaves will have points (model-theoretically generic structures do really exist). From this point of view, set-theoretic forcing over countable standard models of set theory (rather than over the universe) is more similar to model-theoretic forcing, i.e., to J_R , than to $J_{\neg\neg}$. An exposition emphasizing this similarity is given in [10].

References

- [1] K.J. Barwise and A. Robinson, Completing theories by forcing, *Ann. Math. Logic* 2 (1970) 119–142.
- [2] A. Blass and A. Ščedrov, Boolean classifying topoi, *J. Pure Appl. Algebra* 28 (1983) 15–35.
- [3] A. Boileau and A. Joyal, La logique des topos, *J. Symbolic Logic* 46 (1981) 6–16.
- [4] C.C. Chang and H.J. Keisler, *Model Theory* (North-Holland, Amsterdam, 1973).
- [5] R. Diaconescu, Change of base for toposes with generators, *J. Pure Appl. Algebra* 6 (1975) 191–218.
- [6] A. Grothendieck and J.L. Verdier, *Séminaire de Géométrie Algébrique IV, Tome 1, Lecture Notes in Mathematics* 269 (Springer, Berlin–New York, 1972).
- [7] J. Hirschfeld and W. Wheeler, *Forcing, Arithmetic, and Division Rings, Lecture Notes in Mathematics* 454 (Springer, Berlin–New York, 1975).
- [8] P. Johnstone, *Topos Theory* (Academic Press, New York, 1977).
- [9] A. Joyal and G. Reyes, *Forcing and generic models in categorical logic*, preprint (1978).
- [10] H.J. Keisler, Forcing and the omitting types theorem, in: M. Morley, Ed., *Studies in Model Theory, MAA Studies in Mathematics*, Vol. 8 (1973) 96–133.
- [11] A. Macintyre, Omitting quantifier-free types in generic structures, *J. Symbolic Logic* 37 (1972) 512–520.
- [12] S. MacLane, *Categories for the Working Mathematician* (Springer, Berlin–New York, 1971).

- [13] M. Makkai and G. Reyes, *First Order Categorical Logic*, Lecture Notes in Mathematics 611, (Springer, Berlin–New York, 1977).
- [14] A. Ščedrov, *Sheaves and forcing and their metamathematical applications*, Thesis, SUNY at Buffalo (1981); to appear in *Memoirs Amer. Math. Soc.*
- [15] M. Tierney, *Forcing topologies and classifying topoi*, in: A. Heller and M. Tierney, Eds., *Algebra, Topology, and Category Theory: a collection of papers in honor of Samuel Eilenberg* (Academic Press, New York, 1976) 189–210.