

BOOLEAN CLASSIFYING TOPOI

Andreas BLASS*

Department of Mathematics, University of Michigan, Ann Arbor, MI 48109, USA

Andrej ŠČEDROV

Department of Mathematics, University of Pennsylvania, Philadelphia, PA 19104, USA

Communicated by P.J. Freyd

Received 19 October 1982

Let \mathcal{L} be a finitary geometric theory and \mathcal{C} its classifying topos. We prove that \mathcal{C} is Boolean if and only if (1) every first-order formula in the language of \mathcal{L} is \mathcal{L} -provably equivalent to a geometric formula and (2) for any finite list of variables, x , there are, up to \mathcal{L} -provable equivalence, only finitely many formulas, in the language of \mathcal{L} , with free variables among x . We use this characterization to show that, when \mathcal{C} is Boolean, it is an atomic topos and can be viewed as a finite coproduct of topos of continuous G -sets for topological groups G satisfying a certain finiteness condition.

AMS 1980 Classification: 03C35, 03G30, 18B25.

Key words and phrases: Classifying topos, Boolean topos, Atomic topos, Model-complete theory, \aleph_0 -categorical theory.

The primary purpose of the research reported here is to specify the model-theoretic properties that cause certain (geometric) theories, e.g. dense linear orders without endpoints, to have Boolean classifying topos while other closely related theories, e.g. dense linear orders, do not. The specification, given in Theorem 1 below, involves model completeness, a positivity requirement, a weak form of completeness, and a variant of \aleph_0 -categoricity. These conditions on the theory turn out to be so restrictive that they permit a quite detailed analysis of the classifying topos. The results of this analysis, summarized in Theorem 2, amount to a structure theory for coherent Boolean topos. In particular, they imply that all such topos are atomic.

* Partially supported by NSF grant MCS-8101560.

1. Preliminaries

Throughout this paper, ‘topos’ always means Grothendieck topos. The first-order languages, formulas, and theories that we consider are always assumed to be finitary. If \mathcal{T} is a theory, then ‘ \mathcal{T} -provable’, ‘consistent with \mathcal{T} ’, and related concepts are always to be understood as referring to classical logic. Thus, although the classifying topos of \mathcal{T} is defined in terms of intuitionistic models of \mathcal{T} in various topoi, our criterion for its Booleanness is expressed in terms of the traditional classical model theory of \mathcal{T} . We temporarily assume, for notational simplicity, that we are dealing with a single-sorted language. After the proof of Theorem 1, we indicate how our arguments can be applied to multi-sorted theories.

The *geometric* formulas of a first-order language L with equality are those obtainable from atomic formulas by finite conjunction (including the empty conjunction, *true*), finite disjunction (including the empty disjunction, *false*), and existential quantification. Geometric formulas can be characterized model-theoretically as the formulas whose satisfaction is preserved by arbitrary (not necessarily surjective) homomorphisms of L -structures; see [3, §5.2]. Topos-theoretically, their key property is that, if M is an L -structure in a topos \mathcal{E} , if $f: \mathcal{F} \rightarrow \mathcal{E}$ is a geometric morphism, and if ϕ is a geometric formula, then the truth values of ϕ in M and f^*M satisfy $\|\phi\|_{f^*M} = f^*\|\phi\|_M$, i.e. f^* preserves the truth value of ϕ .

A *geometric sequent* is a sentence of the form $\forall \mathbf{x} (\phi \rightarrow \psi)$, where ϕ and ψ are geometric formulas and every variable free in ϕ or ψ occurs in the list \mathbf{x} . (We systematically use boldface letters like \mathbf{x} for finite sequences.) A *geometric theory* is a theory axiomatized by geometric sequents. It follows from the preservation of geometric formulas under geometric morphisms that, if $f: \mathcal{F} \rightarrow \mathcal{E}$ is a geometric morphism of topoi and M is a model in \mathcal{E} of a geometric theory \mathcal{T} , then f^*M is a model of \mathcal{T} in \mathcal{F} . Barr’s theorem [5, §7.5] implies that, if one geometric sequent is deducible from certain others (in classical logic, according to our convention), then the deduction can be carried out in intuitionistic logic as well. Thus, the geometric sequents provable in a geometric theory \mathcal{T} hold in all models of \mathcal{T} in arbitrary topoi.

If \mathcal{T} is a geometric theory, then there exists a topos \mathcal{E} , called the *classifying topos* of \mathcal{T} , and there exists a model M of \mathcal{T} in \mathcal{E} , called the *universal model* of \mathcal{T} , such that, for any topos \mathcal{F} , the category of geometric morphisms $f: \mathcal{F} \rightarrow \mathcal{E}$ and natural transformations $\eta: f^* \rightarrow g^*$ is equivalent to the category of models of \mathcal{T} in \mathcal{F} and homomorphisms, the identity morphism of \mathcal{E} corresponding to the model M . Thus, M has the universal property that every model of \mathcal{T} in any topos \mathcal{F} can be obtained (up to isomorphism) as f^*M for a unique (up to natural isomorphism) geometric morphism $f: \mathcal{F} \rightarrow \mathcal{E}$. Of the various constructions of classifying topoi, the one of Joyal and Reyes described in [5, §7.4] will be useful in what follows, so we give an outline of it.

Given a geometric theory \mathcal{T} , we shall define its *syntactic site* (\mathcal{C}, J) ; the classifying

topos will be the topos of sheaves on this site. An object of \mathcal{C} is a formal class term $\{x \mid \phi(x)\}$ where $\phi(x)$ is a geometric formula whose free variables are among x . We adopt the convention that a formula or class term is not changed if its bound variables (including the x in $\{x \mid \phi(x)\}$) are renamed subject to the usual precautions for avoiding clashes. To define morphisms from $\{x \mid \phi(x)\}$ to $\{y \mid \psi(y)\}$, we may and do assume that the lists x and y are disjoint; a morphism is then an equivalence class, with respect to \mathcal{F} -provable equivalence, of geometric formulas $\theta(x, y)$ such that the following geometric sequents are provable in \mathcal{F} :

$$\forall x \forall y (\theta(x, y) \rightarrow \phi(x) \wedge \psi(y)),$$

$$\forall x (\phi(x) \rightarrow \exists y \theta(x, y)),$$

$$\forall x \forall y \forall z (\theta(x, y) \wedge \theta(x, z) \rightarrow y = z).$$

(If y is y_1, \dots, y_n and z is z_1, \dots, z_n , then $\exists y$ is $\exists y_1 \dots \exists y_n$ and $y = z$ is the conjunction of the n formulas $y_i = z_i$.) The morphism defined by $\theta(x, y)$ will be called $[x \rightarrow y \mid \theta(x, y)]$. The composite of $[x \rightarrow y \mid \theta(x, y)]$ and $[y \rightarrow z \mid \eta(y, z)]$ is defined to be $[x \rightarrow z \mid \exists y (\theta(x, y) \wedge \eta(y, z))]$. It is easy to check that this defines a category \mathcal{C} . We make it a site by defining a sieve on $\{y \mid \psi(y)\}$ to be J -covering if and only if it contains a finite family of morphisms

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

for which the geometric sequent

$$\forall y (\psi(y) \rightarrow \bigvee_i \exists x_i \theta_i(x_i, y))$$

is provable in \mathcal{F} ; it is easy to check that J is a Grothendieck topology. Finally, according to Theorem 7.45 of [5], the topos $\mathcal{E}(\mathcal{F})$ of sheaves on the site (\mathcal{C}, J) is a classifying topos for the theory \mathcal{F} .

Observe that $\{y \mid \psi(y)\}$ is covered by the empty sieve if and only if $\psi(y)$ is inconsistent with \mathcal{F} , i.e. \mathcal{F} proves the geometric sequent $\forall y (\psi(y) \rightarrow \text{false})$. Such objects can be deleted from the site without changing the topos of sheaves, according to [4, III. 4.1]. More precisely, we let \mathcal{C}' be the full subcategory of \mathcal{C} whose objects are $\{x \mid \phi(x)\}$ for geometric $\phi(x)$ consistent with \mathcal{F} , and we let J' be the topology on \mathcal{C}' induced by J . Then $\mathcal{E}(\mathcal{F})$ is also the topos of sheaves on (\mathcal{C}', J') . We note for future reference that, since J' -coverings are never empty, J' is included in the double-negation topology on \mathcal{C}' . We also note that, if

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

is any morphism of \mathcal{C}' , then $\theta(x, y)$ is consistent with \mathcal{F} . This follows immediately from the facts that $\phi(x)$ is consistent with \mathcal{F} , by definition of \mathcal{C}' , and that \mathcal{F} proves $\forall x (\phi(x) \rightarrow \exists y \theta(x, y))$, by definition of morphisms.

2. Theories classified by Boolean topoi

Theorem 1. *The classifying topos $\mathcal{E}(\mathcal{T})$ of a geometric theory \mathcal{T} is Boolean if and only if both of the following conditions are satisfied.*

(a) *Every formula in the language of \mathcal{T} is \mathcal{T} -provably equivalent to a geometric formula.*

(b) *For every finite list \mathbf{x} of variables, there are, up to \mathcal{T} -provable equivalence, only finitely many formulas with free variables among \mathbf{x} .*

Before proving the theorem, we comment on the conditions (a) and (b). Recall our convention that ‘ \mathcal{T} -provable’ refers to classical logic. The preservation theorems in [3, §5.2] imply that (a) is equivalent to requiring that every formula is \mathcal{T} -provably equivalent to a positive formula and to an existential one. The positivity requirement trivially reduces to

(a₁) *For each predicate symbol R of L (including equality), there is a positive formula $\phi(\mathbf{x})$ \mathcal{T} -provably equivalent to $\neg R(\mathbf{x})$.*

And equivalence of arbitrary formulas to existential formulas means

(a₂) *\mathcal{T} is model-complete.*

Condition (b) can also be expressed in more familiar terms. First, taking the list \mathbf{x} to be empty, we infer from (b) that there are only finitely many inequivalent sentences, i.e.

(b₁) *\mathcal{T} has only finitely many completions.*

Condition (b) for a theory \mathcal{T} clearly implies the same condition for all the completions of \mathcal{T} . In the presence of (b₁), the converse implication holds as well, for the \mathcal{T} -provable-equivalence class of a formula ϕ is uniquely determined by the list of \mathcal{T}' -provable-equivalence classes of ϕ as \mathcal{T}' ranges over all completions of \mathcal{T} . Thus, condition (b) for \mathcal{T} is equivalent to the conjunction of (b₁) with ‘‘all completions of \mathcal{T} satisfy (b)’’. For countable complete theories, (b) is Ryll–Nardzewski’s criterion [3, Theorem 2.3.13] for \aleph_0 -categoricity (where we include among \aleph_0 -categorical theories the trivial ones whose models are finite). In the case of an uncountable complete theory, condition (b) still implies \aleph_0 -categoricity, since this half of Ryll–Nardzewski’s theorem does not require countability of the theory. (\aleph_0 -categoricity may, of course, be vacuous; the theory need not have any countable models.) We shall refer to condition (b), for complete theories, as *persistent \aleph_0 -categoricity*, and we digress briefly to justify this terminology.

If a complete theory \mathcal{T} satisfies (b), then it continues to do so in every Boolean extension $V^\#$ of the universe V of sets. More precisely the truth value of ‘‘ \mathcal{T} satisfies (b)’’ is 1, and therefore, by our previous remarks, so is ‘‘ \mathcal{T} is \aleph_0 -categorical’’. Thus, the \aleph_0 -categoricity of \mathcal{T} persists when we pass to $V^\#$. Conversely, if \mathcal{T} remains \aleph_0 -categorical in every $V^\#$, then in particular we can choose $\#$ so that the cardinality of \mathcal{T} is collapsed to \aleph_0 in $V^\#$. Then Ryll–Nardzewski’s theorem, applied in $V^\#$ to the countable (in $V^\#$) \aleph_0 -categorical theory \mathcal{T} , tells us that \mathcal{T} satisfies (b) in $V^\#$. But (b) is clearly absolute, so \mathcal{T} really satisfies (b). This

shows that (b) is equivalent to the assertion that \mathcal{F} is \aleph_0 -categorical in $V^{\mathcal{B}}$ for all \mathcal{B} ; hence the terminology ‘persistently \aleph_0 -categorical’.

With this terminology, we have that (b) is equivalent to the conjunction of (b₁) and

(b₂) *For each completion of \mathcal{F} is persistently \aleph_0 -categorical.*

We emphasize that, for countable languages, the word ‘persistently’ becomes vacuous. Summarizing this discussion, we have:

Corollary 1. *$\mathcal{E}(\mathcal{F})$ is Boolean if and only if \mathcal{F} satisfies the four condition (a₁), (a₂), (b₁), and (b₂).*

The preservation theorems previously cited also yield the following reformulation of (a) as a strong form of model-completeness.

(a’) *Every homomorphism from one model of \mathcal{F} into another is an elementary embedding.*

From this point of view, (a₁) says that (a’) minus the word ‘elementary’ holds, and (a₂) reinstates the omitted word.

Proof of Theorem 1. Assume that $\mathcal{E}(\mathcal{F})$ is Boolean. In accordance with our discussion in Section 1, we represent $\mathcal{E}(\mathcal{F})$ as the topos of sheaves on the site (\mathcal{C}', J') , where J' is included in the double negation topology. Our first objective is to infer that J' coincides with the double negation topology. We begin with a general fact about the double-negation topologies; it is probably folklore, but we give a proof since we have not seen one in the literature.

Lemma 1.1. *Let \mathcal{F} be any topos and let j be a topology in \mathcal{F} such that $j \leq \neg\neg$. Then the topoi of double-negation sheaves in \mathcal{F} and in $\text{sh}_j(\mathcal{F})$ are equivalent.*

Proof. In the diagram

$$\text{sh}_{\neg\neg}(\text{sh}_j(\mathcal{F})) \xrightarrow{g} \text{sh}_j(\mathcal{F}) \xrightarrow{f} \mathcal{F},$$

where $\neg\neg$ is the double-negation topology in $\text{sh}_j(\mathcal{F})$, the composite fg of the two inclusions, being an inclusion, is, by Proposition 4.15 of [5], equivalent to the inclusion of $\text{sh}_k(\mathcal{F})$ in \mathcal{F} , where k is some topology in \mathcal{F} . We must show that k is the double-negation topology. Half of this is easy, in view of the following characterization of the double-negation topology, which is essentially Proposition 5.18 of [5]: A topology is smaller than the double-negation topology if and only if 0 is not dense in any non-zero object, i.e. if and only if the associated sheaf functor sends non-zero objects to non-zero objects. Now, by two applications of this criterion, a non-zero object of \mathcal{F} goes, under f^* , to a non-zero object of $\text{sh}_j(\mathcal{F})$, which in turn goes, under g^* , to a non-zero object of $\text{sh}_{\neg\neg}(\text{sh}_j(\mathcal{F})) \cong \text{sh}_k(\mathcal{F})$. Another application of

the criterion yields that $k \leq \neg\neg$. To prove the reverse inequality, we consider an arbitrary subobject A of an arbitrary object B in \mathcal{F} , and we prove that $\neg\neg A \subseteq k(A)$. Pulling back from B to $\neg\neg A$, we may assume, without loss of generality, that $A \hookrightarrow B$ is $\neg\neg$ -dense, and we must show that it is also k -dense, i.e. that g^*f^* sends it to an isomorphism. This means, in view of the definition of g , that we must show that $f^*A \hookrightarrow f^*B$ is $\neg\neg$ -dense in $\text{sh}_j(\mathcal{F})$. So we consider an arbitrary subobject of f^*B disjoint from f^*A and show that it is 0. This arbitrary subobject is $f^*C \hookrightarrow f^*B$ for some $C \hookrightarrow B$ in \mathcal{F} (namely, the C whose classifying map is $B \rightarrow \Omega_j \subseteq \Omega_\delta$, where $\Omega_j = f_*\Omega_{\mathcal{F}}$ and the map $B \rightarrow f_*\Omega_{\mathcal{F}}$ comes, via adjointness, from the classifying map $f^*B \rightarrow \Omega_{\mathcal{F}}$ of the given subobject). Disjointness of f^*C from f^*A means, since f^* preserves intersections, that $f^*(A \cap C) = 0$. But $j \leq \neg\neg$ so, by our criterion for topologies smaller than $\neg\neg$, we infer that $A \cap C = 0$. As A is $\neg\neg$ -dense in B , we have $C = 0$, so $f^*C = 0$, as required. This completes the proof of Lemma 1.1.

Lemma 1.2. *Let \mathcal{F} and j be as in the preceding lemma. If $\text{sh}_j(\mathcal{F})$ is Boolean, then $j = \neg\neg$.*

Proof. If $\text{sh}_j(\mathcal{F})$ is Boolean, then it is its own double-negation sheaf subtopos. By Lemma 1.1, $\text{sh}_j(\mathcal{F})$ and $\text{sh}_{\neg\neg}(\mathcal{F})$ are equivalent (as subtopoi of \mathcal{F} , by the proof of Lemma 1.1). Therefore, $j = \neg\neg$, as required.

We return to the proof of Theorem 1 and apply Lemma 1.2 with \mathcal{F} being the topos of presheaves on \mathcal{X}' and j being the topology in \mathcal{F} determined by the topology J' on \mathcal{X}' . Then $\text{sh}_j(\mathcal{F}) = \delta(\mathcal{F})$ is Boolean, so $j = \neg\neg$, which means that J' is the double-negation topology on \mathcal{X}' .

Consider now an arbitrary model M of \mathcal{F} (in the topos \mathcal{U} of sets) and an arbitrary tuple \mathbf{a} of elements of M . Fix a list \mathbf{x} of variables, of the same length as \mathbf{a} , and let Φ be the set of all geometric formulas $\phi(\mathbf{x})$, with free variables among \mathbf{x} , that are *not* true of \mathbf{a} in M ,

$$\Phi = \{\phi(\mathbf{x}) \mid M \models \neg\phi(\mathbf{a})\}.$$

The morphisms of \mathcal{X}'

$$[\mathbf{x} \mapsto \mathbf{y} \mid \phi(\mathbf{x}) \wedge \mathbf{x} = \mathbf{y}]: \quad \{\mathbf{x} \mid \phi(\mathbf{x})\} \rightarrow \{\mathbf{y} \mid \text{true}\}, \quad (1)$$

for $\phi \in \Phi$, do not J' -cover $\{\mathbf{y} \mid \text{true}\}$, for, if they did, there would be finitely many $\phi_i \in \Phi$ such that

$$\forall \mathbf{y} [\text{true} \rightarrow \bigvee_i \exists \mathbf{x} (\phi_i(\mathbf{x}) \wedge \mathbf{x} = \mathbf{y})],$$

i.e.

$$\forall \mathbf{y} \bigvee_i \phi_i(\mathbf{y}),$$

is provable in \mathcal{F} , hence true in M . But this contradicts the definition of Φ .

Since J' is the double-negation topology, the collection of morphisms (1) is not a $\neg\neg$ -cover of $\{\mathbf{y} \mid \text{true}\}$, so we can find a morphism in \mathcal{X}'

$$\alpha = [z \mapsto y \mid \theta(z, y)]: \{z \mid \psi(z)\} \rightarrow \{y \mid \text{true}\}$$

such that no morphism into $\{y \mid \text{true}\}$ factors through α and also through one of the morphisms (1). Now $\theta(z, y)$ is consistent with \mathcal{F} , by the observation at the end of Section 1. We assert that, for $\phi \in \Phi$, $\theta(z, y) \wedge \phi(y)$ is *not* consistent with \mathcal{F} .

To prove this assertion, suppose it were false for a certain $\phi \in \Phi$ and consider the morphism

$$[v, w \mapsto y \mid y = w \wedge \theta(v, w) \wedge \phi(w)]: \{v, w \mid \theta(v, w) \wedge \phi(w)\} \rightarrow \{y \mid \text{true}\}.$$

It is easy to check that it factors through α , via

$$[v, w \mapsto z \mid v = z \wedge \theta(v, w) \wedge \phi(w)]: \{v, w \mid \theta(v, w) \wedge \phi(w)\} \rightarrow \{z \mid \psi(z)\},$$

and factors through the morphism (1) associated with ϕ , via

$$[v, w \mapsto x \mid x = w \wedge \theta(v, w) \wedge \phi(w)]: \{v, w \mid \theta(v, w) \wedge \phi(w)\} \rightarrow \{x \mid \phi(x)\}.$$

This contradicts the choice of α , so the assertion is proved.

We have thus found a geometric formula, $\exists z \theta(z, y)$, henceforth abbreviated as $\gamma(y)$, which is consistent with \mathcal{F} and \mathcal{F} -provably implies $\neg\phi(y)$ for every $\phi \in \Phi$. Recall that Φ consisted of the geometric formulas not satisfied by a specific tuple \mathbf{a} in a specific model M of \mathcal{F} . Clearly, $\gamma(\mathbf{a})$ must hold in M , as otherwise $\gamma(x)$ would be in Φ and would therefore be inconsistent with itself and \mathcal{F} . Thus we have, for any \mathbf{a} and M as above, a geometric formula $\gamma(x)$, true of \mathbf{a} in M , that \mathcal{F} -provably implies every negated-geometric formula true of \mathbf{a} in M .

We are now ready to prove assertion (a) of the theorem, by induction on formulas. Since (a) refers to classical provability and since the class of geometric formulas contains the atomic formulas and is closed under conjunction, disjunction, and existential quantification, the only point requiring proof is that the negation of a geometric formula is \mathcal{F} -provably equivalent to a geometric formula. So let $\phi(x)$ be any geometric formula, and let Ψ be the collection of all geometric formulas $\psi(x)$, with the same free variables, that \mathcal{F} -provably imply $\neg\phi(x)$. Suppose, toward a contradiction, that in some model M of \mathcal{F} there were elements \mathbf{a} satisfying $\neg\phi(\mathbf{a})$ and simultaneously satisfying $\neg\psi(\mathbf{a})$ for all $\psi(x) \in \Psi$. By what we proved above, there is a geometric formula $\gamma(x)$, satisfied by \mathbf{a} , and \mathcal{F} -provably implying $\neg\phi(x)$. By the latter, $\gamma(x) \in \Psi$, but then $\neg\gamma(\mathbf{a})$ holds, a contradiction. This contradiction shows that the set of formulas

$$\{\neg\phi(x)\} \cup \{\neg\psi(x) \mid \psi(x) \in \Psi\}$$

is inconsistent with \mathcal{F} . By the compactness theorem, a finite subset, say

$$\{\neg\phi(x)\} \cup \{\neg\psi_i(x) \mid 1 \leq i \leq n\}$$

is inconsistent with \mathcal{F} . This means that $\neg\phi(x)$ \mathcal{F} -provably implies the geometric formula $\bigvee_{i=1}^n \psi_i(x)$; the converse implication also holds because each $\psi_i(x)$ is in Ψ . This completes the proof of (a).

Our earlier construction of the formulas $\gamma(\mathbf{x})$ yields, in view of (a), the following fact. For any elements \mathbf{a} in any model M of \mathcal{T} , there is a formula $\gamma(\mathbf{x})$, satisfied by \mathbf{a} that \mathcal{T} -provably implies every formula satisfied by \mathbf{a} . In other words, the type realized by \mathbf{a} is principal. Since every n -type (=ultrafilter in the Lindenbaum algebra of \mathcal{T} -provable-equivalence classes of formulas with free variables among x_1, \dots, x_n) is realized in some model of \mathcal{T} , every n -type is principal. But this means that the Lindenbaum algebra is finite, so (b) holds. This completes the proof of the ‘only if’ part of Theorem 1.

To prove the ‘if’ part, let \mathcal{T} be a geometric theory satisfying (a) and (b). To show that $\mathcal{L}(\mathcal{T})$ is Boolean, we show that the topology J' on \mathcal{C}' coincides with the double-negation topology. That J' is included in the double-negation topology was already observed, for arbitrary \mathcal{T} , in Section 1. To prove the reverse inclusion, consider any $\neg\neg$ -covering, say the family

$$[\mathbf{x}_i \multimap \mathbf{y} \mid \theta_i(\mathbf{x}_i, \mathbf{y})]: \{\mathbf{x}_i \mid \phi_i(\mathbf{x}_i)\} \rightarrow \{\mathbf{y} \mid \psi(\mathbf{y})\} \quad (2)$$

where i ranges over some index set I . We shall show that it is also a J' -covering.

By (b), there are, up to \mathcal{T} -provable equivalence, only finitely many formulas of the form

$$\exists \mathbf{x}_i \theta_i(\mathbf{x}_i, \mathbf{y}), \quad \text{with } i \in I. \quad (3)$$

Fix a finite set $I_0 \subseteq I$ such that each formula (3) is \mathcal{T} -provably equivalent to one with $i \in I_0$. We shall show that the morphisms (2) for $i \in I_0$ have the property required in the definition of J' (or J), i.e. that \mathcal{T} proves

$$\forall \mathbf{y} (\psi(\mathbf{y}) \rightarrow \bigvee_{i \in I_0} \exists \mathbf{x}_i \theta_i(\mathbf{x}_i, \mathbf{y})). \quad (4)$$

To show this, which will complete the proof, we assume it is false and derive a contradiction. So assume that (4) is not \mathcal{T} -provable. Thus

$$\psi(\mathbf{y}) \wedge \neg \bigvee_{i \in I_0} \exists \mathbf{x}_i \theta_i(\mathbf{x}_i, \mathbf{y}) \quad (5)$$

is consistent with \mathcal{T} . By (a), find a geometric formula $\eta(\mathbf{z})$ \mathcal{T} -provably equivalent to (5), and consider the morphism

$$[\mathbf{z} \multimap \mathbf{y} \mid \eta(\mathbf{z}) \wedge \mathbf{z} = \mathbf{y}]: \{\mathbf{z} \mid \eta(\mathbf{z})\} \rightarrow \{\mathbf{y} \mid \psi(\mathbf{y})\}. \quad (6)$$

Since the family (2) is a $\neg\neg$ -covering, there must be a morphism, say

$$[\mathbf{v} \multimap \mathbf{y} \mid \gamma(\mathbf{v}, \mathbf{y})]: \{\mathbf{v} \mid \mu(\mathbf{v})\} \rightarrow \{\mathbf{y} \mid \psi(\mathbf{y})\}, \quad (7)$$

that factors through both (6) and one of the morphisms (2), as in the diagram.

$$\begin{array}{ccc} \{\mathbf{v} \mid \mu(\mathbf{v})\} & \xrightarrow{[\mathbf{v} \multimap \mathbf{z} \mid \alpha(\mathbf{v}, \mathbf{z})]} & \{\mathbf{z} \mid \eta(\mathbf{z})\} \\ \downarrow [\mathbf{v} \multimap \mathbf{x} \mid \beta(\mathbf{v}, \mathbf{x})] & \searrow [\mathbf{v} \multimap \mathbf{y} \mid \gamma(\mathbf{v}, \mathbf{y})] & \downarrow [\mathbf{z} \multimap \mathbf{y} \mid \mathbf{y} = \mathbf{z} \wedge \eta(\mathbf{z})] \\ \{\mathbf{x}_i \mid \phi_i(\mathbf{x}_i)\} & \xrightarrow{[\mathbf{x}_i \multimap \mathbf{y} \mid \theta_i(\mathbf{x}_i, \mathbf{y})]} & \{\mathbf{y} \mid \psi(\mathbf{y})\} \end{array}$$

We pointed out, at the end of Section 1, that $\gamma(v, y)$ must be consistent with \mathcal{F} . To complete the proof, we shall derive, in \mathcal{F} , two contradictory consequences from $\gamma(v, y)$.

Working in \mathcal{F} , assume $\gamma(v, y)$. By the commutativity of the bottom triangle, we have

$$\exists x_i (\beta(v, w) \wedge \theta_i(x_i, y))$$

and, in particular, $\exists x_i \theta_i(x_i, y)$. By our choice of I_0 , we also have

$$\forall_{i \in I_0} \exists x_i \theta_i(x_i, y).$$

Therefore $\neg \eta(y)$. On the other hand, by commutativity of the top triangle, we have

$$\exists z (\alpha(v, z) \wedge y = z \wedge \eta(z))$$

which implies $\eta(y)$. This contradiction completes the proof of Theorem 1.

It is useful, for applications in the next section, to observe that the preceding work applies to many-sorted theories. The only changes needed are, first, the insertion of requirements that the sorts of the variables and elements used in the proof match properly, and, second, the remarks that countability for many-sorted languages includes countability of the set of sorts and countability of a many-sorted structure means countability of the union of all the domains.

We close this section with a collection of examples showing that the conditions in Corollary 1 are independent.

Example 1. Let L be a 0-sorted (i.e. propositional) language with just one (necessarily 0-ary) relation symbol P (a propositional variable). Let \mathcal{F} be the theory with no axioms. Then conditions (a₂) and (b) are satisfied, but (a₁) fails. The classifying topos is the Sierpiński topos [5, Example 4.37(ii)].

Example 2. Let L be the 1-sorted language with a single binary relation symbol. Let \mathcal{F} be the theory of dense linear orderings. Then (a₁) holds, as the negations of $x < y$ and of $x = y$ are equivalent to $x = y \vee y < x$ and to $x < y \vee y < x$ respectively. (b₁) holds, as there are precisely five completions, obtained by specifying which endpoints exist and, if both exist, whether they are equal. It is well known that these completions are \aleph_0 -categorical, so (b₂) holds. But (a₂) fails, since the formula $\forall y (x < y \vee x = y)$, asserting that x is a left endpoint, is not preserved by embeddings and is therefore not \mathcal{F} -provably equivalent to an existential formula. If \mathcal{A} is the category of finite pointed linearly ordered sets and strictly order-preserving maps preserving the distinguished points, and if \mathcal{B} is the subcategory with the same objects but only those morphisms that preserve endpoints, then the methods of [6] show that the classifying topos of \mathcal{F} is the topos of sheaves on \mathcal{A}^{op} with respect to the topology generated by coverings that consist of a single morphism in \mathcal{B} .

For the remaining two examples, we shall need a process, sometimes called

Morleyization, for making all formulas equivalent to atomic ones by adding new predicate symbols. The idea is that, beginning with a theory \mathcal{T}_0 in a language L_0 , one adds, for each list \mathbf{x} of n variables and each formula $\phi(\mathbf{x})$ with free variables among \mathbf{x} , a new n -ary predicate symbol $P_{\mathbf{x},\phi(\mathbf{x})}$ and a new axiom $\forall \mathbf{x} (P_{\mathbf{x},\phi(\mathbf{x})}(\mathbf{x}) \leftrightarrow \phi(\mathbf{x}))$. In this form, the process involves axioms that are not geometric sequents; since we want geometric theories we modify the axioms as follows. For simplicity we write $\forall \mathbf{x} (\alpha \leftrightarrow \beta)$ instead of the two geometric sequents $\forall \mathbf{x} (\alpha \rightarrow \beta)$ and $\forall \mathbf{x} (\beta \rightarrow \alpha)$. With this convention, the new axioms are

$$\begin{aligned} &\forall \mathbf{x} (P_{\mathbf{x},\phi(\mathbf{x})}(\mathbf{x}) \leftrightarrow \phi(\mathbf{x})) \quad \text{for atomic } \phi, \\ &\forall \mathbf{x} (P_{\mathbf{x},\phi(\mathbf{x}) \wedge \psi(\mathbf{x})}(\mathbf{x}) \leftrightarrow P_{\mathbf{x},\phi(\mathbf{x})}(\mathbf{x}) \wedge P_{\mathbf{x},\psi(\mathbf{x})}(\mathbf{x})), \\ &\text{a similar clause for } \vee, \\ &\forall \mathbf{x} (P_{\mathbf{x},\exists y \phi(\mathbf{x},y)}(\mathbf{x}) \leftrightarrow \exists y P_{\mathbf{x},y,\phi(\mathbf{x},y)}(\mathbf{x},y)), \\ &\forall \mathbf{x} (\text{true} \rightarrow P_{\mathbf{x},\phi(\mathbf{x})}(\mathbf{x}) \vee P_{\mathbf{x},\neg\phi(\mathbf{x})}(\mathbf{x})), \\ &\forall \mathbf{x} (P_{\mathbf{x},\phi(\mathbf{x})}(\mathbf{x}) \wedge P_{\mathbf{x},\neg\phi(\mathbf{x})}(\mathbf{x}) \rightarrow \text{false}). \end{aligned}$$

(The universal quantifier and other connectives are to be treated as defined symbols.) An easy induction on formulas ϕ shows that $\forall \mathbf{x} (P_{\mathbf{x},\phi(\mathbf{x})}(\mathbf{x}) \leftrightarrow \phi(\mathbf{x}))$ follows from these axioms. Clearly, the theory obtained by Morleyization of \mathcal{T}_0 always satisfies condition (a) of Theorem 1, and it satisfies (b₁) or (b₂) if and only if \mathcal{T}_0 does. Note also that the Morleyization of \mathcal{T}_0 is always geometric, even if \mathcal{T}_0 is not, since each axiom α of \mathcal{T}_0 can be replaced by P_α .

Example 3. Let L_0 be a 1-sorted language without nonlogical symbols. Let \mathcal{T}_0 have no axioms. This theory, pure equality theory, has infinitely many completions (specifying a finite cardinality for the universe or saying that is infinite), all of which are \aleph_0 -categorical. Therefore, the Morleyization satisfies (a) and (b₂) but not (b₁).

Example 4. Let \mathcal{T}_0 be any complete countable first-order theory that is not \aleph_0 -categorical, for instance the theory of the natural numbers with the successor function. Its Morleyization satisfies (a) and (b₁) but not (b₂).

In [2], we work with the classifying topoi of universal geometric theories, i.e. theories in whose axioms $\forall \mathbf{x} (\phi \rightarrow \psi)$ the geometric formulas ϕ and ψ contain no quantifiers. The following corollary of Theorem 1 tells us that these classifying topoi are practically never Boolean.

Corollary 2. *The classifying topos δ of a universal geometric theory \mathcal{T} is Boolean if and only if \mathcal{T} is the theory of a finite collection $\{M_1, \dots, M_n\}$ of finite models such that, for $i \neq j$, no homomorphism from M_i into M_j exists.*

Proof. Suppose first that the universal geometric theory \mathcal{T} is the theory of $\{M_1, \dots, M_n\}$, with the M_i as in the statement of the corollary. Observe that since the isomorphism class of a finite structure is an elementary class, every model of \mathcal{T} is isomorphic to one of the M_i . In particular, every substructure of M_j , being a model of \mathcal{T} because \mathcal{T} is universal, must be among the M_i ; the assumption about non-existence of homomorphisms then requires the substructure to be M_j itself. Thus, in each M_j , there are no proper substructures, so every element is the denotation of a closed term. This conclusion, to which we refer as the existence of enough names, will be used repeatedly in what follows. Observe that it implies that the only homomorphism from M_i to itself is the identity.

For $i \neq j$, there must be an atomic sentence true in M_i but false in M_j , because otherwise we could, thanks to the existence of enough names, define a homomorphism from M_i to M_j by sending the denotation in M_i of any closed term to the denotation in M_j of the same term. If we fix i and let j vary, the conjunction δ_i of these atomic sentences will be true in M_i and false in all the other M_j .

Let \mathcal{T}' be the theory whose axioms are:

- (a) $true \rightarrow \bigvee_{i=1}^n \delta_i$,
- (b) $\delta_i \wedge \delta_j \rightarrow false$, for $i \neq j$,
- (c) $\delta_i \rightarrow \alpha$, for α an atomic sentence true in M_i ,
- (d) $\delta_i \wedge \beta \rightarrow false$, for β an atomic sentence false in M_i ,
- (e) $\forall x (true \rightarrow \bigvee_{t \in N} (x = t))$,

where N is a finite set of closed terms large enough to contain a name for each element of each M_i . It is obvious that each M_i is a model of \mathcal{T}' . We shall prove that every model of \mathcal{T}' is isomorphic to one of the M_i , so that \mathcal{T}' is equivalent to \mathcal{T} .

In fact, we shall prove somewhat more, namely that for any model M of \mathcal{T}' in any topos \mathcal{F} , there exists a partition of 1 into open sub-objects U_1, \dots, U_n in \mathcal{F} such that, over U_i , M is isomorphic to M_i . (More precisely, if $\mathcal{F}/U_i \xrightarrow{f} \mathcal{F} \xrightarrow{p} \mathcal{G}$ are the obvious geometric morphisms then $f^*M \cong f^*p^*M_i$.) To see this, let M be given and define U_i to be the truth value of δ_i ; this defines a partition of 1 because of axioms (a) and (b) of \mathcal{T}' . For the rest of the argument, we fix an i and work in \mathcal{F}/U_i , where M satisfies δ_i . By axioms (c) and (d), any atomic sentence true (resp. false) in M_i is also true (resp. false) in M . This means, since there are enough names, that we can embed M_i into M by sending the value in M_i of any closed term to the value in M of the same closed term. This embedding is an isomorphism because M satisfies axiom (e).

This description of the \mathcal{F} -models of \mathcal{T}' (which we now know to be equivalent to \mathcal{T}), along with the observation that a homomorphism between two such models can exist only when the associated partitions are the same (because the δ_i are positive sentences) and must then be the obvious isomorphism (induced by the isomorphisms to the M_i , because there are enough names), tells us that the category of \mathcal{F} -models of \mathcal{T} is equivalent to the discrete category of partitions of 1 into n labeled pieces in \mathcal{F} . The classifying topos for such partitions is clearly \mathcal{G}/n , which is Boolean.

To prove the converse, we assume that \mathcal{F} is countable. This assumption involves

no loss of generality, for the hypothesis that $\mathcal{E}(\mathcal{T})$ is Boolean, when expressed as (a) and (b) of Theorem 1, is clearly preserved when we pass to any Boolean extension of the universe, e.g. an extension in which \mathcal{T} is countable, and the desired conclusion is clearly preserved when we return to the original universe.

So let \mathcal{T} be a countable universal geometric theory satisfying the conditions in Theorem 1 and thus also the conditions in Corollary 1.

Consider an arbitrary countable model M of \mathcal{T} . The substructure consisting of denotations of closed terms is a model of \mathcal{T} (as \mathcal{T} is universal), hence an elementary substructure of M (as \mathcal{T} is model-complete), hence isomorphic to M (as all completions of \mathcal{T} are \aleph_0 -categorical). We infer that every element of M is the denotation of a closed term, since this is the case for the substructure that we have seen is isomorphic to M . If M were infinite, then we would obtain a contradiction by applying the preceding discussion to a countable proper elementary extension of M . So all models of \mathcal{T} are finite.

\mathcal{T} can have only finitely many non-isomorphic models, because non-isomorphic finite models are not elementarily equivalent and \mathcal{T} has only finitely many completions. Finally, any homomorphism between models of \mathcal{T} is an elementary embedding (by (a')) and therefore an isomorphism (by finiteness). This completes the proof of Corollary 2.

Notice that Example 1 above exhibits a theory whose classifying topos fails to be Boolean despite the fact that the theory has only finitely many models all of which are finite. The conditions in Corollary 2 are not satisfied because the model in which P is false has a (vacuous) homomorphism to the one in which P is true.

3. Coherent Boolean topoi

Using Theorem 1, we shall obtain a rather complete description of coherent Boolean topoi. Before stating this description, Theorem 2 below, we recall a few definitions and introduce one new definition.

Coherent topoi can be defined in two equivalent ways. First, they are the classifying topoi of geometric theories. (Recall that all of our theories are finitary.) Second, they are the topoi of sheaves on sites where finite limits exist and every covering sieve has a finite subset that generates a covering sieve. For the equivalence of these two definitions, see [5, §7.4]; in one direction the proof uses the description of $\mathcal{E}(\mathcal{T})$ as the topos of sheaves on the site (\mathcal{C}, J) defined in Section 1.

Atomic topoi [1] are the topoi of sheaves on *atomic sites*, i.e. sites whose covering sieves are precisely all the nonempty sieves. These topoi are characterized [1] by the property that the ‘constant sheaf’ functor from the topos of sets is logical. Note that, since an atomic site need not have finite limits, we cannot conclude that all atomic topoi are coherent. In fact, a simple counter-example is obtained by taking the (underlying category of the) atomic site to be the monoid of one-to-one functions from natural numbers to natural numbers.

Recall that the coproduct of two topoi, in the sense of geometric morphisms, is their product as categories. If the two topoi are given as the topos of sheaves over two sites, then their coproduct is the topos of sheaves on the disjoint union of the two sites (with the obvious topology).

We define a topological group G to be *coherent* if, for every open subgroup H , the number of double cosets HgH , with $g \in G$, is finite.

Theorem 2. *For any topos \mathcal{E} , the following are equivalent.*

- (i) \mathcal{E} is coherent and atomic.
- (ii) \mathcal{E} is coherent and Boolean.
- (iii) \mathcal{E} is the classifying topos of a theory \mathcal{T} with the properties in Theorem 1.
- (iv) \mathcal{E} is the coproduct of finitely many topoi each of which is the topos of continuous G -sets for some coherent topological group G .

Proof. Theorem 1 and the first definition of coherent topoi immediately yield that (ii) implies (iii). Also, (i) implies (ii) trivially, since all atomic topoi are Boolean. To complete the proof, we show that (iii) implies (iv) and that (iv) implies (i); in fact, our proof of the former also establishes directly that (iii) implies (i) and thus establishes the equivalence of (i), (ii), and (iii) without reference to the more explicit characterization (iv).¹

We begin with the proof that (iv) implies (i). Since the desired conclusion, (i), is preserved by coproducts, we assume, without loss of generality, that \mathcal{E} is the topos of continuous G -sets for a coherent topological group G . That \mathcal{E} is atomic is well known and follows (without any need for the coherence of G) either from the observation that its ‘constant sheaf’ functor, which gives each set the trivial action of G , is logical or from the equally easy observation that the transitive continuous G -sets, G/H for H an open subgroup of G , form a set of generators for \mathcal{E} and that the topology induced by the canonical topology of \mathcal{E} makes the full subcategory \mathcal{A} of transitive continuous G -sets an atomic site (because every G -equivariant map from G/H to G/K is surjective). It remains, therefore, to prove that \mathcal{E} is coherent. To do this, we shall use a slightly larger site of definition than \mathcal{A} , because \mathcal{A} is unlikely to have finite limits. Let \mathcal{B} be the full subcategory of \mathcal{E} consisting of the objects that contain only finitely many G -orbits, i.e. the closure of \mathcal{A} under finite coproducts. It is easy to see that any covering of an object of \mathcal{B} has a finite subcovering; indeed, if the object consists of n orbits, then the subcovering can always be taken to consist of at most n maps, one to cover each orbit. Since \mathcal{B} includes \mathcal{A} , it serves as a site of definition for \mathcal{E} , and to prove the coherence of \mathcal{E} it suffices to check that \mathcal{B} is closed under finite limits in \mathcal{E} . But \mathcal{B} is obviously closed under subobjects and obviously contains 1, so we need only check closure under binary products. Furthermore, since binary products distribute over coproducts, we

¹ In response to our announcement of the equivalence of (i) and (ii), J.M.E. Hyland and M. Barr provided different, purely topos-theoretic proofs of it.

need only check that $(G/H_1) \times (G/H_2)$ is in \mathcal{B} for all open subgroups H_1 and H_2 of G . Taking H to be $H_1 \cap H_2$, we observe that G/H maps onto both G/H_1 and G/H_2 (by $gH \mapsto gH_i$), so we need only check that $(G/H) \times (G/H)$ is in \mathcal{B} . But the G -orbits in $(G/H) \times (G/H)$ correspond bijectively to the H -orbits in G/H (by intersection with $\{H\} \times (G/H)$) which in turn correspond bijectively to the double cosets HgH . The assumption that G is coherent thus suffices to complete the proof that \mathcal{E} is coherent. (The terminology ‘coherent’ for groups is motivated by the fact that this property is equivalent to the coherence of the topos of continuous G -sets.)

We turn to the proof that (iii) implies (iv). Assume (iii), and let $\mathcal{F}_1, \dots, \mathcal{F}_n$ be the completions of \mathcal{F} ; n is finite by (b₁). In the syntactic site (\mathcal{C}', J') described in Section 1, let \mathcal{C}'' be the full subcategory consisting of those $\{\mathbf{x} \mid \phi(\mathbf{x})\}$ for which $\phi(\mathbf{x})$ is an atom in the Lindenbaum algebra of formulas with free variables among \mathbf{x} . (For the sake of brevity, we confuse a formula with its \mathcal{F} -provable-equivalence class.) For any object $\{\mathbf{y} \mid \psi(\mathbf{y})\}$ of \mathcal{C}' , we obtain a covering by objects of \mathcal{C}'' as follows. The finiteness of the Lindenbaum algebra (condition (b)) lets us express $\psi(\mathbf{y})$ as a finite disjunction of atoms, say $\bigvee_{i=1}^n \phi_i(\mathbf{y})$. By condition (a), we may take each ϕ_i to be geometric, so each $\{\mathbf{x} \mid \phi_i(\mathbf{x})\}$ is an object of \mathcal{C}'' . These objects clearly cover $\{\mathbf{y} \mid \psi(\mathbf{y})\}$ via the morphisms $[\mathbf{x} \mapsto \mathbf{y} \mid \phi_i(\mathbf{x}) \wedge \mathbf{x} = \mathbf{y}]$. By the comparison lemma [4, III. 4.1], the topos $\mathcal{E}(\mathcal{F})$ of sheaves on (\mathcal{C}', J') is also the topos of sheaves on (\mathcal{C}'', J'') , where J'' is the topology on \mathcal{C}'' induced by J' . For each atom $\phi(\mathbf{x})$, the formulas $\psi(\mathbf{x})$ that it \mathcal{F} -provably implies constitute a complete type, and the sentences that it \mathcal{F} -provably implies therefore constitute one of the completions \mathcal{F}_k of \mathcal{F} . This \mathcal{F}_k is the unique completion of \mathcal{F} with which $\phi(\mathbf{x})$ is consistent. If there is a morphism $[\mathbf{x} \mapsto \mathbf{y} \mid \theta(\mathbf{x}, \mathbf{y})]$ in \mathcal{C}'' from $\{\mathbf{x} \mid \phi(\mathbf{x})\}$ to $\{\mathbf{y} \mid \psi(\mathbf{y})\}$, then, since $\theta(\mathbf{x}, \mathbf{y})$ \mathcal{F} -provably implies $\phi(\mathbf{x})$ and $\psi(\mathbf{y})$, all three of these formulas are consistent with the same completion of \mathcal{F} . Thus, \mathcal{C}'' is the disjoint union of n full subcategories \mathcal{C}''_k , each consisting of the $\{\mathbf{x} \mid \phi(\mathbf{x})\}$ where $\phi(\mathbf{x})$ is consistent with a particular \mathcal{F}_k . It follows that \mathcal{E} is the coproduct of the topoi of sheaves on these components \mathcal{C}''_k (for the induced topologies). These components are just the sites (\mathcal{C}''_k, J''_k) associated to the theories \mathcal{F}_k , so we assume, without loss of generality, that \mathcal{F} itself is complete, i.e. that $n=1$.

In \mathcal{C}'' , each morphism $[\mathbf{x} \mapsto \mathbf{y} \mid \theta(\mathbf{x}, \mathbf{y})] : \{\mathbf{x} \mid \phi(\mathbf{x})\} \rightarrow \{\mathbf{y} \mid \psi(\mathbf{y})\}$ is a covering. To see this, we note that the consistent (with \mathcal{F}) formula $\phi(\mathbf{x})$ \mathcal{F} -provably implies $\exists \mathbf{y} (\theta(\mathbf{x}, \mathbf{y}) \wedge \psi(\mathbf{y}))$, by definition of morphism. So $\exists \mathbf{x} \theta(\mathbf{x}, \mathbf{y})$ is \mathcal{F} -consistent with $\psi(\mathbf{y})$. But $\psi(\mathbf{y})$, being at atom, \mathcal{F} -provably implies everything \mathcal{F} -consistent with it, so $\psi(\mathbf{y}) \rightarrow \exists \mathbf{x} \theta(\mathbf{x}, \mathbf{y})$ is \mathcal{F} -provable, as required. (Note that this shows that \mathcal{E} is atomic.)

To construct a group G as in (iv), we let M be a model of \mathcal{F} with the following homogeneity property (for all n): If \mathbf{a} and \mathbf{b} are two n -tuples from M that satisfy exactly the same formulas (equivalently: that satisfy the same atom of the Lindenbaum algebra for n variables) then M has an automorphism sending \mathbf{a} to \mathbf{b} . Such a model exists; if \mathcal{F} is countable its unique countable model will do, and in any case a special model will do [3, Theorem 5.1.17]. Let G be the automorphism group of M , and let the topology of G be defined by declaring a basis of neighborhoods of the

identity e to consist of the subgroups

$$H_a = \{g \in G \mid g \text{ fixes each element of } a\}$$

for finite tuples a from M .

Since every open subgroup includes some H_a , the objects G/H_a generate the topos \mathcal{F} of continuous G -sets. We shall show that \mathcal{E} and \mathcal{F} are equivalent by showing that the sites of definition, (\mathcal{C}, J) for \mathcal{E} and the full subcategory \mathcal{A} of objects G/H_a for \mathcal{F} , are equivalent. Since all morphisms in \mathcal{F} between objects of \mathcal{A} are epic, \mathcal{A} is an atomic site. Since (\mathcal{C}, J) is also an atomic site, we need only check that \mathcal{A} and \mathcal{C} are equivalent as categories. To do this, we define a functor F sending an arbitrary object $\{x \mid \phi(x)\}$ of \mathcal{C} to G/H_a , where a is some (selected) tuple in M satisfying $\phi(x)$. If $[x \rightarrow y \mid \theta(x, y)] : \{x \mid \phi(x)\} \rightarrow \{y \mid \psi(y)\}$ is a morphism in \mathcal{C} and if a and b are the selected solutions of $\phi(x)$ and $\psi(y)$, then F takes this morphism to the map of G -sets, $\alpha : G/H_a \rightarrow G/H_b$ defined as follows. From the definition of morphism in \mathcal{C} and the fact that M is a model of \mathcal{F} , it follows that there is a unique b' in M satisfying $\theta(a, b')$ and therefore also $\psi(b')$. From the homogeneity property of M , it follows that there exists $g \in G$ mapping b to b' . We then have $gH_b g^{-1} = H_{g(b)} = H_{b'}$ and, because b' is definable from a by θ , $H_a \subseteq H_{b'}$. We define α by $\alpha(qH_a) = qgH_b$ and leave to the reader the straightforward verification that this is well defined. We also leave to the reader the equally straightforward but more tedious verification that the F we have defined is a functor. To see that it is faithful, suppose that two morphisms, given by formulas θ^1 and θ^2 , lead to the same α , and let b^1 and b^2 , g^1 and g^2 be the corresponding b' and g as in the preceding discussion. Since $\alpha(H_a) = g^1 H_b = g^2 H_b$, we infer that $g^1 = g^2 \cdot h$ for some $h \in H_b$, which implies $b^1 = b^2$, so we may revert to the notation b' . Now we have both $\theta^1(a, b')$ and $\theta^2(a, b')$ holding in M , so $\exists y (\theta^1(x, y) \wedge \theta^2(x, y))$ is satisfied by a , hence \mathcal{F} -consistent with $\phi(x)$, hence \mathcal{F} -provable from $\phi(x)$ as $\phi(x)$ is an atom. It immediately follows that θ^1 and θ^2 define the same morphism. Thus, F is faithful.

To see that F is full, let any G -equivariant map $\alpha : G/H_a \rightarrow G/H_b$ be given. We attempt to reverse the steps in the definition of F in order to find a morphism $[x \rightarrow y \mid \theta(x, y)]$ that F maps to α . Choose g so that $gH_b = \alpha(H_a)$; then equivariance yields $qgH_b = \alpha(qH_a)$. Choose $b' = g(b)$. Choose $\theta(x, y)$ to be an atom satisfied in M by a, b' ; this is possible since the Lindenbaum algebra is finite. We need only check that θ defines a morphism from $\{x \mid \phi(x)\}$ to $\{y \mid \psi(y)\}$, i.e. that

$$\forall x \forall y (\theta(x, y) \rightarrow \phi(x) \wedge \psi(y)),$$

$$\forall x (\phi(x) \rightarrow \exists y \theta(x, y)),$$

$$\forall x \forall y \forall z (\theta(x, y) \wedge \theta(x, z) \rightarrow y = z)$$

are \mathcal{F} -provable. The first two are easy because, in each of these, the antecedent is an atom, so it suffices to show that the antecedent and consequent are \mathcal{F} -consistent with each other, and the consistency is clear since a and b' satisfy these clauses. The third sequent is a bit harder, but we note that it suffices to prove that it holds in M

(since \mathcal{A} is complete), and for this it suffices to prove that, in M

$$\forall y \forall z (\theta(\mathbf{a}, y) \wedge \theta(\mathbf{a}, z) \rightarrow y = z). \quad (8)$$

Indeed, any x satisfying $\theta(x, y) \wedge \theta(x, z)$ would satisfy $\phi(x)$ and would therefore, as $\phi(x)$ is an atom, satisfy exactly the same formulas as \mathbf{a} . To establish (8), suppose that there were a \mathbf{b}'' , distinct from \mathbf{b}' , satisfying $\theta(\mathbf{a}, \mathbf{b}'')$. Since the tuples $(\mathbf{a}, \mathbf{b}')$ and $(\mathbf{a}, \mathbf{b}'')$ satisfy the same atom θ , there is an automorphism $z \in G$ sending the first to the second. Thus, $z \in H_{\mathbf{a}}$ but $z \notin H_{\mathbf{b}'}$. Since $z \in H_{\mathbf{a}}$, z stabilizes the element $\alpha(H_{\mathbf{a}}) = gH_{\mathbf{b}}$ in $G/H_{\mathbf{b}}$. This means $z \in gH_{\mathbf{b}}g^{-1} = H_{\mathbf{b}'}$, a contradiction. This completes the proof that F is full.

Finally, to see that F is essentially surjective on objects, consider any object $G/H_{\mathbf{c}}$ of \mathcal{A} . Let $\phi(x)$ be an atom satisfied in M by \mathbf{c} , and let \mathbf{a} be the selected solution of $\phi(x)$, so F takes $\{x \mid \phi(x)\}$ to $G/H_{\mathbf{a}}$. The homogeneity property of M provides a $g \in G$ that sends \mathbf{a} to \mathbf{c} , so $H_{\mathbf{c}} = gH_{\mathbf{a}}g^{-1}$. It is easy to check that $G/H_{\mathbf{c}}$ is isomorphic to $G/H_{\mathbf{a}}$ by the map sending $qH_{\mathbf{c}}$ to $qgH_{\mathbf{a}}$. This completes the proof that F is an equivalence and thus also the proof that \mathcal{A} is equivalent to the topos of continuous G -sets.

To complete the proof of Theorem 2, we verify that G is coherent. Since each open subgroup includes $H_{\mathbf{a}}$ for some finite sequence \mathbf{a} in M , we need only check that, for each fixed n -tuple \mathbf{a} , the number of double cosets $H_{\mathbf{a}}gH_{\mathbf{a}}$ is finite. We assert that such a double coset is completely determined by an atom $\phi(x, y)$ (in $2n$ variables) satisfied in M by \mathbf{a} , $g(\mathbf{a})$. This will finish the proof since there are only finitely many such atoms. To prove the assertion, we consider two elements $g_1, g_2 \in G$ such that $\mathbf{a}, g_1(\mathbf{a})$ and $\mathbf{a}, g_2(\mathbf{a})$ satisfy the same atom. Then there is an automorphism $h \in G$ sending the first of these sequences to the second. Since it fixes \mathbf{a} , h belongs to $H_{\mathbf{a}}$. Also, the automorphism $h' = g_1^{-1}h^{-1}g_2$ fixes \mathbf{a} , hence belongs to $H_{\mathbf{a}}$. But $g_2 = hg_1h'$, so $H_{\mathbf{a}}g_1H_{\mathbf{a}} = H_{\mathbf{a}}g_2H_{\mathbf{a}}$. This completes the proof of Theorem 2.

References

- [1] M. Barr and R. Diaconescu, Atomic toposes, *J. Pure and Appl. Algebra* 17 (1980) 1-24.
- [2] A. Blass and A. Ščedrov, Classifying topoi and finite forcing, *J. Pure Appl. Algebra* 28 (1983), to appear.
- [3] C.C. Chang and H.J. Keisler, *Model Theory* (North-Holland, Amsterdam, 1973).
- [4] A. Grothendieck and J.L. Verdier, *Séminaire de Géométrie Algébrique IV, Tome 1 Lecture Notes in Mathematics* No. 269 (Springer, Berlin-New York, 1972).
- [5] P. Johnstone, *Topos Theory* (Academic Press, New York, 1977).
- [6] A. Ščedrov, Sheaves and forcing and their metamathematical applications, Thesis, SUNY at Buffalo (1981), to appear in *Memoirs Amer. Math. Soc.*