The Logic in Computer Science Column

by

Yuri Gurevich

Microsoft Research
One Microsoft Way, Redmond WA 98052, USA
gurevich@microsoft.com

Background of Computation

Andreas Blass[*] Yuri Gurevich

Abstract

In a computational process, certain entities (for example sets or arrays) and operations on them may be automatically available, for example by being provided by the programming language. We define background classes to formalize this idea, and we study some of their basic properties. The present notion of background class is more general than the one we introduced in an earlier paper, and it thereby corrects one of the examples in that paper. The greater generality requires a non-trivial notion of equivalence of background classes, which we explain and use. Roughly speaking, a background class assigns to each set (of atoms) a structure (for example of sets or arrays or combinations of these and similar entities), and it assigns to each embedding of one set of atoms into another a standard embedding between the associated background structures. We discuss several, frequently useful, properties that background classes may have, for example that each element of a background structure depends (in some sense) on only finitely many atoms, or that there are explicit operations by which all elements of background structures can be produced from atoms.

[*]Mathematics Department, University of Michigan, Ann Arbor, MI 48109–1043, U.S.A., ablass@umich.edu

1 Introduction

Background structures were introduced in [2] to serve two purposes. One, which was of primary importance to us in writing the present paper, is to describe the world in which an algorithm operates. The other, which provided some of our motivation for [2], is to facilitate dealing with new elements imported into the underlying structures of computations. We describe these two purposes in turn.

Programming languages provide certain sorts of constructions automatically, while others have to be programmed separately. For example, C provides arrays (of elements of any given type) automatically, but a programmer who wants to use multisets will have to represent them somehow and code any needed operations on them. A C compiler "knows" about arrays but not about multisets. In AsmL [1], sets and multisets are automatically available, as are maps and sequences. Arrays, however, are not first-class citizens of AsmL. On the abstraction level of AsmL they can be represented by sequences or maps. (An implementation can make arrays available, possibly in a limited way.)

We refer to the automatically available notions as the background of a programming language. This would include not only the entities, like sets, but also operations on them. Different programming languages could provide sets but with different pre-defined operations; these would constitute different backgrounds.

Here are some other examples: A database system's background would include numbers and operations like average and median. Metafinite structures [3] usually also involve a background of numbers. Even such a low-level computing system as the Turing machine has a background containing the tape cells and the operations of moving to the left and right.

Remark 1. Algorithms also have backgrounds, consisting of whatever is presupposed by the algorithm, i.e., whatever would have to be supplied by an operating system or by a programming language in order to actually run the algorithm. In general, if one lowers the level of abstraction of an algorithm, by making more of the implementation explicit, then one decreases the required background.

This notion of background occurs also in situations other than programming languages. A mathematical text, for example, presupposes that certain notions are already known to the reader; these can be considered the background of the text. Even in everyday life, conversations usually depend on some shared background information. The computational situation, however, differs from these others in that it requires the background to be completely explicit. In the case of programming languages, it must be built into the compiler. In mathematics, one is usually content with the knowledge that the background could, if necessary, be made explicit, for example by formalizing proofs in a set-theoretic foundational system. And in everyday life, the background is usually even less precisely delimited.

A secondary purpose of backgrounds was to solve (in [2]) a problem in the description of computations whose states are first-order structures. In the theory of abstract state machines [5], it is convenient and customary to assume that the underlying set of a structure does not change as the algorithm proceeds. To reconcile this assumption with the frequent need to introduce a new vertex into a graph, to create a new object of a given class, or to append a new cell to a finite Turing machine tape, one assumes that the structures contain infinitely many reserve elements. These are not involved in the computation but can be imported, i.e., removed from the reserve and allowed to participate in the computation, when the need arises. Intuitively, the reserve is an entirely unstructured part of the base set. But this point of view causes difficulties if the algorithm also uses, for example, sets or arrays of elements. Whenever an element a is imported from the reserve, it would be necessary to import many additional elements to serve as sets, arrays, etc., involving a, and it would be necessary to define the relevant relations and operations on these elements. It is more reasonable to assume that these sets, arrays, etc., were already present, even while a was still in the reserve. But of course, such an assumption must be made precise — what exactly is attached to the reserve elements? Furthermore, these additional entities should not represent additional internal structure on the reserve; for example, they should not allow a definition of a linear ordering of the reserve elements.

These considerations were a large part of the motivation for [2]. Background classes were introduced there to formalize the idea of things that can be built on top of a set (e.g., the reserve) without thereby introducing any additional structure on the set itself. The starting set here is usually called a set of atoms, and a background structure is built over a set of atoms.

The purpose of the present paper is to extend and generalize the theory of background structures developed in the first part of [2]. The generalization also serves to correct an error in one of the examples in [2].

After a review of preliminary information in Section 2, we describe in Section 3 the error in one of the examples of [2, Section 5.1] and we describe our approach to correcting it. This approach is formalized in Section 4 where we define background classes by means of axioms. Then, in Section 5 we discuss the natural notion of equivalence of background classes and show that a background class can be described, up to equivalence, by a relatively limited amount of information. In the next two sections, we discuss two additional properties enjoyed by some but not all background classes. One property (Section 6) is that each element of a background structure arises from finitely many atoms. The other (Section 7) is that each element has a smallest source from which it arises. (In [2], this "smallest source" requirement was included in the definition of background class; here we drop the requirement, but the property in question remains an interesting one.) In Section 8, we list several ways in which the elements of a background structure could be said to be constructed out of the atoms, and we give examples to show how these interpretations of "constructed" differ. We also briefly discuss the prospects for making a background class more constructive by adding additional operations to serve as constructors.

2 Preliminaries

We adopt the following conventions from [2] about vocabularies and structures.

A *vocabulary* Υ is a collection of function symbols of prescribed arities; some of these symbols may be marked as *relational*. Every vocabulary contains the *logic symbols*: binary =; nullary `true`, `false`, and `undef`; unary `Boole`; and the usual propositional connectives. All of these except `undef` are relational. In addition, every vocabulary contains the unary relational symbol `Atomic`. The logical symbols and `Atomic` are called the *obligatory symbols*, because every vocabulary must contain them.

An Υ-*structure* X consists of a nonempty *base set*, usually denoted by the same symbol X, with interpretations of all the function symbols of Υ as functions. The interpretation of a symbol f is officially written f^X, but we often omit the superscript. It is required that the interpretation of `true` be different from those of `false` and `undef`, that the interpretations of relational symbols take values in {`true`, `false`}, that `Boole` map `true` and `false` to `true` and everything else to `false`, that equality have its standard interpretation, and that the propositional connectives have their standard interpretations (with value `false` if any argument is not `true` or `false`).

In any structure, the elements that satisfy `Atomic` (i.e., that are mapped to `true` by the interpretation of `Atomic`) are called the *atoms* of the structure. For consistency with [2], we use the notation Atoms(X) as a synonym for $\texttt{Atomic}^X$, the set of atoms of the structure X. The interpretations in any structure of the nullary obligatory symbols `true`, `false`, and `undef` are called the *obligatory elements* of the structure.

It is occasionally convenient to work in the foundational framework of a set theory that allows atoms (also called urelements). The intuitive picture of such a framework is a universe containing atoms (which are not sets), sets of atoms, sets of (atoms and sets of atoms), etc. That is, one begins with atoms and then repeatedly forms all sets of previously formed things (atoms and sets).

In such a framework, it is helpful to think of Υ-structures in which the predicate `Atomic` applies to exactly those elements that are atoms in the foundational framework. Then the foundational and the structure-based meanings of "atom" coincide. But it is important to remember that not all structures have "standard atoms" in this sense. We shall usually be dealing with classes of structures closed under isomorphism; in such a situation, we may describe the class by telling which structures with standard atoms it contains, and it is then to be understood that it also contains the isomorphic copies of these structures.

Analogous comments apply when a vocabulary contains symbols for other notions used in foundations, such as the notion of set. For such a vocabulary, we may specify a class of structures by describing structures in which the interpretation of "set" (or similar notions) coincides with the meaning given by the foundational system, but we must remember that isomorphic copies of the structures are also included.

When we work with structures whose interpretations of "atom", "set", "member", or other such notions disagree with those of the foundational system, we intend these words to be understood in the sense of the structure unless we explicitly indicate the contrary.

3 The Problem and Two Solutions

In Section 5 of [2], we gave several examples of background classes, the first of which was a set background that contained, for each set U, the structure of hereditarily finite sets over the set of atoms U. The only non-obligatory symbol in the vocabulary was the symbol $\in$ for membership. (This example was accompanied by variations with a larger vocabulary, but the main version, with only $\in$, is the one that caused a problem.)

Tatiana Yavorskaya pointed out that this example is not really a background class because it does not satisfy the requirement BC2 of the definition of background class [2, Definition 4.1]. That requirement says that, given two background structures X and Y, and given an embedding of sets $\zeta : \text{Atoms}(X) \to \text{Atoms}(Y)$, there is a unique extension of ζ to an embedding $X \to Y$. In the example under consideration, the uniqueness part of this requirement fails. Consider, for example, the structures X and Y of hereditarily finite sets over $\{1, 2\}$ and $\{1, 2, 3\}$, respectively, and the embedding $\zeta : \{1, 2\} \to \{1, 2, 3\}$ sending each of 1 and 2 to itself, i.e., the inclusion map. There is an obvious extension of ζ to an embedding $X \to Y$, namely a bigger inclusion map — send every set in X to itself considered as an element of Y. But there are also other extensions that make use of the additional atom 3. For example, one could define an embedding η to send the empty set $\varnothing$ to $\{3\}$ and then map all other sets in the natural way, $\eta(s) = \{\eta(x) : x \in s\}$. Or one could throw in 3 at every opportunity, defining $\eta(s) = \{\eta(x) : x \in s\} \cup \{3\}$. There are many more such possibilities. (The two variations suggested at the end of this example in [2] avoid the problem because the extraneous embeddings would not preserve the additional basic functions.)

One possible solution to the problem is simply to give up. Drop the set-background from the list of examples. If some version of a set-background is needed, equip it with not only $\in$ but enough additional structure to block unwanted embeddings. Perhaps even insist that all background structures be explicitly atom-generated, as defined in [2]; that is, no background structure can have a proper substructure (background or not) that contains all the atoms. This assumption clearly implies the uniqueness part of requirement BC2, since, if we had two embeddings $\eta, \eta' : X \to Y$ extending ζ then $\{x \in X : \eta(x) = \eta'(x)\}$ would be a substructure of X containing all the atoms.

In fact, for some purposes it is useful for a background class to have even stronger properties than being explicitly atom-generated. For example, if each background structure is obtained as an Herbrand universe of terms built from the atoms and (the interpretations of) certain function symbols, then these terms can provide a convenient representation of background entities, both for reasoning about computation and for internal use within the computation. We discuss such properties of background classes in Section 8.

On the other hand, this solution may force us to an undesirably low level of abstraction. Specifically, it requires us to choose the additional functions to add in order to generate all the background elements or at least to block the unwanted embeddings. There are algorithms whose natural formulation uses sets but does not rely on any particular additional functions to generate them. An example is given by a familiar algorithm for deciding whether two vertices of a graph, say s and t, are in the same component: Initialize a variable set V to $\{s\}$ and then repeatedly adjoin to V all neighbors of vertices already in V, checking whether t has entered V by the time V stops growing. As indicated in [2], there are several reasonable options to generate a set-background structure from its atoms. No one of these, however, is naturally called for by the algorithm just described. Furthermore, intuitively, set theory should have just the one primitive concept $\in$, and enlarging the vocabulary represents a departure from the basic intuition of the set-background. The enlarged vocabulary might be viewed as providing a particular implementation of the original, more abstract notion of background.

We therefore present in this paper a different solution of the problem, a solution that allows us to retain the simple vocabulary of just $\in$ for the set-background. Of course, this approach must, in view of Yavorskaya's observation, alter the definition of background class.

The main idea of the alteration is to specify, along with the structures in the background class, certain embeddings between them, which we call the standard embeddings. Thus, in the example above of the set-background structures X and Y over the sets $\{1, 2\}$ and $\{1, 2, 3\}$, the inclusion map of X into Y would be a standard embedding but the other embeddings, using the new atom 3, would not be standard. (The idea of specifying the desired embeddings along with the desired structures was suggested by the theory of abstract elementary classes. Such a class consists of a class of structures together with a specified notion of so-called strong substructures, intended to play the role of elementary substructures in traditional model theory. For an introduction to abstract elementary classes, see [4].)

It should be pointed out that, although this difficulty affected only one example in [2], there are many other, similar situations where the same difficulty would arise. Notice first that the difficulty did not depend at all on having sets of higher rank, i.e., sets with other sets as elements. Even with only sets of atoms, Yavorskaya's observation would still apply. Another very similar example is a background of multisets and numbers, with a "multiplicity of membership" function (taking an element x and a multiset y to the multiplicity of x in y) along with arithmetic operations on numbers. Another has words formed from atoms, with a relation "contiguous subword". In these examples, as in the set example above, the standard embeddings would be those that simply don't use the new elements (the 3 in the preceding discussion), but there are non-standard embeddings that insert 3 into the images of multisets or words. An example with a slightly different flavor has words and the relation "not-necessarily-contiguous subword". Again, an embedding can introduce extraneous atoms into words, but some care is needed in order to preserve this more general subword relation. An approach that works is to insert a particular new atom between every two consecutive atoms in a word. (One can also insert the new atom at the beginning or end or both, but if one refrains from doing so then these embeddings also preserve the "first letter" and "last letter" relations.) Thus, our modification of the axioms from [2] addresses a problem more extensive than the one example that motivated it.

Although the introduction of standard embeddings is the only essential change here from the framework in [2], we take this opportunity to reorganize the axioms somewhat. In [2], we included, in the definition of background class, the requirement that, for each element x of a background structure, there is a smallest background substructure containing x. On the other hand, we did not include in the definition a requirement that every element of a background structure "come from" a finite number of atoms; we called background structures with this additional property "finitary". We now believe that these conventions misrepresent the relative importance of these two requirements. Accordingly, we shall introduce a basic set of axioms, incorporating neither of these two requirements but only the bare minimum of what seems to be needed for a sensible notion of background. These basic axioms correspond to BC0, BC1, and BC2 of [2], modified to take standard embeddings into account. All additional properties will be dealt with separately.

4 Basic Axioms and Their Consequences

In this section, we introduce the basic axioms for background classes, present some examples, discuss the relationship with the axioms in [2], and deduce some immediate consequences of the axioms.

Definition 2. A *background class* consists of

- a vocabulary Υ that contains all the obligatory symbols,
- a class $\mathcal{B}$ of Υ-structures, called the *background* structures, and
- a class of embeddings (i.e., structure-preserving one-to-one functions) between background structures, called the *standard embeddings*,

subject to the following axioms.

1. The identity embedding from any background structure to itself is standard.
2. The composite $\zeta \circ \eta$ of two standard embeddings is standard.
3. Every structure isomorphic to a background structure is also a background structure.
4. Every set is $\mathrm{Atoms}(X)$ for some background structure X.
5. If X and Y are background structures, then every embedding (of sets) $\mathrm{Atoms}(X) \to \mathrm{Atoms}(Y)$ has a unique extension to a standard embedding $X \to Y$.

Definition 3. In the situation of Axiom 5, if $\zeta : \mathrm{Atoms}(X) \to \mathrm{Atoms}(Y)$ is an embedding, we write $\hat{\zeta}$ for the standard embedding $X \to Y$ that extends it. (Strictly speaking, the notation $\hat{\zeta}$ should also mention X and Y, but they will always be clear from the context.)

Note that the definition of background class requires standard embeddings to be embeddings in the usual sense — one-to-one maps of the base sets of structures, respecting the interpretations of the function symbols of Υ. In the special case where a standard embedding is an isomorphism or an automorphism, we shall speak of a *standard isomorphism* or *standard automorphism*.

Remark 4. Following [2], our semantic conventions in Section 2 did not specify whether the values of `false` and `undef` should be equal or not, though `true` must be different from them. We shall occasionally use the convention that `false` $\neq$ `undef`, but everything we do under this convention can be easily modified to work under the opposite convention. In this connection, it is worth noting that, within any background class, all the background structures agree as to whether

$\texttt{false} = \texttt{undef}$ or not. Indeed, whatever happens in a background structure with the empty set of atoms will be preserved by the standard embeddings into arbitrary background structures.

Remark 5. The non-obligatory symbols of the vocabulary Υ and their interpretations in background structures play a lesser role under our current definition of background classes than they did in [2]. They determine which maps are embeddings, but most of the role of embeddings in [2] is now played by the standard embeddings. These are not determined but only somewhat constrained by the non-obligatory functions, because all standard embeddings are embeddings but not conversely. Indeed, in a part of our analysis of background structures, the non-obligatory functions will be relegated to the sidelines. In applications, though, the non-obligatory functions are very important, because they determine what operations are available to an algorithm working in a given background.

Axioms 1 and 2 imply that $\widehat{\mathrm{Id}}_{\mathrm{Atoms}(X)} = \mathrm{Id}_X$ for all background structures X and that $\widehat{\zeta \circ \eta} = \hat{\zeta} \circ \hat{\eta}$ whenever the compositions make sense. This implies the following proposition, just as in [2, Lemma 4.4].

Proposition 6. *If $\zeta : Atoms(X) \to Atoms(Y)$ is a bijection, then $\hat{\zeta} : X \to Y$ is an isomorphism.*

Proof. Being a bijection, ζ has an inverse function η and

$$\hat{\zeta} \circ \hat{\eta} = \widehat{\zeta \circ \eta} = \widehat{\mathrm{Id}}_{\mathrm{Atoms}(Y)} = \mathrm{Id}_Y.$$

Similarly $\hat{\eta} \circ \hat{\zeta} = \mathrm{Id}_X$. So $\hat{\zeta}$ and $\hat{\eta}$ are inverse isomorphisms. □

To avoid possible confusion, note that Axiom 1 refers only to identity maps from a structure into itself, not to inclusion maps from a substructure into a superstructure.

Axiom 3 requires that if X is a background structure and $i : X \to Y$ is an isomorphism, then Y is also a background structure, but it does not require i to be a standard embedding. In fact, we have the following characterization of the background classes for which all isomorphisms between background structures are standard.

Proposition 7. *For any background class, the following are equivalent:*

1. *Every isomorphism between background structures is standard.*

2. *Whenever an automorphism i of a background structure X leaves all the atoms of X fixed, it leaves all the elements of X fixed.*

Proof. Assume statement (1). Then in particular, every automorphism i of any background structure X is standard. If i leaves all the atoms of X fixed, then both it and the identity map (which is standard by Axiom 1) are standard embeddings of X into itself, extending the same embedding of the atoms, namely the identity. So i must be the identity map of X, because of the uniqueness requirement in Axiom 5.

Conversely, assume (2) and let $i : X \to Y$ be an isomorphism of background structures. By Axiom 5, there is a standard embedding $j : X \to Y$ that agrees with i on atoms. As observed above, since i is a bijection on atoms, j is an isomorphism. Then $i^{-1} \circ j$ is an automorphism of X fixing all the atoms. By our assumption (2), it is the identity automorphism, which means that $i = j$, and thus i is standard. □

Every background class in the sense of [2] becomes a background class in the sense of the present definition if we call all embeddings standard. Indeed, with this notion of standard, our Axioms 1 and 2 are trivially satisfied, and our other three axioms become Axioms BC0, BC1, and BC2 of [2].

The set background in [2, Section 5.1], which fails to satisfy the uniqueness clause of BC2, nevertheless satisfies our new definition of background class, provided the standard embeddings are chosen to be the ones that were intended in [2]. These are the embeddings from the hereditarily finite sets over A to the hereditarily finite sets over B that one obtains by starting with an arbitrary embedding $\zeta : A \to B$ and extending it to act on sets by recursively defining $\hat{\zeta}(x) = \{\hat{\zeta}(y) : y \in x\}$. (The set notation here refers to sets and membership as interpreted in the structures.) Note that these standard embeddings are exactly the ones that preserve the additional functions used in the two alternative versions of the set background described at the end of Section 5.1 of [2].

Axiom 4 requires the existence of a background structure with any prescribed set of atoms. The following proposition describes the sense in which this background structure is unique.

Proposition 8. *If X and X' are background structures with the same set of atoms, then they are isomorphic, by an isomorphism that leaves every atom fixed. Furthermore, exactly one such isomorphism is standard.*

Proof. By Axiom 5, the identity map of the atoms extends to a unique standard embedding $X \to X'$, and by Proposition 6 this is an isomorphism. □

There are, in the general context of background classes, two reasonable but inequivalent ways to formalize the notion of "the set of atoms involved in x" where x is an element of a background structure. We give these two formalizations, which we call the source and the core of x, and we point out some immediate consequences of the definition as well as one not so immediate consequence. Both sources and cores will play a role in subsequent sections.[1]

Definition 9. Let x be an element of a background structure X, and let $S \subseteq \text{Atoms}(X)$. Then S is a *source* for x if x is in the image of the standard embedding $\hat{\imath}$ that extends the inclusion map $i : S \to \text{Atoms}(X)$.

There is a harmless ambiguity in this definition, namely that we did not specify which background structure, with S as its set of atoms, is to serve as the domain of $\hat{\imath}$. It follows easily from Proposition 8 and Axiom 2 that all choices here produce the same image of $\hat{\imath}$, so there is no need to specify a choice.

Definition 10. Let x be an element of a background structure X, and let $C \subseteq \text{Atoms}(X)$. Then C is a *core* of x if, for all background structures Y and all standard embeddings $\zeta, \eta : X \to Y$, if ζ and η agree on C, then $\zeta(x) = \eta(x)$.

The following proposition lists some immediate consequences of the definitions.

Proposition 11. *For any background structure X and any $x \in X$:*

1. *If S is a source for x and $S \subseteq T \subseteq \text{Atoms}(X)$, then T is also a source for X.*
2. *If S is a core for x and $S \subseteq T \subseteq \text{Atoms}(X)$, then T is also a core for X.*
3. *Every source for x is also a core for x.*
4. *If $\zeta : X \to Y$ is a standard embedding and $S \subseteq \text{Atoms}(X)$ is a source for x in X, then $\zeta[S]$ is a source for $\zeta(x)$ in Y.*
5. *If $\zeta : X \to Y$ is a standard embedding and $C \subseteq \text{Atoms}(X)$ is a core for x in X, then $\zeta[C]$ is a core for $\zeta(x)$ in Y.*
6. *If $\zeta : X \to Y$ is a standard embedding, $x \in X$ and $S \subseteq \zeta[\text{Atoms}(X)]$ is a source for $\zeta(x)$, then $\zeta^{-1}[S]$ is a source for x.*
7. *If $\zeta : X \to Y$ is a standard embedding, $x \in X$, and $C \subseteq \text{Atoms}(Y)$ is a core for $\zeta(x)$ in Y, then $\zeta^{-1}[C]$ is a core for x in X.*
8. *Every set of atoms is a source, and therefore also a core, for the values of any constant symbols in Υ. In particular, the empty set is a source for such elements, including the obligatory elements* `true`, `false`, *and* `undef`.

[1] In [2], we used the word "support" for essentially what is here called "source". It is, however, standard in the set-theoretic literature to use "support" for essentially what is here called "core". It therefore seems appropriate to avoid the word "support" when, as here, both of its past meanings are under discussion.

9. *Each atom a has $\{a\}$ as is smallest source and smallest core.*

Proof. For assertion (1), observe that the inclusion map $i : S \to \mathrm{Atoms}(X)$ is the composite $i = j \circ k$ of the inclusion maps $k : S \to T$ and $j : T \to \mathrm{Atoms}(X)$. Therefore $\hat{\imath} = \hat{\jmath} \circ \hat{k}$ and the image of $\hat{\imath}$ is included in that of $\hat{\jmath}$.

Assertion (2) follows immediately from the fact that any two embeddings $X \to Y$ that agree on T also agree on its subset S.

For assertion (3), let $i : S \to \mathrm{Atoms}(X)$ be the inclusion map of a source S for x, so $x = \hat{\imath}(s)$ for some s in a background structure with atoms S, and assume $\zeta, \eta : X \to Y$ are standard embeddings that agree on S. Then $\zeta \circ i = \eta \circ i$ and so $\zeta \circ \hat{\imath} = \eta \circ \hat{\imath}$. Apply both sides of this equation to s to get $\zeta(x) = \eta(x)$.

For assertion (4), let U and V be background structures with S and $\zeta[S]$ as their respective sets of atoms, let $\eta : U \to V$ be the standard embedding extending $\zeta \restriction S$, and let $i : S \to \mathrm{Atoms}(X)$ and $j : \zeta[S] \to \mathrm{Atoms}(Y)$ be the inclusion maps. Notice that $\zeta \circ \hat{\imath} = \hat{\jmath} \circ \eta$. Since S is a source for x, fix some $u \in U$ with $\hat{\imath}(u) = x$, and compute that

$$\zeta(x) = \zeta(\hat{\imath}(u)) = \hat{\jmath}(\eta(u)),$$

which is in the image of $\hat{\jmath}$ as required.

For assertion (5), suppose $\eta, \theta : Y \to Z$ are two standard embeddings that agree on $\zeta[C]$. Then $\eta \circ \zeta$ and $\theta \circ \zeta$ agree on C. So they send x to the same value. That means $\eta(\zeta(x)) = \theta(\zeta(x))$, as required.

For assertion (6), assume the hypotheses, and let $i : S \to \mathrm{Atoms}(Y)$ and $j : \zeta^{-1}[S] \to \mathrm{Atoms}(X)$ be the inclusion maps. Because $S \subseteq \zeta[\mathrm{Atoms}(X)]$, we have $i = \zeta \circ j \circ \eta$ where η is the restriction of ζ^{-1} to S. Fix background structures U and V with $\mathrm{Atoms}(U) = S$ and $\mathrm{Atoms}(V) = \zeta^{-1}[S]$. Because S is a source for $\zeta(x)$, there is $u \in U$ with $\hat{\imath}(u) = \zeta(x)$. Thus,

$$\zeta(x) = \hat{\imath}(u) = \zeta(\hat{\jmath}(\hat{\eta}(u))).$$

Since ζ is one-to-one, we conclude that $x = \hat{\jmath}(\hat{\eta}(u))$. So x is in the range of $\hat{\jmath}$, which means that $\zeta^{-1}[S]$ is a source for x.

For assertion (7), assume the hypotheses, and let $\eta, \theta : X \to Z$ be standard embeddings that agree on $\zeta^{-1}[C]$. Temporarily suppose $\mathrm{Atoms}(Z)$ is large enough so that $\mathrm{Atoms}(Y) - \zeta[\mathrm{Atoms}(X)]$ admits a one-to-one map λ into $\mathrm{Atoms}(Z) - \eta[\mathrm{Atoms}(X)] - \theta[\mathrm{Atoms}(X)]$. Then we can define a one-to-one map from Y into Z agreeing with $\eta \circ \zeta^{-1}$ on $\zeta[\mathrm{Atoms}(X)]$ and with λ on the rest of $\mathrm{Atoms}(Y)$. Let $\eta' : Y \to Z$ be a standard embedding extending this map. Define $\theta' : Y \to Z$ analogously. Then η' and θ' agree on C, so $\eta'(\zeta(x)) = \theta'(\zeta(x))$. But $\eta' \circ \zeta = \eta$ and $\theta' \circ \zeta = \theta$, so $\eta(x) = \theta(x)$, as required.

It remains to eliminate the assumption that $\mathrm{Atoms}(Z)$ is large enough. If it isn't, then let Z' be any background structure whose $\mathrm{Atoms}(Z')$ is a large enough superset of Atoms(Z). Let $\hat{\imath} : Z \to Z'$ be the standard embedding that extends the inclusion map on the atoms. Then the preceding argument applies to $\hat{\imath} \circ \eta$ and $\hat{\imath} \circ \theta$, so these map x to the same element. As $\hat{\imath}$ is one-to-one, it follows that $\eta(x) = \theta(x)$.

For assertion (8) just notice that any constant symbol has some value in a background structure over the empty set of atoms, and this value is sent, by standard embeddings (extending the embeddings of $\varnothing$ into arbitrary sets), to the values of that constant symbol in all other background structures. So $\varnothing$ is a source of each such value, and therefore so are all sets of atoms.

For assertion (9), it is immediate from the definition that $\{a\}$ is a source and therefore a core for a. That it is the smallest follows from the fact that a can be sent to any atom in any background structure by some standard embedding ζ. Applying this to a background structure with at least two atoms, we see that $\zeta(a)$ cannot be determined by $\zeta \upharpoonright \varnothing$, so $\varnothing$ is not a core, and therefore not a source for a. Thus, $\{a\}$ is the smallest source and the smallest core for $\{a\}$. □

The converse of part (3) of the proposition is false in general. The simplest counterexample is the following.

Example 12. Define a background class by having each background structure consist of the atoms, the obligatory elements, and one other element provided there is at least one atom, but no other element when there are no atoms. (All embeddings are standard.) Then the one non-obligatory non-atom has the empty set as a core but not as a source.

Other examples are easily constructed using the same idea. For example, modify the set background by omitting all the sets when there are fewer than 17 atoms.

There is an important difference between the two ζ^{-1} parts of Proposition 11. In part (6), S is required to be included in the range of ζ. In part (7), there is no such requirement on C. The requirement cannot be removed from part (6). Indeed, in Example 12, if $\zeta : X \to Y$ is not surjective on atoms, then letting a be any element of Atoms(Y) $- \zeta$[Atoms(X)] and letting x be the non-obligatory non-atom of X, we have that $\{a\}$ is a source for $\zeta(x)$ but $\zeta^{-1}(\{a\}) = \varnothing$ is not a source for x.

Proposition 13. *Let X be a background structure, x an element of X, and C_1 and C_2 two cores for x. Then $C_1 \cap C_2$ is also a core for x.*

Proof. Let $\hat{\zeta}, \hat{\eta} : X \to Y$ be two standard embeddings that agree on $C_1 \cap C_2$. (The use of the “hat” notation serves to provide short names ζ and η for the restrictions of the embeddings to atoms.) We shall first prove that $\hat{\zeta}(x) = \hat{\eta}(x)$ under the assumption that Y is large enough, meaning that there are at least |Atoms(X)| atoms in Atoms(Y) $- \zeta$(Atoms(X)) $- \eta$(Atoms(X)), i.e., atoms of Y not used by either of the given embeddings. Afterward, we shall eliminate this additional assumption.

Because Y is large enough, there is an embedding $\zeta' : \text{Atoms}(X) \to \text{Atoms}(Y)$ that agrees with ζ on C_1 but maps the rest of Atoms(X) to atoms not used by ζ and η. Because ζ' agrees with ζ on a core C_1 of x, we have $\hat{\zeta}'(x) = \hat{\zeta}(x)$.

Obtain ζ'' from ζ' by modifying ζ' on $C_1 - C_2$ to agree with η there. This ζ'' is an embedding, i.e., it is one-to-one. Indeed, the modified values cannot clash with the values on $\text{Atoms}(X) - C_1$ because we arranged for ζ' to take values on $\text{Atoms}(X) - C_1$ that are not in the range of η. Nor can the modified values clash with the values on $C_1 \cap C_2$ because ζ', ζ, and η all agree there.

Because ζ'' agrees with ζ' on C_2, we have $\hat{\zeta}''(x) = \hat{\zeta}'(x)$.

Finally, because ζ agrees with η on $C_1 \cap C_2$, so does ζ'' and, as a result, ζ'' agrees with η on C_1. Therefore,

$$\hat{\eta}(x) = \hat{\zeta}''(x) = \hat{\zeta}'(x) = \hat{\zeta}(x),$$

as required. This completes the proof when Y is large enough.

It remains to handle the possibility that Y is not large enough. In this case, let i be an embedding of Y into a set Z that is large enough, with respect to the embeddings $i \circ \zeta$ and $i \circ \eta$. Since these embeddings agree on $C_1 \cap C_2$, the preceding argument applies and gives

$$\hat{\imath} \circ \hat{\zeta}(x) = \widehat{i \circ \zeta}(x) = \widehat{i \circ \eta}(x) = \hat{\imath} \circ \hat{\eta}(x).$$

Since $\hat{\imath}$ is one-to-one, the desired conclusion, $\hat{\zeta}(x) = \hat{\eta}(x)$, follows. □

The analogous statement for sources fails; the intersection of two sources of x need not be a source for x. In Example 12, the one non-obligatory non-atom has all nonempty sets of atoms as its sources, but the intersection of two of these may be empty and thus not a source.

Remark 14. The proof of Proposition 13 used the fact that, given ζ, η, and Y, we can find Z that is large enough in the appropriate sense. If Atoms(Y) is infinite, then we can in fact find a suitable Atoms(Z) of the same cardinality as Y; just add $|\text{Atoms}(Y)|$ new atoms and let i be the inclusion map. If Atoms(Y) is finite then Atoms(Z) may need to be of larger cardinality than Atoms(Y), but it can still be finite. Thus, the proposition would remain correct even if we developed the whole theory in some contexts with limited numbers of atoms in our structures. For example, we could limit ourselves to finite sets of atoms, or to countable ones.

Remark 15. Results like Proposition 13 are common in the branch of set theory concerned with (non)-implications between various consequences of the axiom of choice. The idea goes back at least to [6, item 88], where such a result is proved in a more difficult context.

5 Equivalence of Background Classes

When are two background classes essentially the same? An intuitively natural criterion is that each structure in either class is isomorphic to a structure with the same atoms in the other class, by isomorphisms that fix the atoms and respect the standard embeddings. Here "respect the standard embeddings" means that if we take any standard embedding $\zeta : X \to Y$ of either class and compose it with the given isomorphisms $i : X \to X'$ and $j : Y \to Y'$ to background structures of the other class, then the result $j \circ \zeta \circ i^{-1} : X' \to Y'$ is a standard embedding of the other class.

This criterion can be simplified, because background classes are closed under isomorphisms, and we adopt the following simplified version as the official definition.

Definition 16. Two background classes $\mathcal{K}$ and $\mathcal{K}'$ are *equivalent* if they have the same background structures and there is a family of automorphisms α_X, one for each background structure X, such that

- α_X fixes all the atoms of X and
- an embedding $\zeta : X \to Y$ is standard in $\mathcal{K}$ if and only if $\alpha_Y \circ \zeta \circ \alpha_X{}^{-1} : X \to Y$ is standard in $\mathcal{K}'$.

The terminology is honest, in the sense that equivalence as defined here is clearly an equivalence relation.

This notion of equivalence is clearly a special case of the notion described in the first paragraph of this section, namely the special case where the corresponding structures in the two classes are identical.

Conversely, if two background classes $\mathcal{K}$ and $\mathcal{K}'$ satisfy the criterion in the first paragraph, say with isomorphisms $i_X : X \to X'$, then we claim that they also satisfy our official definition. For each background structure X, invoke Proposition 8 to let $\nu_X : X' \to X$ be the standard isomorphism in the sense of $\mathcal{K}'$ that fixes all the atoms. Then $\alpha_X = \nu_X \circ i_X$ is an automorphism of X, and we shall check that these automorphisms satisfy the definition of equivalence. It is clear that they fix all atoms, since both ν_X andi_X do. If $\zeta : X \to Y$ is standard in $\mathcal{K}$, then since $i_Y \circ \zeta \circ i_X{}^{-1}$, ν_Y, and $\nu_X{}^{-1}$ are standard in $\mathcal{K}'$, so is

$$\nu_Y \circ (i_Y \circ \zeta \circ i_X{}^{-1}) \circ \nu_X{}^{-1} = \alpha_Y \circ \zeta \circ \alpha_X^{-1}.$$

Conversely, if the isomorphism just displayed is standard in $\mathcal{K}'$, then so is $i_Y \circ \zeta \circ i_X{}^{-1}$ (as it is obtained by composing with the standard isomorphisms $\nu_Y{}^{-1}$ and ν_X), and therefore ζ is standard in $\mathcal{K}$.

It will often be convenient to define a background class by associating to each set S a specific background structure with S as its set of atoms. For example, in the set background, we might want to say that the background structure with S as atoms consists of (S and the obligatory elements and) the hereditarily finite sets over S. Of course, by Axiom 3, a background class cannot contain only these specific structures; it must also contain their isomorphic copies. But this enlargement of the class is, in some sense, inessential. The following proposition formalizes this observation.

Proposition 17. *Let Υ be a vocabulary, and assume given, for each set S, an Υ-structure $\tilde{S}$ with Atoms($\tilde{S}$) $= S$. Also assume given, for each embedding of sets $\zeta : S \to T$, an extension $\tilde{\zeta} : \tilde{S} \to \tilde{T}$ that is an embedding of Υ-structures. Finally, assume that $\widetilde{Id}_S = Id_{\tilde{S}}$ for all S and that $\widetilde{\zeta \circ \eta} = \tilde{\zeta} \circ \tilde{\eta}$ whenever the compositions here are defined. Then there is a background class for which each $\tilde{S}$ is a background structure and each $\tilde{\zeta}$ is a standard embedding. This background class is unique up to equivalence.*

Proof. We begin the definition of the desired background class by declaring the background structures to be the structures of the form $\tilde{S}$ and all their isomorphic copies. Thus we trivially satisfy Axiom 3 for background classes and the part of the proposition that says each $\tilde{S}$ is a background structure.

In order to define the standard embeddings, we must first choose, for each of our background structures X an isomorphism $i_X : X \to \tilde{S}$, where $S = \text{Atoms}(X)$, with i_X leaving all the atoms fixed. To see that such an isomorphism exists, note first that, since X is a background structure, we certainly have an isomorphism $j : X \cong \tilde{T}$ for some set T. This j restricts to a bijection between the sets of atoms, and we write $k : T \to S$ for the inverse of this bijection. Then $\tilde{k} : \tilde{T} \to \tilde{S}$ is an isomorphism (for the same reason as in the proof of Proposition 6) and $\tilde{k} \circ j$ is an isomorphism from X to $\tilde{S}$.

When $X = \tilde{S}$, we choose i_X to be the identity map, even if other atom-fixing isomorphisms are available.

Having chosen i_X for each background structure X, we define the standard embeddings to be those embeddings $X \to Y$ of background structures obtainable as follows. Let $S = \text{Atoms}(X)$ and $T = \text{Atoms}(Y)$ and let $\zeta : S \to T$ be an arbitrary embedding of sets. Then form the composite $\hat{\zeta} = i_Y^{-1} \circ \tilde{\zeta} \circ i_X$. As the notation $\hat{\zeta}$ suggests, this is an embedding $X \to Y$ that extends ζ, because i_X and i_Y leave all atoms fixed.

Our choice of i_X to be the identity map when $X = \tilde{S}$ ensures that, when both $X = \tilde{S}$ and $Y = \tilde{T}$, we have $\hat{\zeta} = \tilde{\zeta}$. Thus, we have satisfied the proposition's requirement that each $\tilde{\zeta}$ be a standard embedding.

We have already observed that Axiom 3 for background classes is satisfied by our construction. Axiom 4 is also clear, because $\tilde{S}$ has S as its set of atoms. For

Axiom 1, we compute that, when $S = T$ and $\zeta = \mathrm{Id}_S$ in the definition above,

$$\widehat{\mathrm{Id}}_S = i_X^{-1} \circ \tilde{\mathrm{Id}} \circ i_X = i_X^{-1} \circ i_X = \mathrm{Id}_X,$$

where we have used that $\tilde{\mathrm{Id}}_S = \mathrm{Id}_{\tilde{S}}$. Similarly, for Axiom 2, we use the assumption that $\widetilde{\zeta \circ \eta} = \tilde{\zeta} \circ \tilde{\eta}$ to compute, for $\hat{\eta} : X \to Y$ and $\hat{\zeta} : Y \to Z$, that

$$\hat{\zeta} \circ \hat{\eta} = i_Z^{-1} \circ \tilde{\zeta} \circ i_Y \circ i_Y^{-1} \circ \tilde{\eta} \circ i_X = i_Z^{-1} \circ \widetilde{\zeta \circ \eta} \circ i_X = \widehat{\zeta \circ \eta}.$$

The existence part of Axiom 5 is clear, as $\hat{\zeta}$ is a standard embedding extending ζ. To prove the uniqueness, notice that every standard embedding $X \to Y$ is, by our definition of standard, of the form $\hat{\eta}$ for some $\eta : \mathrm{Atoms}(X) \to \mathrm{Atoms}(Y)$. On atoms, $\hat{\eta}$ agrees with η. So the only way $\hat{\eta}$ can extend ζ is if $\eta = \zeta$, and then of course $\hat{\eta} = \hat{\zeta}$.

This completes the proof of the existence part of the proposition. We turn next to uniqueness, so suppose that we have some alternative background class that satisfies the conditions of the proposition; we must show that it coincides, up to equivalence, with the background class defined above.

In the alternative background class, every $\tilde{S}$ is a background structure, and so are all its isomorphic copies, by Axiom 3. On the other hand, if X is a background structure in the alternative sense and if S is its set of atoms, then by Axiom 5 the identity map of S must extend to a standard embedding $j_X : X \to \tilde{S}$, and by Proposition 6 this is an isomorphism. Thus, the alternative background class has exactly the same background structures as ours.

It remains to compare our standard embeddings with those of the alternative background class. Since Axiom 5 holds for both background classes, it suffices to consider an arbitrary embedding $\zeta : S \to T$ between the sets of atoms of two background structures X and Y, and to compare its (unique) extensions to standard embeddings, our $\hat{\zeta} : X \to Y$ and the corresponding extension $\check{\zeta} : X \to Y$ in the other background class.

Among the standard embeddings of the alternative background class are the $j_X : X \to \tilde{S}$ that we found above, the analogous $j_Y : Y \to \tilde{T}$, the inverse $j_Y^{-1} : \tilde{T} \to Y$ (by the proof of Proposition 6), and $\tilde{\zeta} : \tilde{S} \to \tilde{T}$ (by the requirements in the proposition). By Axiom 2, the composite $j_Y^{-1} \circ \tilde{\zeta} \circ j_X : X \to Y$ is also standard. It agrees with ζ on atoms, because j_X and j_Y leave all atoms fixed. So, by Axiom 5, we must have

$$\check{\zeta} = j_Y^{-1} \circ \tilde{\zeta} \circ j_X.$$

Notice that this formula is just like the definition of $\hat{\zeta}$ except that it uses j's instead of i's.

For each background structure X, define $\alpha_X = i_X^{-1} \circ j_X$, and notice that α_X is an automorphism of X (because i_X and j_X are isomorphisms) fixing all atoms (because i_X and j_X do). Furthermore, the formulas for $\hat{\zeta}$ and $\check{\zeta}$ combine to give

$$\hat{\zeta} = \alpha_Y \circ \check{\zeta} \circ {\alpha_X}^{-1}.$$

Thus, the automorphisms α_X witness the claimed equivalence of the two background classes. □

Remark 18. We indicate briefly, without proofs, a category-theoretic view of background classes. This material and subsequent category-theoretic remarks can be skipped without damage to understanding the rest of the paper.

The unary function symbol `Atomic` (like any unary function symbol) determines a functor from the category Υ-Str of Υ-structures and embeddings to the category Set of sets and embeddings (i.e., one-to-one functions). This functor sends each Υ-structure X to its set Atoms(X), and it sends each embedding ζ of Υ-structures to its restriction to the atoms.

The first two axioms for background classes say that a background class and its standard embeddings constitute a subcategory $\mathcal{K}$ of Υ-Str. Axiom 3 says that $\mathcal{K}$ is a replete subcategory of Υ-Str. Axiom 4 says that the restriction of the functor Atoms to $\mathcal{K}$ is surjective on objects, and Axiom 5 says that it is fully faithful.

Notice that this category-theoretic expression of our axioms is almost identical to that in Remark 4.12 of [2]. The differences are (1) that we no longer require background classes to be closed under intersections (but see Section 7 and particularly Remark 29 below) and (2) that the subcategory $\mathcal{K}$ need not be a full subcategory of Υ-Str, because not all embeddings are required to be standard.

Axioms 4 and 5 together imply that Atoms is an equivalence of categories between $\mathcal{K}$ and Set, i.e., it is fully faithful and essentially surjective on objects. ("Essentially surjective" means that every object of Set is isomorphic to one in the image of Atoms.) In fact, they imply slightly more, because Axiom 4 requires genuine surjectivity on objects, not merely essential surjectivity.[2] As a result, we get slightly better properties for the inverse functor $\mathcal{B}$ than we would get by just using that Atoms is an equivalence. Specifically, we get $\mathcal{B} : \text{Set} \to \mathcal{K}$ such that $\mathcal{B} \circ \text{Atoms}$ is naturally equivalent to the identity on $\mathcal{K}$ and $\text{Atoms} \circ \mathcal{B}$ is equal (not merely naturally equivalent) to the identity on Set. Note that such a functor $\mathcal{B}$ is exactly what was assumed (and denoted by tildes) in Proposition 17.

Although $\mathcal{B}$ is not uniquely determined by $\mathcal{K}$, it is determined up to natural equivalence. Furthermore, two background classes $\mathcal{K}$ and $\mathcal{K}'$ are equivalent if and only if the corresponding functors $\mathcal{B}$ and $\mathcal{B}'$ are naturally equivalent when considered as functors from Set to Υ-Str.

[2] Because of Axiom 3, it would make no difference if we assumed only essential surjectivity in Axiom 4. But this change, though quite natural from a category-theoretic viewpoint, would have made the non-category-theoretic formulation of the axioms slightly more complicated.

Remark 19. The category-theoretic formulation suggests a generalization of the idea of backgrounds. Instead of starting with entirely unstructured sets, like the reserve in abstract state machines, and describing what can be built over such a set without introducing any structure on the set itself, we could begin with some sort of structures, for example groups, and ask what can be built over them without introducing any *additional* structure.

Taking groups as a typical example, we can consider vocabularies Υ that have, among the obligatory symbols, the multiplication, inversion, and identity symbols of group theory. More precisely, these symbols are intended to denote a group structure on the atoms and to produce `undef` when applied to non-atoms. Let Υ-Mod be the category of Υ-structures whose atomic parts are groups, with embeddings as morphisms, so that Atoms is a functor from Υ-Mod to the category Grp of groups.

Then a generalized background class for groups would consist of a replete subcategory $\mathcal{K}$ of Υ-Mod such that the restriction of Atoms to $\mathcal{K}$ is a natural equivalence from $\mathcal{K}$ to Grp.

6 Finitary Backgrounds

In this section we describe background classes where each element of a background structure is "caused" by a finite set of atoms.

Definition 20. A background class is *finitary* if every element in every background structure has a finite source.

To show that a background class is finitary, it suffices to find finite sources for the non-obligatory non-atom elements of background structures, because, according to parts (7) and (8) of Proposition 11, the obligatory elements and the atoms always have finite sources.

Proposition 21. *For a finitary background class, every element of a background structure has a smallest (with respect to $\subseteq$) core, and this core is finite.*

Proof. Let x be an element of a background structure X. Finitariness and part (3) of Proposition 11 immediately imply that x has a finite core. Let C be a core of x with the smallest possible number of elements. If C' is any other core of x, then by Proposition 13 $C \cap C'$ is also a core of x. Since it cannot have smaller cardinality than C and since C is finite, it follows that $C \subseteq C'$. Thus, C is the smallest core of x. □

The proposition does not apply to sources in place of cores. Although an element of a background structure will have sources of smallest possible (finite) size, there may be many such minimal sources for the same element, so that none of the sources is the smallest. For instance, in Example 12, the minimal sources of the non-obligatory non-atom are all singletons of atoms; as soon as a background structure has at least two atoms, none of these sources is smallest.

We now describe some particular finitary background classes. Later, we shall show that any finitary background class can be viewed, as far as the elements and obligatory functions of the background structures are concerned, as being built from these particular examples. It is convenient to assume here that all background structures in the background classes under consideration satisfy `false` $\neq$ `undef`. All our considerations would also work under the opposite convention, but if we adopted no convention at all then the background classes $\mathcal{K}[n, G, k]$ introduced in the following example would need a fourth parameter (in addition to n, G, and k) to specify how `false` and `undef` are to behave.

Example 22. Let n and k be natural numbers, and let G be a subgroup of the group of all permutations of $[n] = \{1, 2, \ldots, n\}$. For any set S, define $[n, G, k](S)$ to be the structure, for the obligatory vocabulary, that has, in addition to the obligatory elements `true`, `false`, and `undef`, the set S of atoms, and, if $|S| \geq n + k$, the equivalence classes of one-to-one maps $x : [n] \to S$ under the orbit equivalence relation induced by G. That is, if $g \in G$, then $x \circ g$ is equivalent to x. We write x/G for the equivalence class of x. If $\zeta : S \to T$ is an embedding of sets, then its extension $\tilde{\zeta} : [n, G, k](S) \to [n, G, k](T)$ is defined in the obvious way on atoms and obligatory elements and by $\tilde{\zeta}(x/G) = (\zeta \circ x)/G$ on the "interesting" elements. It is trivial to check that this extension is well-defined.

It is also trivial to check that by assigning to each set S the structure $[n, G, k](S)$ and to each embedding ζ of sets the embedding $\tilde{\zeta}$ of the corresponding structures, we satisfy the hypotheses of Proposition 17. According to that proposition, there is a unique (up to equivalence) background class, which we call $\mathcal{K}[n, G, k]$, in which each $[n, G, k](S)$ is a background structure and each $\tilde{\zeta}$ is a standard embedding.

For an interesting element x/G in $[n, G, k](S)$, where $x : [n] \to S$, the smallest core is the range of x, and the sources are all supersets of the core that have at least $n + k$ elements. Since n and k are finite, $\mathcal{K}[n, G, k]$ is a finitary background class.

Our goal in the rest of this section is to show how every finitary background class can be regarded as built up from these particular examples $\mathcal{K}[n, G, k]$.

Consider any finitary background class $\mathcal{B}$. We shall write $\mathcal{B}(S)$ for a background structure in $\mathcal{B}$ with S as its set of atoms. Recall that such a structure exists by Axiom 4 and is unique up to unique standard isomorphism by Proposition 8. The particular choice of $\mathcal{B}(S)$ in this isomorphism class can be made arbitrarily; it will not matter for any of our work. In fact, we shall work directly only with the background structures $\mathcal{B}(S)$; our results can be transferred, via the standard isomorphisms, to all the other background structures of $\mathcal{B}$. We intend to show that the elements of $\mathcal{B}(S)$ are in one-to-one correspondence with the elements of a pseudo-disjoint union $\bigcup_{i\in I}[n_i, G_i, k_i](S)$ for a certain index set I, certain natural numbers n_i and k_i, and certain subgroups G_i of the group of all permutations of $[n_i]$. By "pseudo-disjoint" we mean that the obligatory elements and the atoms of the various $[n_i, G_i, k_i](S)$ are to be identified, but the rest of the structures are disjoint. All the parameters used in the pseudo-disjoint union representation, I, n_i, G_i, and k_i, depend only on the given background class, not on the particular S. Furthermore, the correspondences between the background structures $\mathcal{B}(S)$ of the given class and the pseudo-disjoint unions (though not quite canonical — there are some arbitrary choices involved) commute with the standard embeddings of the background class $\mathcal{B}$ and the various background classes $\mathcal{K}[n_i, G_i, k_i]$. In this situation, we sometimes refer to the given background class itself (and not just the individual structures in it) as having a pseudo-disjoint union representation.

According to pseudo-disjointness, $\bigcup_{i\in I}[n_i, G_i, k_i](S)$ contains a single `true`, a single `false`, a single `undef`, and a single set S of atoms. Our correspondences identify these with the `true`, `false`, `undef`, and set S of atoms in $\mathcal{B}(S)$. So from now on, we tacitly ignore obligatory elements and atoms, and we concentrate on the interesting parts of the background structures.

Temporarily fix natural numbers n and k. Consider the set $M(n, k)$ consisting of those elements $u \in \mathcal{B}([n + k])$ that have no smaller source than $[n + k]$ and that have smallest core $[n]$. The symmetric group $\mathrm{Sym}(n)$ of all permutations of $[n]$ acts on $M(n, k)$ as follows. Given a permutation π of $[n]$, extend it arbitrarily to a permutation, still called π, of $[n + k]$, and let it act on any $u \in M(n, k)$ by sending it to $\hat{\pi}(u)$. This $\hat{\pi}(u)$ does not depend on how we extended π to a permutation of $[n + k]$ because $[n]$ is a core of u. Furthermore, $\hat{\pi}(u) \in M(n, k)$; it has the right sources and cores by parts (4) and (5) of Proposition 11. It is easy to verify that we have defined an action of $\mathrm{Sym}(n)$ on $M(n, k)$. Let $I(n, k)$ be the set of orbits of this action. For each $i \in I(n, k)$, choose an element $u_i \in i$. (The notation i for a set may look strange here, but these i's will be the indices in the pseudo-disjoint union promised in the preceding paragraph.) Let G_i be the subgroup of $\mathrm{Sym}(n)$ that maps u_i to itself.

Now let n and k vary. Define I to be the disjoint union over all n and k of the sets $I(n, k)$. For any $i \in I$, write n_i and k_i for the numbers such that $i \in I(n_i, k_i)$. We now have all the ingredients involved in the pseudo-disjoint union $\bigcup_{i\in I}[n_i, G_i, k_i](S)$ that is to be in one-to-one correspondence with $\mathcal{B}(S)$. It remains to exhibit the correspondence C on the non-obligatory, non-atom elements and to check that it commutes with standard embeddings.

Consider any non-obligatory non-atom element of $[n_i, G_i, k_i](S)$ for any $i \in I$. It has the form x/G_i for some one-to-one map $x : [n_i] \to S$, and $|S| \geq n_i + k_i$, according to the definition of $[n_i, G_i, k_i](S)$. Thus, we can extend x to an embedding, which we still call x, from $[n_i + k_i]$ into S. Define $C(x/G_i)$ to be $\hat{x}(u_i)$. We must check that this is well-defined. First, it does not depend on how we extended x from its original domain $[n_i]$ to $[n_i + k_i]$, because $[n_i]$ is a core for u_i. Second, $C(x/G_i)$ depends only on the G_i-orbit x/G_i. To see this, we compare x with any other element $x \circ \pi$ of x/G_i; here π is an arbitrary element of G_i. We have already extended x to have domain $[n_i + k_i]$, and we have also, in defining the action of $\mathrm{Sym}(n_i)$ on $M(n_i, k_i)$, extended π to a permutation of $[n_i + k_i]$. The composition of these extensions provides an extension of $x \circ \pi$ to an embedding of $[n_i + k_i]$ into S. Using this in the definition of $C((x \circ \pi)/G_i)$, we obtain

$$\widehat{x \circ \pi}(u_i) = \hat{x}(\hat{\pi}(u_i)) = \hat{x}(u_i),$$

because π is in the subgroup G_i of $\mathrm{Sym}(n_i)$ fixing u_i. Thus, $C((x \circ \pi)/G_i) = C(x/G_i)$, as required for well-definedness of C.

As a first step toward showing that C is one-to-one, we consider an arbitrary non-obligatory non-atom x/G_i in the i^{th} component $[n_i, G_i, k_i](S)$ of our disjoint union, and we show how knowledge of only the element $z = C(x/G_i) \in \mathcal{B}(S)$ suffices to recover the numbers n_i and k_i. By definition, $z = \hat{x}(u_i)$ and u_i has smallest core $[n_i]$. By part (5) of Proposition 11, the image[3] $x[n_i]$ of this under $\hat{x}$ is a core of z, and, by part (7) of the same proposition, no proper subset of $x[n_i]$ can be a core of z. Thus, n_i is determined as the size of the smallest core of z.

Furthermore, the very definition of source gives that the whole range of x, i.e., $x[n_i + k_i]$ is a source of z. We claim that z has no source of smaller cardinality than $n_i + k_i$. To see this, suppose, toward a contradiction, that $A \subseteq S$ is a source for z and has size $< n_i + k_i$. Being a core of z (by Proposition 11, part (3)), A includes the smallest core $x[n_i]$. Recall that, in defining $C(x/G_i)$, we extended $x : [n_i] \to S$ arbitrarily to an embedding $x : [n_i + k_i] \to S$ (and we showed that the choice of this extension doesn't affect $C(x/G_i)$). Taking advantage of this arbitrariness, we can arrange that A is a proper subset of $x[n_i + k_i]$. But now part (6) of Proposition 11 applies and tells us that $x^{-1}[A]$ is a source for u_i. Since $x^{-1}[A]$ is a proper subset of $[n_i + k_i]$, this contradicts the fact that $u_i \in M(n_i, k_i)$. This contradiction establishes the claim that $n_i + k_i$ is the smallest size of any source of z. Thus z determines $n_i + k_i$ and therefore, since we already saw that it determines n_i, it also determines k_i.

In the course of the preceding argument, we also showed that z determines $x[n_i]$, namely as the smallest core of z. Any other one-to-one map of $y : [n_i] \to S$ with $y[n_i] = x[n_i]$ would be of the form $x \circ \pi$ for some permutation of $[n_i]$, so it would be in the same orbit $i \in I(n_i, k_i)$ as x. This shows that z determines the index i and thus also the group G_i.

---

[3]Our standard notation for images of sets, using square brackets, would give $x[[n]]$ here, but we use only one pair of brackets for aesthetic reasons.

Furthermore, if some other $y/G_i \in [n_i, G_i, k_i](S)$ is mapped by C to the same z, then consider the permutation $\pi = x^{-1} \circ y$ of $[n_i]$ (which makes sense because x and y biject $[n_i]$ to the same subset of S, namely the smallest core of z). Since $x \circ \pi = y$, we have

$$\hat{x}(\hat{\pi}(u_i)) = \hat{y}(u_i) = z = \hat{x}(u_i),$$

and since $\hat{x}$ is one-to-one, $\hat{\pi}(u_i) = u_i$. This means that $\hat{\pi} \in G_i$ and therefore x and $y = x \circ \pi$ are in the same orbit $x/G_i = y/G_i$ in $[n_i, G_i, k_i](S)$. This concludes the proof that C is one-to-one.

It remains to prove that it is surjective. For this purpose, consider an arbitrary $z \in \mathcal{B}(S)$. We must find some $i \in I$ and some $x/G_i \in [n_i, G_i, k_i](S)$ with $z = C(x/G_i) = \hat{x}(u_i)$. The proof that C is one-to-one gives us a good deal of information about where to look for i and x. We know that n_i must be the size of the smallest core of z and that $n_i + k_i$ must be the smallest size of any source of z. But we don't yet know i, so let us write simply n and $n + k$ for these two numbers. Fix a source A of z of size $n + k$; it includes the smallest core B of z. By definition of source, z is in the image of $\hat{j}$ where j is the inclusion map of A into S. Fix a bijection between A and $[n + k]$ that sends B onto $[n]$. This bijection and its inverse, along with the inclusion map j, show that z is in the image of $\hat{y}$ for some embedding $y : [n + k] \to S$ that sends $[n]$ onto B. Let $v \in \mathcal{B}[n + k]$ be such that $\hat{y}(v) = z$. By parts 4, 5, and 6 of Proposition 11, v has smallest core $[n]$ and has no smaller source than $[n + k]$. That is, $v \in M(n, k)$. Let $i \in I(n, k)$ be its orbit. So $n_i = n$ and $k_i = k$. We chose u_i from this orbit, so there is a permutation $\pi \in \mathrm{Sym}(n)$ such that $v = \hat{\pi}(u_i)$. Then

$$z = \hat{y}(v) = \hat{y}(\hat{\pi}(u_i)) = C((y \circ \pi)/G_i),$$

which establishes surjectivity.

Remark 23. Let $\mathcal{K}$ be a background class and $\mathcal{B} : \mathrm{Set} \to \mathcal{K}$ an inverse for Atoms as in Remark 18. Then $\mathcal{K}$ is finitary if and only if $\mathcal{B}$ preserves directed colimits.

7 Smallest Sources

In [2], the definition of background class included the requirement that each element of a background structure have a smallest source. In this section, we adapt this requirement to our new, broader notion of background classes; that is, we take into account the notion of standardness of embeddings. In contrast to [2], we do not incorporate this requirement into the definition of background classes but consider it as an additional property that a background class may or may not possess.

Definition 24. By a *least source* of an element x in a background structure X we mean a source $S \subseteq \text{Atoms}(X)$ for x in X that is a subset of every source for x in X. A background structure X *has least sources* if each of its elements has a least source. A background class *has least sources* if each of its background structures has least sources.

We begin by relating this notion of least sources to cores.

Proposition 25. *Suppose an element x of a background structure X has a least source S and that S is not all of $\text{Atoms}(X)$. Then S is also the least core of x.*

Proof. According to part 3 of Proposition 11, S is a core for x. If there were another core that is not a superset of S, then, intersecting that core with S, we would get, by Proposition 13, a core C that is a proper subset of S. Choose an element $a \in S - C$ and an element $b \in \text{Atoms}(X) - S$. Let π be the permutation of $\text{Atoms}(X)$ that interchanges a with b and fixes all the other atoms. The standard automorphism $\hat{\pi}$ of X fixes x because π fixes all elements of the core C. By part 4 of Proposition 11, $\pi[S]$ is a source for $\hat{\pi}(x) = x$. But $a \in S - \pi[S]$, so this contradicts the assumption that S is the smallest source for x. □

Example 26. Recall that an interesting element x/G of a background structure $[n, G, k](S)$ has the range of x as its smallest core and has as sources all supersets of the core of size at least $n + k$. Thus, x/G has a least source if and only if either $k = 0$ or $|S| = n + k$. Since this criterion is independent of the particular x, it is also the criterion for $[n, G, k](S)$ to have least sources. The background class $\mathcal{K}[n, G, k]$ has least sources if and only if $k = 0$.

It follows easily that a finitary background class has least sources if and only if, when it is represented by a pseudo-disjoint union of background classes $\mathcal{K}[n_i, G_i, k_i]$ as in the preceding section, all the k_i are zero.

To understand what is needed for the existence of least sources in a structure, it is convenient to fix, for the time being, a background class $\mathcal{B}$ and a specific background structure X in $\mathcal{B}$. We shall consider elements of X and their sources. For any subset S of $\text{Atoms}(X)$, fix a background structure $\mathcal{B}(S)$ with S as its set of atoms. (Recall that such a structure exists and is unique up to unique standard isomorphism.) The inclusion map $i : S \to \text{Atoms}(X)$ induces a standard embedding $\hat{\imath} : \mathcal{B}(S) \to X$. The image of this map, which is a submodel of X isomorphic (via $\hat{\imath}$) to $\mathcal{B}(S)$, will be denoted by $\text{Smota}(S)$. (The name is an acronym for "**s**ub**m**odel **o**ver **t**he **a**toms" as well as being "atoms" backward.) Notice that $\text{Atoms}(\text{Smota}(S)) = S$ for all subsets of $\text{Atoms}(X)$. A substructure Y of X satisfies $Y = \text{Smota}(\text{Atoms}(Y))$ if and only if it is the image of $\hat{\imath}$ for some inclusion map i. We shall refer to such substructures Y, those of the form $\text{Smota}(S)$, as the *natural* substructures of X.

Remark 27. It would seem reasonable to call them the *standard* substructures of X, since they are obtained by means of standard embeddings. We refrain from doing so, because that terminology could suggest that the inclusion map $\mathrm{Smota}(S) \to X$ is a standard embedding, and this need not be the case. In fact, this inclusion will be standard if and only if $\hat{\imath}$, which is of course a standard embedding of $\mathcal{B}(S)$ into X, is also standard when considered as an isomorphism from $\mathcal{B}(S)$ to $\mathrm{Smota}(S)$. For an example (the simplest possible one) where these are not standard, use the example in Subsection 8.1 below of a non-rigid background class where each background structure has two non-obligatory non-atoms a and b, and take advantage of the freedom (since a and b are not named in the vocabulary) to assign these labels differently in X and in a substructure Y.

It follows immediately from the definitions that the elements of $\mathrm{Smota}(S)$ are exactly those $x \in X$ for which S is a source.

Notice also that Smota is monotone, i.e., for $S \subseteq T \subseteq \mathrm{Atoms}(X)$, we have $\mathrm{Smota}(S) \subseteq \mathrm{Smota}(T)$, because the inclusion map $i : S \to \mathrm{Atoms}(X)$ factors through the inclusion $j : T \to \mathrm{Atoms}(X)$ and so $\hat{\imath}$ factors through $\hat{\jmath}$. (In view of the connection between Smota and sources, this amounts to part (1) of Proposition 11.)

Proposition 28. *The following are equivalent.*

- *X has least sources.*
- *The intersection of any family of natural substructures of X is a natural substructure.*
- *The operation Smota, from subsets of Atoms(X) to natural substructures of X, preserves intersections.*

Proof. First, let us assume (3). Then we obviously have (2), and we obtain (1) as follows. Consider any $x \in X$ and let S_0 be the intersection of all the sources S for x. Since $x \in \mathrm{Smota}(S)$ for each of these sources, (3) implies that $x \in \mathrm{Smota}(S_0)$. That is, S_0 is a source for x, and it is, because of its definition, clearly the least source for x.

Conversely, let us assume (1) and prove (3). So let I be the intersection of some family $\mathcal{F}$ of subsets of X; we must prove that $\mathrm{Smota}(I) = \bigcap_{S\in\mathcal{F}} \mathrm{Smota}(S)$. The $\subseteq$ half of this follows from monotonicity of Smota, so it remains to prove the $\supseteq$ half. Consider any $x \in \bigcap_{S\in\mathcal{F}} \mathrm{Smota}(S)$, and, by assumption (1), let S_0 be its smallest source. Each $S \in \mathcal{F}$ is a source for x, so $S_0 \subseteq S$. Thus $S_0 \subseteq I$, which means that I is a source for x, as required.

Finally, we show that (2) implies (3). With I and $\mathcal{F}$ as above, we must again show that $\mathrm{Smota}(I) = \bigcap_{S\in\mathcal{F}} \mathrm{Smota}(S)$, this time using (2) rather than (1). By (2),

$\bigcap_{S\in\mathcal{F}} \mathrm{Smota}(S)$ is a natural substructure, say Smota(J), and all we must show is that $J = I$. But, since the operation Atoms clearly preserves intersections, we have

$$J = \mathrm{Atoms}(\mathrm{Smota}(J)) = \mathrm{Atoms}(\bigcap_{S\in\mathcal{F}} \mathrm{Smota}(S)) = \bigcap_{S\in\mathcal{F}} \mathrm{Atoms}(\mathrm{Smota}(S)) = \bigcap_{S\in\mathcal{F}} S = I, \quad (1)$$

and the proof is complete. □

Remark 29. Let $\mathcal{K}$ be a background class and $\mathcal{B} : \mathrm{Set} \to \mathcal{K}$ an inverse for Atoms as in Remark 18. Then $\mathcal{K}$ has least sources if and only if $\mathcal{B}$ preserves intersections.

8 Constructivity

We use "constructivity" as a general term for the idea that the elements of a background structure are somehow constructed from the atoms. In a very abstract sense, such an idea is already present in Axiom 5 for background classes, according to which a standard embedding between two background structures is completely determined by its action on atoms. In this section, we are concerned with stricter and more concrete notions of constructivity, notions that can be expressed with reference only to the structures, not to the standard embeddings. (Such notions may, however, have equivalent formulations in terms of standard embeddings.) The idea is that the functions of the vocabulary should suffice to construct all the elements of a background structure from the atoms, or at least should show how the atoms determine the remaining elements. We present several such notions of constructivity, in order of increasing strength, indicating some equivalent characterizations as well as examples showing that the strength is really increasing.

8.1 Rigid

We call a structure *rigid relative to atoms* if the only automorphism that fixes all the atoms is the identity map. Equivalently, if we expand the vocabulary by adding names for all the atoms, then the resulting structure is rigid, meaning that it has no non-trivial automorphisms. We call a background class *rigid relative to atoms* if each of its background structures is rigid relative to atoms.

Although the definition of rigidity relative to atoms refers only to the background structures, not to the standard embeddings, Proposition 7 provides an equivalent characterization in terms of standard embeddings, namely that all isomorphisms between background structures are standard.

Rigidity relative to atoms is the weakest of the constructivity notions that we shall consider, but it is not trivial. That is, not all background classes are rigid relative to atoms. For a counterexample, consider a background class whose background structures each have, in addition to atoms and obligatory elements, exactly two other elements, arbitrarily labeled a and b. The vocabulary, however, contains only the obligatory symbols, not names for a and b. Define the standard embeddings to be those embeddings between background structures that map the a of the domain to the a of the range (and therefore also map b to b). This defines a background class. It is not rigid, because every background structure has a non-trivial (and non-standard) automorphism that fixes all the atoms and obligatory elements but interchanges a with b.

The need for a non-trivial notion of equivalence (Section 5) is due to the possible lack of rigidity relative to atoms. The following proposition exhibits the connection between the two.

Proposition 30. *Let $\mathcal{K}$ and $\mathcal{K}'$ be two background classes, at least one of which is rigid. Then $\mathcal{K}$ and $\mathcal{K}'$ are equivalent if and only if they are equal.*

Proof. Of course equality implies equivalence, so we prove the converse, assuming without loss of generality that $\mathcal{K}$ is rigid. As the two background classes are equivalent, they have the same background structures; we must show that they have the same standard embeddings. Let α_X for $X \in \mathcal{K}$ be automorphisms as required by the definition of equivalence; thus $\zeta : X \to Y$ is standard in $\mathcal{K}$ if and only if $\alpha_Y \circ \zeta \circ \alpha_X{}^{-1}$ is standard in $\mathcal{K}'$. By rigidity, we know that each α_X is standard. But, as required in the definition of equivalence, α_X fixes all the atoms of X. By Axioms 1 and 5, it follows that $\alpha_X = \mathrm{Id}_X$ for all X, and therefore $\alpha_Y \circ \zeta \circ \alpha_X{}^{-1} = \zeta$. Thus, ζ is standard in $\mathcal{K}$ if and only if ζ is standard in $\mathcal{K}'$. □

8.2 All embeddings standard

A stronger notion of constructivity is obtained by requiring not only all automorphisms but all embeddings between background structures to be standard. A background class with this property satisfies the axioms for background classes in [2] provided it has least supports.

An equivalent characterization that doesn't mention standardness is that any embedding between background structures is determined by its restriction to the atoms. In other words, no two distinct embeddings between the same background structures can agree on all atoms. That this characterization is equivalent to standardness of all embeddings follows immediately from the existence and uniqueness of a standard embedding between background structures extending any given embedding between the atoms.

To see that this form of constructivity is strictly stronger than rigidity relative to atoms, it suffices to consider the example from [2] discussed in Section 3 above, where the non-obligatory non-atom elements of a background structure are the hereditarily finite sets over the atoms. As discussed above, this class fails to satisfy the axioms of [2], but it can be made to satisfy the axioms of the present paper by a suitable, natural choice of standard embeddings. Some non-standard embeddings between background structures, noticed by Yavorskaya, were described in Section 3. Yet it is easy to check that the only automorphisms of these background structures are the standard ones. Thus, this example is rigid relative to atoms but not all embeddings are standard.

In the following subsections, we shall consider even stronger notions of constructivity. So we can, in these subsections, omit "standard" as all embeddings will be standard.

8.3 Atom generated

As in [2], we call a structure X *explicitly atom-generated* if the smallest (and therefore the only) substructure containing all the atoms is X itself. Recall that, in any structure X, the smallest substructure that includes some given subset A consists of all the values, in X, of terms when the variables are assigned values in A. Equivalently, it consists of the values of closed terms after we enrich the vocabulary by adding names for all elements of A. In particular, a structure X is explicitly atom-generated if all its elements can be obtained as values of terms when the variables are assigned atoms as values.

We call a background class *explicitly atom-generated* if all its background structures are. This property easily implies the property considered in the preceding subsection, namely that an embedding between background structures is determined by what it does to the atoms. Indeed, if two embeddings $\zeta, \eta : X \to Y$ agree on the atoms, then $\{x \in X : \zeta(x) = \eta(x)\}$ is (the base set of) a substructure of X containing all the atoms, so it must be all of X.

The converse does not hold. There are background classes in which all embeddings are standard but not all background structures are explicitly atom-generated. A simple example is given by Example 12. (It makes no difference for our current purpose that the empty set of atoms is treated differently in this example; we could equally well have one non-obligatory non-atom in all background structures, even when there are no atoms.) The non-obligatory non-atom is not the value of any term with atoms as values of the variables. Yet all embeddings respect it, and so all embeddings are determined by their restrictions to atoms.

For a less trivial example, define a background class by letting the non-obligatory non-atom elements of any background structure be in one-to-one correspondence with the atoms, and let the vocabulary include a unary function symbol f denoting a bijection from these elements to the atoms. Then the background structures are not explicitly atom-generated (except when the set of atoms is empty), essentially because the vocabulary doesn't provide a symbol for f^{-1}. Yet an embedding between two such background structures is determined by what it does to the atoms, because it must preserve f.

The same idea provides an even less trivial example, in which the non-obligatory non-atoms are ordered pairs of atoms, and the vocabulary has non-obligatory symbols only for the projection functions $\langle x, y\rangle \mapsto x$ and $\langle x, y\rangle \mapsto y$, not for the pairing function $x, y \mapsto \langle x, y\rangle$. Similar examples can be constructed along the same lines, by including deconstructors but not constructors in the vocabulary.

8.4 Freely atom generated

In the situation of the preceding subsection, each element of a background structure X can be named by a term of the vocabulary enriched with constant symbols for all the atoms of X. This does not, however, fully describe what the elements of X are. The problem is that a single element may have many names; for a full description of the base set of X we would need to say under what circumstances two names denote the same element. In this subsection, we consider one simple and useful form in which such information could be available, namely identities satisfied by our names, or by a suitable subfamily of them.

As in universal algebra, we use "identity" to mean a formula of the form $t = t'$, where t and t' are terms. By an *instance* of this identity, we mean any formula obtained by substituting other terms for the variables; it is a *closed* instance if it contains no variables. Semantically, an identity is interpreted as universally quantified. That is, an identity holds in a structure if and only if it is true for all assignments of values to its variables. Thus, for example, an identity semantically entails all its instances.

We say that an Υ-structure X is *freely atom-generated* via a subset Γ of the vocabulary Υ and a set Σ of identities over the vocabulary Γ if

- all elements of X are values of closed terms in the vocabulary Γ enriched with names for all atoms, and
- two such terms have the same value in X if and only if this equality is a consequence of the identities in Σ.

In particular, such a structure is explicitly atom-generated (even using only Γ rather than all of Υ). In addition, Σ provides a criterion for when two names should denote the same element. Although this criterion was formulated in terms of the semantic notion of consequence, there is an equivalent syntactic formulation. Specifically, an identity is a consequence of Σ if it can be obtained from Σ by instantiation and rules of inference that embody the information that equality is an equivalence relation and is preserved by all the functions.

We say that a background class is *freely atom-generated* via Γ and Σ if all its background structures are. The practical significance of this concept is that it provides quite explicit control over the elements of background structures, at least when Σ is reasonably simple.

Not every explicitly atom-generated background class is freely atom-generated. Consider, for example, the following admittedly artificial background class. The non-obligatory non-atoms in any background structure are the exactly three-element sets of atoms and the ordered pairs of atoms. (Strictly speaking, this describes one representative of the relevant isomorphism class; the rest is obtained as in Proposition 17.) The vocabulary has one non-obligatory symbol, a ternary function f. Its values in any background structure are `undef` except in the following two situations.

- If x, y, z are three distinct atoms, then $f(x, y, z) = \{x, y, z\}$.
- For any atoms x and y, $f(x, x, y) = \langle x, y\rangle$.

It is clear that these background structures are explicitly atom-generated, since f provides access to all the non-obligatory non-atoms. If we try to find Γ and Σ to make this class freely atom-generated, then we see first that Γ must contain f, for without f we cannot name the sets or pairs. Next, whenever x, y, z are distinct atoms, then $f(x, y, z) = f(x, z, y)$ must be a consequence of Σ. The only way for this to happen is for this identity to be a consequence of Σ even with variables in place of the atoms x, y, z. But then it is also a consequence of Σ that $f(x, y, x) = f(x, x, y)$ for any atoms x, y. When $x \neq y$, this gives `undef` $= \langle x, y\rangle$, which isn't true in our background structures. Thus, this background class is not freely atom-generated; in fact, none of its background structures with three or more atoms is freely atom-generated.

In this example, the failure of free atom-generation can be traced to what might be viewed as a defect in the semantics of identities, namely that one must treat all instances together, not for example just those instances where the variables are given distinct values. More generally, one might want to impose other sorts of conditions on the instantiations, for example that the values of variables be atoms. The example would disappear if we could impose the identity $f(x, y, z) = f(x, z, y)$ and others like it subject to the condition that x, y, z are distinct atoms. Unfortunately, such an interpretation does not lend itself to as simple a deductive calculus as the usual interpretation of identities does. In particular, one would need a way of deducing non-equations, to be used in verifying the distinctness hypotheses that could be imposed on identities.

Other variants of free atom-generation may also be desirable from the point of view of doing (or analyzing) computations in background sructures. If Γ is not all of Υ, then it would be useful to have, in addition to the names and information about equality between their values, a description, in terms of our Γ-names, of the functions in $\Upsilon - \Gamma$. That is, it would be useful to have a calculus for deducing information of the form $f(t_1, \ldots, t_k) = t_0$, where f is a k-ary function symbol from $\Upsilon - \Gamma$ and where the t_i are terms over Γ (possibly with free variables). We do not pursue this additional desideratum here, since our primary interest in this section is constructing the elements of background structures from the atoms.

8.5 Absolutely freely atom generated

Absolutely free atom-generation is the special case of free atom-generation where the set Σ of identities is empty. In this case, two names over Γ enriched by constant symbols for the atoms will have the same value if and only if they are syntactically identical. A background structure can then be regarded (up to isomorphism) as simply consisting of the names. Such a structure is often called an Herbrand universe (with respect to the enriched Γ).

There are natural examples of freely but not absolutely freely generated background classes. For one example, let the non-obligatory non-atom elements of any background structure be obtained by iterating the "list of" construction, starting with the atoms and obligatory elements. That is, the structure contains the atoms, the obligatory elements, and all finite sequences of its elements. Let the vocabulary contain a nullary function `nil` (to be interpreted as the empty list), a unary function f (to be interpreted as sending any x to the one-element sequence $\langle x \rangle$) and a binary operation $*$, to be interpreted as concatenation. Then each background structure is freely atom-generated via $\Gamma = \{\mathtt{true}, \mathtt{false}, \mathtt{undef}, \mathtt{nil}, f, *\}$ and Σ consisting of the associative law for $*$ and the laws saying that `nil` is a two-sided identity element for $*$. Absolutely free generation would require getting rid of all these laws, so that instead of lists we would have ordered binary trees.

8.6 Constructivization

The constructivity properties of some background classes can be improved by adding suitable new functions to the vocabulary, with suitable interpretations in the background structures. In this section, we discuss the possibility of achieving explicit atom-generation in this manner. What makes this task, which we call constructivization, non-trivial is that the same new function symbols must be interpreted in all of the background structures, and the interpretations must be preserved by the standard embeddings (so that these remain embeddings for the enlarged vocabulary).

If a background class admits a constructivization of this sort, then it must be finitary. Indeed, any element x in any background structure X is the value of some term t for some assignment of atoms as values to the (finitely many) variables in t. Then the set of atoms assigned to the variables is a source for x.

Any finitary background class admits a constructivization, provided we allow an infinite vocabulary. To see this, represent the background class as the pseudo-disjoint union, indexed by a set I, of background classes $\mathcal{K}[n_i, G_i, k_i]$ and let the vocabulary contain, for each $i \in I$, an $(n_i + k_i)$-ary function symbol f_i, interpreted in background structures as follows. Let $(\vec{a}, \vec{b})$ be any $(n_i + k_i)$-tuple of elements of a background structure X, with the first n_i components constituting $\vec{a}$ and the remaining k_i constituting $\vec{b}$. If at least one component in $(\vec{a}, \vec{b})$ is not an atom or if some two components are equal, then $f_i^X(\vec{a}, \vec{b}) = \mathtt{undef}$. Otherwise, regard the n_i-component vector $\vec{a}$ as a (one-to-one) map $[n_i] \to \text{Atoms}(X)$, and define $f_i^X(\vec{a}, \vec{b})$ as $\vec{a}/G_i \in [n_i, G_i, k_i](\text{Atoms}(X))$, or, more precisely, as the element of X that corresponds to $\vec{a}/G_i$ in the pseudo-disjoint union representation of X. From the fact, proved in Section 6, that standard embeddings commute with this representation, it follows that they also commute with these new functions f_i. So adding the f_i to the vocabulary does not require any changes in the notion of standard embedding; we still have a background class, and the f_i ensure that it is explicitly atom-generated.

If, on the other hand, we insist that vocabularies be finite, then the preceding constructivization is not available when I is infinite. Indeed, in such a case, we cannot arrange for all elements of background structures to be obtained by applying functions of the structure directly to atoms. Nevertheless, in favorable circumstances, we can achieve explicit atom-generation with finitely many functions by using compositions of these functions.

A typical and useful example is given by the set background, the example of Section 5.1 in [2] as corrected by using our new notion of background class. The non-obligatory non-atoms of a background structure are the hereditarily finite sets over the atoms, and the standard embeddings act on sets by acting (recursively) on their elements. This background is finitary, since a hereditarily finite set involves only finitely many atoms. But its pseudo-disjoint representation in terms of $\mathcal{K}[n, G, k]$ classes involves infinitely many components. Indeed, there are infinitely many components just of the form $\mathcal{K}[0, \text{Sym}(0), 0]$, one for every pure set. (Pure sets are sets involving no atoms, like $\varnothing$, $\{\varnothing\}$, etc.) Nevertheless, this background class can be constructivized, for example by adding a constant symbol for $\varnothing$ and a binary function symbol for $x \cup \{y\}$. Iterated application of this binary function to atoms and $\varnothing$ suffices to generate all the hereditarily finite sets.

Not all finitary background classes admit constructivization with finitely many new functions. Here is a counterexample. Define the background structures to consist of the obligatory elements, the atoms, and those tuples $\vec{a}$ of atoms such that

- all components of $\vec{a}$ are distinct, and
- the length of $\vec{a}$ is of the form $n!$ for some positive integer n.

There are no non-obligatory function symbols in the vocabulary. Standard embeddings act componentwise on tuples. This is a finitary background class with least sources; the minimal source and core of a tuple is the set of its components. So the cardinalities of all nonempty cores are of the form $n!$. Now suppose we had a constructivization involving only finitely many functions. Let r be the largest of their arities. Consider what happens when any one of these functions, say f, is applied to arguments whose least sources have size $r!$ or less. Since f has arity at most r, there is a set of size at most $r \cdot r!$ that is a source for all of the arguments simultaneously and therefore also for the value of f. Since $r \cdot r! < (r+1)!$, the least source of the value of f has size $< (r+1)!$ and therefore $\leq r!$. That is, in any background structure, the set of elements whose least source has size at most $r!$ is closed under all of our functions. Therefore, even by composing these functions, we can never produce, from the atoms, any of the tuples of length $> r!$.

Acknowledgment

We thank Tatiana Yavorskaya for informing us about the error in [2].

References

[1] The AsmL webpage, `http://research.microsoft.com/foundations/AsmL/`.

[2] Andreas Blass and Yuri Gurevich, "Background, reserve, and Gandy machines," *Computer Science Logic*, P. G. Clote and H. Schwichtenberg, eds., Springer-Verlag, Lecture Notes in Computer Science 1862 (2000) 1–17.

[3] Erich Grädel and Yuri Gurevich, "Metafinite model theory," *Inform. and Comput.* 140:1 (1998) 26–81.

[4] Rami Grossberg, "Classification theory for non-elementary classes," *Logic and Algebra*, Y. Zhang,ed., American Mathematical Society, Contemporary Math. 302 (2002) 165–204.

[5] Yuri Gurevich, "Evolving algebra 1993: Lipari guide," in *Specification and Validation Methods*, E. Börger, ed., Oxford University Press (1995) 9–36.

[6] Andrzej Mostowski, “Über die Unabhängigkeit des Wohlordnungssatzes vom Ordnungsprinzip,” *Fund. Math.* 32 (1939) 201–252.