An application of first-order compactness in canonicity of relation algebras

Ian Hodkinson*

June 17, 2019

Abstract

The classical compactness theorem is a central theorem in first-order model theory. It sometimes appears in other areas of logic, and in perhaps surprising ways. In this paper, we survey one such appearance in algebraic logic. We show how first-order compactness can be used to simplify slightly the proof of Hodkinson and Venema (2005) that the variety of representable relation algebras, although canonical, has no canonical axiomatisation, and indeed every first-order axiomatisation of it has infinitely many non-canonical sentences.

1 Introduction

The *compactness theorem* is a fundamental result in first-order model theory. It says that every first-order theory (set of first-order sentences) that is consistent, in the sense that every finite subset of it has a model, has a model as a whole. Equivalently, if T, U are first-order theories, $T \models U$ (meaning that every model of T is a model of U), and U is finite, then there is a finite subset $T_0 \subseteq T$ with $T_0 \models U$.

Though not nowadays the most sophisticated technique in model theory, compactness is still very powerful and has a firm place in my affections — as I believe it does in Mara's, since she devoted an entire chapter of her *Model Theory* book to it [26, chapter 5]. See [26, theorem 5.24] for the theorem itself.

This paper was intended to be a short and snappy survey of a couple of applications of compactness in *algebraic logic.* To keep it short, it has turned out in the end to cover only one application, but perhaps it still has some value. The application may be at least a little surprising and entertaining, since at first sight it doesn't seem to have much to do with compactness. The proof, though a minor variant of one in the literature, has not actually appeared before. And

*Department of Computing, Imperial College London, SW7 2AZ, UK.
This article is dedicated to Maria Manzano on the occasion of her retirement.

given Mara's longstanding interest in teaching logic, she might be happy that it uses work done by a student, Jannis Bulian, in his 3rd-year undergraduate project [2], a Distinguished Project in my department in 2010–2011.[1]

And so to business. We are going to show, in §3, that every first-order axiomatisation of the class of representable relation algebras contains infinitely many non-canonical axioms (the technical terms here will of course be explained later). This was originally proved in joint work with Yde Venema [14]. Bulian used a similar but slightly simpler method to prove analogous results for cylindric algebras in his project [2]. This work was then extended to polyadic and other algebras and published in [3]. The proof that we will sketch in §3 applies the method of [2, 3] to relation algebras.

Notation We will be using several kinds of first-order sentences, and we recall their names here. An *equation* is a first-order sentence of the form $\forall x_1 \ldots x_n(t = u)$, where t and u are terms. A *universal sentence* is one of the form $\forall x_1 \ldots x_n \varphi$, where φ has no quantifiers. Equations are examples of universal sentences, as are quantifier-free sentences (take $n = 0$). An $\forall\exists$-*sentence* is one of the form $\forall x_1 \ldots x_n \exists y_1 \ldots y_m \varphi$, where φ has no quantifiers. Every universal sentence is an $\forall\exists$-sentence (take $m = 0$).

2 Algebras of relations

We begin with some background information on the part of algebraic logic relevant to our application. It concerns *relations.* A relation on a set U is just a subset of the cartesian power U^n, for some natural number $n \geq 1$ (the *arity* of the relation). When $n = 1$, the relation is *unary,* when $n = 2$ it is *binary,* and so on. Relations are important because they are ubiquitous in mathematics. They are a basic notion in first-order logic, too, as developed in Mara's book [26]. (Many of the terms used below can be found in this book. To reduce distraction, we will not cite all of them explicitly.)

We want to consider not just individual relations, but entire *collections* of relations of fixed arity on a fixed (but arbitrary) set U. (Some approaches, not taken here, allow multiple arities.) And we want to endow the collections with *operations* on relations that are commonly used in practice. We can then try to elucidate the laws governing these operations. Let's see how we do it.

[1] Jannis was a student on the 3-year Mathematics and Computer Science course at Imperial College. In June 2011 the Department of Computing held its annual open day, and about eight students gave talks on their potentially prize-winning final-year projects. At the end of the day, the audience was invited to take part in two votes, for Best Presentation and Best Project. Given that the audience consisted largely of industrialists, it was one of the most surprising (and nicest) moments of my working life when Jannis won the Best Project vote.

2.1 Unary relations

Here we can take as our collection the full power set $\wp(U) = \{a : a \subseteq U\}$ of all unary relations on the set U. As operations, though other choices are possible, we might choose union, intersection, and complement. We might as well include the empty relation and U as distinguished elements, since it is useful to be able to talk about them. This leads us to an *algebra* $(\wp(U), \cup, \cap, -, \emptyset, U)$ — a standard first-order structure in the sense of the very first definition in [26] (definition 1.1, page 6), but called an algebra because ironically, the operations involve no relations but only functions [26, p.7]. (The distinguished elements are just nullary functions.)

We want to study the *class* $\mathcal{U}$, say, of all such algebras, taken over all sets U, just as we might study the class of all symmetric groups consisting of all permutations on an arbitrary set. Now, it is known by Cayley's theorem that *every group embeds into a symmetric group.* That is to say, the closure under isomorphism and subgroups of the class of symmetric groups is the class of all groups — a class with a nice simple definition by equations [26, definition 1.6]. Could we hope for an analogous result for $\mathcal{U}$? After all, a subalgebra of $(\wp(U), \cup, \cap, -, \emptyset, U)$ is a perfectly legitimate algebra of relations — it consists of just *some* of the unary relations on U. So we ask:

($*$) can we characterise the closure of $\mathcal{U}$ under isomorphism and subalgebras?

It turns out, mainly by Stone's theorem [29], that this closure is exactly the class BA of *boolean algebras.* BA is an elementary class, defined by a few straightforward equations expressing standard properties of the operations [26, definition 1.23]. So these equations can be used for sound and complete reasoning about unary relations, for example in the way described in [26, chapter 3], or by purely equational reasoning. We have answered our question ($*$) very satisfactorily.

We make a couple more remarks. Classes defined by equations are called *varieties,* and BA is consequently a variety. By Birkhoff's theorem [1], a class of similar algebras is a variety iff it is closed under subalgebras, products, and homomorphic images. These notions can be found in [26].

Stone's theorem is relevant later, because it actually shows that every boolean algebra $\mathcal{B}$ embeds into a boolean algebra $\mathcal{B}^\sigma$ called the *canonical extension* of $\mathcal{B}$. It is of the form $(\wp(\mathcal{B}_+), \cup, \cap, -, \emptyset, \mathcal{B}_+)$, where $\mathcal{B}_+$ is the set of ultrafilters of $\mathcal{B}$ (see [26, definition 5.77] for ultrafilters). Clearly, $\mathcal{B}^\sigma \in \mathcal{U}$, and it follows that the closure of $\mathcal{U}$ under subalgebras and isomorphism contains BA. The converse inclusion is easily checked. So BA is closed under taking canonical extensions: if $\mathcal{B} \in \mathsf{BA}$ then $\mathcal{B}^\sigma \in \mathsf{BA}$ as well. A class of algebras that is closed under canonical extensions is said to be *canonical.* So BA is a canonical variety.

2.2 Binary relations

Let us now try to do the same for binary relations. We take the set $\wp(U \times U)$ of all binary relations on the set U, and make it an algebra by adding some sensible operations. There is really quite a wide choice of operations here, but a common one is to add the boolean operations $\cup, \cap, -, \emptyset, U \times U$, which still make sense in this context, plus the following distinctively binary-relational ones:

- $Id_U = \{(u, u) : u \in U\}$ (a distinguished element)
- the converse operation $-^{-1}$, where $a^{-1} = \{(u, v) : (v, u) \in a\}$ for $a \subseteq U \times U$,
- relational composition $|$, where, for $a, b \subseteq U \times U$,
$$a \mid b = \{(u, v) : (u, w) \in a \text{ and } (w, v) \in b \text{ for some } w \in U\}.$$

We obtain an algebra of binary relations:

$$\mathfrak{Re}(U) = (\wp(U \times U), \cup, \cap, -, \emptyset, U \times U, Id_U, -^{-1}, |).$$

Again, a subalgebra of $\mathfrak{Re}(U)$ is a perfectly legitimate algebra of binary relations, consisting of just some of the binary relations on U. But this time, the closure of the class $\{\mathfrak{Re}(U) : U \text{ a set}\}$ under isomorphism and subalgebras is not a variety, because (as is not so hard to see) it is not closed under products. However, we can view a product $\prod_{i\in I} \mathfrak{Re}(U_i)$ as a quite sensible algebra of binary relations on the disjoint union $U = \dot{\bigcup}_{i\in I} U_i$ of the U_i, by viewing an element $(a_i : i \in I)$ of this product as the binary relation $\dot{\bigcup}_{i\in I} a_i$ on U. The only price to pay is that the $\subseteq$-largest element of $\prod_{i\in I} \mathfrak{Re}(U_i)$ is not $U \times U$, but rather the equivalence relation $\dot{\bigcup}_{i\in I}(U_i \times U_i)$ on U, and complement $(-)$ is taken relative to this.

Following Tarski, we choose to pay this price. We close $\{\mathfrak{Re}(U) : U \text{ a set}\}$ under isomorphism, subalgebras, *and also products.* The class we obtain is denoted RRA, standing for *Representable Relation Algebras.* In fact we get the same class by closing first under products and then under subalgebras and isomorphism, so each algebra in RRA is isomorphic to an algebra of genuine binary relations. The price is worth paying because RRA turns out to be a variety [31], so we can in principle apply equational reasoning to it, as we did with boolean algebras. For discussion, see [30].

Before we compare RRA with what we arrived at in the unary case (BA), we should probably explain the term 'representable relation algebras' used here. But first, since the algebras in RRA are not necessarily concrete algebras of binary relations, it's inappropriate to use symbols such as $|$ that have a concrete meaning. So we introduce a special alphabet.[2] For historical reasons, its non-logical symbols are the binary function symbols $+$, $\cdot$, and ;, the unary function

[2]See [26, §2.1.1] for 'alphabet' — the non-logical part of it is sometimes called a signature, similarity type, or vocabulary.

symbols $-$ and $\breve{}$, and the distinguished elements 0, 1, and 1'. The binary function symbols are written in infix form ($a\,;b$, etc.), and $\breve{}$ is written in the form $\breve{a}$ or $a\breve{}$, as desired. We pronounce 1', $\breve{}$, and ; as *identity, converse,* and *composition,* respectively.

Algebras in RRA are taken to be structures for this alphabet. In algebras of the form $\mathfrak{Re}(U)$, the symbols are interpreted as follows:

$+$	as	$\cup$	$\cdot$	as	$\cap$	$-$	as	$-$
0	as	$\emptyset$	1	as	$U \times U$			
1'	as	Id_U	$\breve{-}$	as	$-^{-1}$	;	as	$\mid$

For example, $(\mathfrak{Re}(U))(+) = \cup$, $(\mathfrak{Re}(U))(1) = U \times U$, and so on — here we follow [26, p.46] and write $\mathfrak{M}(s)$ for the interpretation of a symbol s in a structure $\mathfrak{M}$. We will sometimes write an arbitrary algebra of this alphabet in the form $\mathcal{A} = (A, +, \cdot, -, 0, 1, 1', \breve{}, ;)$, abusing notation by identifying (notationally) the function symbols above with their interpretation as functions in the algebra.

We can now explain the term 'representable relation algebra'.

A plain *relation algebra* is an algebra $\mathcal{A} = (A, +, \cdot, -, 0, 1, 1', \breve{}, ;)$ of the above alphabet that satisfies a certain finite set of equations put forward in [16]. The equations say in effect that the boolean reduct $(A, +, \cdot, -, 0, 1)$ of $\mathcal{A}$ is a boolean algebra, and some other things[3] that we will not need here. Introductory surveys on relation algebras can be found in [24, 25].

A *representation* of a relation algebra $\mathcal{A}$ is an embedding from $\mathcal{A}$ into an algebra of the form $\prod_{i \in I} \mathfrak{Re}(U_i)$. It *represents* each element a of $\mathcal{A}$ as a binary relation $h(a)$ on $\dot{\bigcup}_{i \in I} U_i$. It is an isomorphism from $\mathcal{A}$ to a genuine algebra of binary relations.

A relation algebra is said to be *representable* if it has a representation. This holds iff the algebra is in RRA as defined above, so justifying the nomenclature.

Now let us compare RRA with BA. The equations from [16] defining relation algebras are chosen to hold in RRA. They are quite powerful and it was hoped for a while that they would in fact *define* RRA, or equivalently, that every relation algebra was representable. This would have put the theory of binary relations on a similar footing to that of unary ones, with relation algebras playing the role of boolean algebras.

The hope turned out to be in vain: Lyndon [21] showed that not every relation algebra is representable. Actually, RRA is quite difficult to capture. Tarski [31] proved it to be a variety by metamathematical means, not by pointing to a known equational axiomatisation of it, as we did with BA. An equational axiomatisation of RRA was given by Lyndon in [22], but it is infinite and complicated. In fact, Monk proved in [28] (corollary 7 below) that RRA is not finitely axiomatisable in first-order logic at all. Numerous facts along these lines are now known. They are often regarded as 'negative' results, but I prefer

[3] Namely, $(A, ;, 1')$ is a monoid, and the conditions $a \cdot (b\,;c) = 0$, $b \cdot (a\,;\breve{c}) = 0$, and $c \cdot (\breve{b}\,;a) = 0$ are equivalent for every $a, b, c \in \mathcal{A}$.

to think of them as illustrating the richness and elusiveness of RRA. Binary relations are very subtle.

In summary, then, binary relations are more complex than unary ones. The closure under isomorphism and subalgebras of the class $\mathcal{U}$ of §2.1 is the finitely axiomatisable variety BA. The closure of $\{\mathfrak{Re}(U) : U \text{ a set}\}$ under these operations is not a variety, so we close under products as well. The resulting class, RRA, is a variety, but is dissimilar to BA in other respects: for example, it is not finitely axiomatisable. The class of relation algebras is a finitely axiomatisable variety and is in many ways a better analogue of BA, but infinitely many further axioms are needed to capture RRA.

3 RRA is barely canonical

In spite of the many 'negative' results about it, RRA does have two striking similarities to BA. First, as we said, it is a variety. Second, it is canonical, under an extended definition of canonical extension due to Jónsson and Tarski [17]. In this immensely influential paper, the canonical extension $\mathcal{A}^\sigma$ of a relation algebra $\mathcal{A}$ (and many other kinds of algebra) was defined abstractly — up to isomorphism — but nowadays it is often given a concrete definition based again on $\wp(\mathcal{A}_+)$, the power set of the set $\mathcal{A}_+$ of ultrafilters of $\mathcal{A}$. The notion of an ultrafilter of $\mathcal{A}$ makes sense because the boolean reduct of $\mathcal{A}$ is a boolean algebra. The non-boolean operations 1', $\breve{}$, and ; are defined on $\wp(\mathcal{A}_+)$ in a special way.

Suppose that $\mathcal{A} \in$ RRA. By [18, theorem 4.21], $\mathcal{A}^\sigma$ is a relation algebra, but it is 'made of' unary relations on $\mathcal{A}_+$, not binary ones, and we cannot immediately conclude that $\mathcal{A}^\sigma$ is representable and so in RRA. Nevertheless, it does turn out that $\mathcal{A}^\sigma \in$ RRA. This was proved by Monk and reported in [27, p.66]; the first published proof is in [23]. It can also be proved by model-theoretic saturation [12, §3.4.4] (see also [7, 9, 8]), illustrating the value of another part of model theory in algebraic logic!

Hence, as we said, RRA is canonical.

At this point, let us interject a quick definition. A first-order sentence φ (for example, an equation) is said to be *canonical* if it is preserved by taking canonical extensions. That is, for any algebra $\mathcal{B}$ of the alphabet of φ and having a canonical extension $\mathcal{B}^\sigma$, if $\mathcal{B} \models \varphi$ then $\mathcal{B}^\sigma \models \varphi$. The study of canonical *equations* is extensive — see, e.g., [17, 5, 15].

Now since RRA is a canonical class, and, being a variety, is defined by equations, one might jump to the conclusion that it can be defined by canonical equations — that it has a *canonical equational axiomatisation.*

That would be most unwise. The true position is strikingly different. Not only does RRA have no canonical equational axiomatisation, but in fact, *every first-order axiomatisation of* RRA *contains infinitely many non-canonical* *sentences* [14]. This applies to equational axiomatisations as a special case.[4] The canonicity of RRA is elusive: it does not reside in the individual axioms defining RRA, but seems to have a quite different source, and emerges only in the limit when all axioms are taken together. So RRA is only 'barely' canonical, where we call a class *barely canonical* if it is canonical and elementary but every first-order axiomatisation of it involves infinitely many non-canonical sentences. I still find it remarkable that such classes exist.

The original proof in [14] that RRA is barely canonical used, among other things, first-order compactness. In the rest of this section, we will sketch a slightly simpler proof using the method in [2, 3]. The simplification is achieved by replacing the finite combinatorics in [14, propositions 6.4 & 6.6] by even more compactness! We will leave out some details from the presentation, for brevity.

3.1 Relation algebras from graphs

The proof uses graphs. According to [26, definition 5.37], a *graph* is a structure $\mathfrak{G} = (G, E)$, where E is an irreflexive and symmetric binary 'edge' relation on G. We will need the notion of an *independent subset of* $\mathfrak{G}$, which is a set $X \subseteq G$ such that for no $x, y \in X$ do we have $\mathfrak{G} \models E(x, y)$. In plain words, there are no edges in X.

Until the end of §3.3, we fix a graph $\mathfrak{G} = (G, E)$. We will write $G \times 3$ for the set $G \times \{0, 1, 2\}$, and $\mathfrak{G} \times 3$ for the graph $(G \times 3, E')$, where for $a, b \in G$ and $i, j \in \{0, 1, 2\}$ we define $E'((a, i), (b, j))$ iff $E(a, b)$ or $i \neq j$. In simple words, $\mathfrak{G} \times 3$ consists of three disjoint copies of $\mathfrak{G}$, with all possible edges added between the copies.

It is possible to construct a certain relation algebra on top of $\mathfrak{G} \times 3$. This algebra was introduced in joint work with Robin Hirsch [13, §4] and has the form

$$\mathcal{A}(\mathfrak{G}) = \big(\wp\big((G \times 3) \cup \{e\}\big), +, \cdot, -, 0, 1, 1\text{'}, \breve{}\,, ;\big).$$

Here, e is a new 'identity' element not in $G \times 3$. The boolean operations $+$, $\cdot$, $-$, 0, and 1 are interpreted as usual in power sets; we interpret 1' as $\{e\}$; and we set $\breve{a} = a$ for all a.

The interpretation of ; in $\mathcal{A}(\mathfrak{G})$ takes a little longer to describe, but the reader may like to see it because it is where the graph structure of $\mathfrak{G} \times 3$ comes in. Let us say that a triple (x, y, z) of elements of $(G \times 3) \cup \{e\}$ is *consistent* if

- one of x, y, z is e and the other two are equal, or
- $e \notin \{x, y, z\}$ and $\{x, y, z\}$ is *not* an independent subset of $\mathfrak{G} \times 3$.

Then, for $X, Y \subseteq (G \times 3) \cup \{e\}$, we define

$$X \,; Y = \{z \in (G \times 3) \cup \{e\} : (x, y, z) \text{ is consistent for some } x \in X,\ y \in Y\}.$$

[4] A harbinger of this result, that RRA has no Sahlqvist axiomatisation, is in [32].

This definition has the following useful consequence:

Lemma 1 *Let $X \subseteq G \times 3$. Then X is an independent subset of $\mathfrak{G} \times 3$ iff $(X\,;X) \cdot X = 0$ in $\mathcal{A}(\mathfrak{G})$.*

Proof. We use that $e \notin X$ without explicit mention.

If X is not independent, pick $x, y \in X$ with $\mathfrak{G} \times 3 \models E'(x, y)$. As $\{x, y\}$ is not independent, (x, y, y) is consistent, so $y \in X\,;X$. Hence, $(X\,;X) \cdot X \neq 0$.

Conversely, if $(X\,;X) \cdot X \neq 0$, pick $z \in (X\,;X) \cap X$. By definition of ;, there are $x, y \in X$ such that (x, y, z) is consistent. So $\{x, y, z\}$ is not an independent subset of $\mathfrak{G} \times 3$. So X cannot be independent either, since $\{x, y, z\} \subseteq X$. $\square$

It turns out that $\mathcal{A}(\mathfrak{G})$ is indeed a relation algebra [14, lemma 6.2]. The number 3 and the extra edges added between the three copies of $\mathfrak{G}$ are to ensure that ; is associative and to establish (†) in §3.3 below.

3.2 Graph colourings seen in $\mathcal{A}(\mathfrak{G})$

For a positive integer m, we say that that $\mathfrak{G}$ *can be coloured with m colours* if G is the union of m (possibly empty) independent sets. This is equivalent to the standard definition, as in [26, definition 5.38]. We say that $\mathfrak{G}$ *can be finitely coloured* if it can be coloured with m colours for some finite m.

We are going to see (in lemma 3) that colourings of $\mathfrak{G}$ are visible in $\mathcal{A}(\mathfrak{G})$. First, we need some definitions. The sets $G \times \{l\}$ (the domains of our three copies of $\mathfrak{G}$) are of course in $\mathcal{A}(\mathfrak{G})$, for each $l = 0, 1, 2$, and it is convenient to consider them as additional distinguished elements b_l of $\mathcal{A}(\mathfrak{G})$, thereby expanding the alphabet of relation algebras a little. We have, for example, $\mathcal{A}(\mathfrak{G}) \models b_0 + b_1 + b_2 + 1' = 1$.

Definition 2 *Using this expanded alphabet, we define the following.*

1. *For each integer $m \geq 1$ and $l = 0, 1, 2$, define the following universal first-order sentence:*

$$\theta^l_m = \forall x_1, \ldots, x_m \Big((x_1 + \cdots + x_m = b_l) \rightarrow \bigvee_{i=1}^{m} [(x_i\,; x_i) \cdot x_i \neq 0] \Big).$$

   Here, $x_1 + \cdots + x_m$ is defined as $((\cdots(x_1 + x_2) + x_3) + \cdots) + x_m$, but actually we will use it only when $+$ is associative.

2. *We will write θ^0_m simply as θ_m.*

3. *Let Θ be the universal first-order theory $\{\theta_m : m \geq 1\}$.*

Lemma 3 1. *Let $m \geq 1$ and $0 \leq l \leq 2$. Then $\mathcal{A}(\mathfrak{G}) \models \theta^l_m$ iff $\mathfrak{G}$ cannot be coloured with m colours.*

2. *$\mathcal{A}(\mathfrak{G}) \models \Theta$ iff $\mathfrak{G}$ cannot be finitely coloured.*

Proof. First, $\mathfrak{G}$ is plainly isomorphic to the induced subgraph of $\mathfrak{G} \times 3$ with domain $G \times \{l\}$ — one of our three copies of $\mathfrak{G}$. The isomorphism is $(x \mapsto (x, l))$. So $\mathfrak{G}$ can be coloured with m colours iff this subgraph can be. By lemma 1, the latter holds iff there are $X_1, \ldots, X_m$ in $\mathcal{A}(\mathfrak{G})$ with $X_1 + \cdots + X_m = b_l$ and $(X_i \,;X_i) \cdot X_i = 0$ for each $i = 1, \ldots, m$. This is plainly iff $\mathcal{A}(\mathfrak{G}) \models \neg\theta^l_m$. That proves (1), and (2) follows from (1) and the definition of Θ. □

3.3 Graph colourings and representability of $\mathcal{A}(\mathfrak{G})$

Colourings of $\mathfrak{G}$ are relevant because they are connected to representability of $\mathcal{A}(\mathfrak{G})$. We have:

(†) $\mathfrak{G}$ cannot be finitely coloured iff $\mathcal{A}(\mathfrak{G})$ is infinite and representable.

For the '⇒' direction, if $\mathfrak{G}$ cannot be finitely coloured then obviously $\mathfrak{G}$ and hence $\mathcal{A}(\mathfrak{G})$ are infinite; representability of $\mathcal{A}(\mathfrak{G})$ can be proved by actually building a representation of it, for example using games. The '⇐' direction can be proved using lemma 1 and Ramsey's theorem. For details, see [12, theorems 14.12–14.13].

Those who see representability as 'good' and non-finite colourability as 'bad' may find (†) surprising, but it is not really counter-intuitive. Representability is defined by universal sentences (since RRA is a variety); and by lemma 3, $\mathfrak{G}$ cannot be finitely coloured iff $\mathcal{A}(\mathfrak{G}) \models \Theta$, and Θ also consists of universal sentences. On both sides, the sentences express the absence of certain (bad or good) elements in the algebra.

3.4 Extending (†) to an elementary class of algebras

The algebras $\mathcal{A}(\mathfrak{G})$ do not form an elementary class. To see this, we could observe that no algebra $\mathcal{A}(\mathfrak{G})$ is countably infinite (because all power sets are finite or uncountable), whereas by the downward Löwenheim–Skolem theorem, every elementary class with countable alphabet and containing infinite structures must contain some countably infinite ones.

To apply compactness, we need to extend (†) above to an elementary class of relation algebras.

Definition 4 *Let U be the first-order theory comprising:*

1. *the $\forall$-theory of the class $\{\mathcal{A}(\mathfrak{G}) : \mathfrak{G}$ a graph$\}$ — that is, the set of all universal sentences of our expanded alphabet that are true in every algebra $\mathcal{A}(\mathfrak{G})$,*
2. *all sentences $\theta^l_m \to \theta^{l'}_{m'}$, for $1 \le m' \le m$ and $0 \le l, l' \le 2$ (see definition 2).*

Each sentence in definition 4(2) is equivalent to an ∀∃ sentence, so overall we can take U to be an ∀∃ theory.

Our desired elementary class will be the class of models of U.[5] We will need that this class does indeed include all algebras $\mathcal{A}(\mathfrak{G})$:

Lemma 5 $\mathcal{A}(\mathfrak{G}) \models U$ *for every graph* $\mathfrak{G}$.

Proof. The sentences of definition 4(1) are obviously true in $\mathcal{A}(\mathfrak{G})$. So consider an arbitrary sentence $\theta^l_m \to \theta^{l'}_{m'}$ as in definition 4(2), and suppose that $\mathcal{A}(\mathfrak{G}) \models \theta^l_m$. By lemma 3, $\mathfrak{G}$ cannot be coloured with m colours. But $m' \leq m$, so $\mathfrak{G}$ cannot be coloured with m' colours either. By lemma 3 again, $\mathcal{A}(\mathfrak{G}) \models \theta^{l'}_{m'}$. We conclude that $\mathcal{A}(\mathfrak{G}) \models \theta^l_m \to \theta^{l'}_{m'}$. Hence, $\mathcal{A}(\mathfrak{G}) \models U$ as required. □

Using lemma 3 and methods in [3, theorems 7.3 & 7.8], (†) can be extended to models of U. For every $\mathcal{A} \models U$, we have:

U1. $\mathcal{A}$ is a relation algebra (because the equations defining relation algebras are universal sentences true in every $\mathcal{A}(\mathfrak{G})$, and so are in U),

U2. if $\mathcal{A} \models \Theta$ then $\mathcal{A}$ is representable (and infinite, but we will not need this),

U3. if $\mathcal{A}$ is infinite and representable, then $\mathcal{A} \models \Theta$.

U1 ensures that U2 and U3 make sense, and they are proved in much the same way as (†), because the main facts needed to prove (†) are included in U.

3.5 Sketch proof that RRA is barely canonical

Now we can present our main theorem. Its statement may seem unrelated to compactness, but the proof uses compactness heavily.

Theorem 6 *Every first-order axiomatisation of* RRA *has infinitely many non-canonical sentences.*

Proof (sketch). Let Θ and U be as in definitions 2 and 4, and let Φ be a first-order theory stating that its models are infinite (see [26, example 5.5]). Suppose for contradiction that RRA is defined by a first-order theory

$$T = T_C \cup T_{NC},$$

where T_C is a set of canonical sentences and T_{NC} is finite. We do not require that T consists of equations. We now have the following facts:

F1. $U \cup \Theta \models T$ (by U2)

F2. $U \cup \Phi \cup T \models \Theta$ (by U3)

[5] Actually, these models are not too different from the $\mathcal{A}(\mathfrak{G})$. Up to isomorphism, they are the subalgebras of algebras $\mathcal{A}(\mathfrak{G})$ that satisfy the sentences of definition 4(2).

F3. $U \cup \{\theta_n\} \models \{\theta_1, \dots, \theta_n\}$ for each $n \geq 1$ (since $\theta_n \to \theta_m \in U$ whenever $m \leq n$).

From them, we can make the following successive deductions:

D1. By F1, F3, and compactness, there is $l \geq 1$ such that $U \cup \{\theta_l\} \models T_{NC}$.

D2. By F2 and compactness, there is a finite subset $S \subseteq T_C$ such that $U \cup \Phi \cup S \cup T_{NC} \models \theta_{l+1}$.

D3. By F1, F3, and compactness, there is $k \geq l$ with $U \cup \{\theta_k\} \models S$.

Now an adaptation given in [14, lemma 4.1] of a famous probabilistic graph construction by Erdős [6] yields an inverse system

$$\mathfrak{G}_0 \leftarrow \mathfrak{G}_1 \leftarrow \mathfrak{G}_2 \leftarrow \cdots$$

of finite graphs $\mathfrak{G}_n$ that cannot be coloured with k colours, and surjective p-morphisms[6] connecting them, whose inverse limit is an infinite graph $\mathfrak{G}$ (say) that can be coloured with $l+1$ but not with l colours.[7]

From this inverse system, duality results of Goldblatt [10, §§1.10–1.11] give us a direct system

$$\mathcal{A}(\mathfrak{G}_0) \to \mathcal{A}(\mathfrak{G}_1) \to \mathcal{A}(\mathfrak{G}_2) \to \cdots$$

of finite relation algebras and embeddings, whose direct limit $\mathcal{A}$ (say) wonderfully satisfies

$$\mathcal{A}^\sigma \cong \mathcal{A}(\mathfrak{G}). \tag{1}$$

Here, $\cong$ denotes isomorphism of algebras. See [14, §6.4] for details.

Let's see what we can get from this. Since the $\mathfrak{G}_n$ cannot be coloured with k colours, lemmas 5 and 3 give $\mathcal{A}(\mathfrak{G}_n) \models U \cup \{\theta_k\}$ for each n. This theory is $\forall\exists$, so preserved by direct limits (well known and an easy exercise). Hence, $\mathcal{A} \models U \cup \{\theta_k\}$ as well. By D3, we obtain $\mathcal{A} \models S$.

Crucially, the sentences in S are canonical, so $\mathcal{A}^\sigma \models S$. By (1), we obtain $\mathcal{A}(\mathfrak{G}) \models S$.

Now $\mathfrak{G}$ cannot be coloured with l colours, so by lemmas 5 and 3 again, $\mathcal{A}(\mathfrak{G}) \models U \cup \{\theta_l\}$. So by D1, $\mathcal{A}(\mathfrak{G}) \models T_{NC}$. Also, $\mathfrak{G}$ is infinite, so $\mathcal{A}(\mathfrak{G}) \models \Phi$. We have arrived at $\mathcal{A}(\mathfrak{G}) \models U \cup \Phi \cup S \cup T_{NC}$. Now, D2 yields $\mathcal{A}(\mathfrak{G}) \models \theta_{l+1}$, which contradicts lemma 3 since $\mathfrak{G}$ *can* be coloured with $l+1$ colours. This contradiction shows that our assumption is false, and that RRA cannot be defined by a first-order theory containing only finitely many non-canonical sentences. □

[6] These are like homomorphisms but a bit stronger: see, e.g., [4, p.30].

[7] This is the 'combinatorial heart' of the proof, but unfortunately we cannot say more about it here. A proof can also be developed from [11]. For direct and inverse systems and their limits, see [10, §1.11] or standard algebra texts.

Although the proof goes by contradiction, for any given k, l the relation algebras $\mathcal{A}(\mathfrak{G}_n)$, $\mathcal{A}$, and $\mathcal{A}(\mathfrak{G})$ are real, and strikingly, it can be arranged that none of them are representable.

The following is immediate.

Corollary 7 (Monk, [28]) RRA *is not finitely axiomatisable in first-order logic.*

In [2, 3], analogous results were proved for other kinds of algebras of n-dimensional relations for $n \geq 3$, including cylindric algebras, diagonal-free cylindric algebras, polyadic algebras, and polyadic equality algebras. The proofs are similar to the above, but more technical because of the intricacies of higher-arity relations.

One final, 'political' comment. I have been accused of dancing on Tarski's grave by proving (with co-authors) things like theorem 6. In fact, this theorem is charting the limits of the very popular concept of canonicity, and RRA is an excellent tool for that. Moreover, the ideas have spread: bare canonicity plays a role in the very striking dichotomy for modal logics developed by Kikot and Zolin [19, 20]. In [19, p.1066], Kikot remarks that 'what was thought pathological [bare canonicity] can now be seen to be the norm'.

4 Acknowledgement

I would like to finish by recording my debt and gratitude to the co-authors Jannis Bulian and Yde Venema of the work I have described; to Rob Goldblatt, for a very valuable remark on the text; to Antonia, for preparing this volume; and lastly to Mara, most especially for her forbearance with me when I was younger. I wish her a delightfully happy and peaceful retirement.

References

[1] G. Birkhoff, *On the structure of abstract algebras*, Proc. Cambr. Philos. Soc. **31** (1935), 433–454.

[2] J. Bulian, *Exploring canonical axiomatisations of representable cylindric algebras*, 2011, final-year individual project report, Department of Computing, Imperial College London, `http://www.doc.ic.ac.uk/teaching/distinguished-projects/2011/j.bulian.pdf`.

[3] J. Bulian and I. Hodkinson, *Bare canonicity of representable cylindric and polyadic algebras*, Ann. Pure. Appl. Logic **164** (2013), 884–906.

[4] A. Chagrov and M. Zakharyaschev, *Modal logic*, Oxford Logic Guides, vol. 35, Clarendon Press, Oxford, 1997.

[5] M. de Rijke and Y. Venema, *Sahlqvist's theorem for boolean algebras with operators*, Studia Logica **54** (1995), 61–78.

[6] P. Erdős, *Graph theory and probability*, Canad. J. Math. **11** (1959), 34–38.

[7] K. Fine, *Some connections between elementary and modal logic*, Proc. 3rd Scandinavian logic symposium, Uppsala, 1973 (S. Kanger, ed.), North Holland, Amsterdam, 1975, pp. 15–31.

[8] M. Gehrke, J. Harding, and Y. Venema, *MacNeille completions and canonical extensions*, Trans. Amer. Math. Soc. **358** (2006), 573–590.

[9] R. Goldblatt, *Varieties of complex algebras*, Ann. Pure. Appl. Logic **44** (1989), 173–242.

[10] ———, *Mathematics of modality*, Lecture notes, vol. 43, CSLI Publications, Stanford, CA, 1993.

[11] P. Hell and J. Nešetřil, *The core of a graph*, Discrete Mathematics **109** (1992), 117–126.

[12] R. Hirsch and I. Hodkinson, *Relation algebras by games*, Studies in Logic and the Foundations of Mathematics, vol. 147, North-Holland, Amsterdam, 2002.

[13] ———, *Strongly representable atom structures of relation algebras*, Proc. Amer. Math. Soc. **130** (2002), 1819–1831.

[14] I. Hodkinson and Y. Venema, *Canonical varieties with no canonical axiomatisation*, Trans. Amer. Math. Soc. **357** (2005), 4579–4605.

[15] B. Jónsson, *On the canonicity of Sahlqvist identities*, Studia Logica **53** (1994), 473–491.

[16] B. Jónsson and A. Tarski, *Representation problems for relation algebras*, Bull. Amer. Math. Soc. **54** (1948), 80, 1192.

[17] ———, *Boolean algebras with operators I*, American Journal of Mathematics **73** (1951), 891–939.

[18] ———, *Boolean algebras with operators II*, American Journal of Mathematics **74** (1952), 127–162.

[19] S. Kikot, *A dichotomy for some elementary generated modal logics*, Studia Logica **103** (2015), 1063–1093.

[20] S. Kikot and E. Zolin, *Modal definability of first-order formulas with free variables and query answering*, Journal of Applied Logic **11** (2013), 190–216.

[21] R. Lyndon, *The representation of relational algebras*, Annals of Mathematics **51** (1950), 707–729.

[22] ———, *The representation of relation algebras, II*, Annals of Mathematics **63** (1956), 294–307.

[23] R. D. Maddux, *A sequent calculus for relation algebras*, Ann. Pure. Appl. Logic **25** (1983), 73–101.

[24] ———, *Introductory course on relation algebras, finite-dimensional cylindric algebras, and their interconnections*, Algebraic logic (Amsterdam) (H. Andréka, J. D. Monk, and I. Németi, eds.), Colloq. Math. Soc. J. Bolyai, vol. 54, North-Holland, 1991, pp. 361–392.

[25] ———, *Subcompletions of representable relation algebras*, Algebra Universalis (2018), 79: 20. `https://doi.org/10.1007/s00012-018-0493-0`.

[26] M. Manzano, *Model theory*, Oxford Logic Guides, vol. 37, Clarendon Press, Oxford, 1999.

[27] R. McKenzie, *The representation of relation algebras*, Ph.D. thesis, University of Colorado at Boulder, 1966.

[28] J. D. Monk, *On representable relation algebras*, Michigan Mathematics Journal **11** (1964), 207–210.

[29] M. H. Stone, *The theory of representations for boolean algebras*, Trans. Amer. Math. Soc. **40** (1936), 37–111.

[30] A. Tarski, *On the calculus of relations*, J. Symbolic Logic **6** (1941), 73–89.

[31] ———, *Contributions to the theory of models, III*, Koninkl. Nederl. Akad. Wetensch Proc. **58 (= Indag. Math. 17)** (1955), 56–64.

[32] Y. Venema, *Atom structures and Sahlqvist equations*, Algebra Universalis **38** (1997), 185–199.