

RECURSIVE AXIOMATISATIONS FROM SEPARATION PROPERTIES

ROB EGROT

Abstract. We define a fragment of monadic infinitary second-order logic corresponding to an abstract separation property. We use this to define the concept of a *separation subclass*. We use model theoretic techniques and games to show that separation subclasses whose axiomatisations are recursively enumerable in our second-order fragment can also be recursively axiomatised in their original first-order language. We pin down the expressive power of this formalism with respect to first-order logic, and investigate some questions relating to decidability and computational complexity. As applications of these results, by showing that certain classes can be straightforwardly defined as separation subclasses, we obtain first-order axiomatisability results for these classes. In particular we apply this technique to graph colourings and a class of partial algebras arising from separation logic.

§1. Introduction. We begin with a motivating example. Precise definitions will be given in the next section. A partially ordered set (poset) is *representable* if it can be embedded into a powerset algebra via a map that preserves existing finite meets and joins. The class of representable posets and its infinitary variations have been studied, not always using this terminology, in [8, 11–15, 21, 29, 39], generalising work done in the setting of semilattices [2, 9, 27, 33], and for distributive lattices and Boolean algebras [1, 3, 4, 6, 7, 16, 31, 35, 36]. At first glance, it is far from obvious that the class of representable posets is elementary. However, it is fairly easy to show that a poset is representable if and only if it has a ‘separating’ set of ‘prime filters’. More precisely, a poset P is representable if and only if whenever $p \not\leq q \in P$ there is a ‘prime filter’ of P containing p and not q . Note that there are several non-equivalent concepts of ‘prime filter of a poset’ in circulation, and we are using one in particular. A more precise definition is given in Example 2.4.

Now, given the description of the class of representable posets in terms of this ‘separation property’, it is possible to show that it can in fact be axiomatised in first-order logic. [11, Theorem 4.5] does this by proving closure under taking isomorphisms, ultraproducts, and ultraroots and appealing to the Keisler–Shelah theorem [28, 34], and similar can be done by proving closure under taking ultraproducts and elementary substructures and appealing to [20, Theorem 2.13]. Such a non-constructive proof of existence may be regarded as being of limited practical use, however, the very fact that an axiomatisation is known to exist can be used in a neat trick to show that a certain constructively generated axiomatisation is correct. This is the main result of [15].

Received August 29, 2019.

2020 *Mathematics Subject Classification*. 03C98, 03B16, 03B70, 05C15, 08A55.

Key words and phrases. first-order axiomatisability, generating recursive axiomatisations, representable posets, graph colourings, harmonious colourings, disjoint union partial algebras

© 2021, Association for Symbolic Logic
0022-4812/21/8603-0015
DOI:10.1017/jsl.2021.19

The method of [15], which is not novel, is to describe the ‘separation property’ of representable posets in terms of a game played between two players. The game is defined so that the number of rounds a certain player can survive in a particular game corresponds, in a sense, to how close a given poset is to being representable. First-order axioms are then written down that correspond to the player ‘having a strategy’ in a game. These axioms are shown to correctly axiomatise the class of representable posets by means of the ‘neat trick’ mentioned previously.

A similar idea appears in [24], where it is used to find an explicit axiomatisation for a certain class of partial algebras of partial functions that appears in connection with separation logic. Again we have a class which is not obviously elementary, but which can fairly easily be shown to be definable in terms of a ‘separation property’. The separation property is then used to show, non-constructively, that a first-order axiomatisation exists, and then to construct explicit axioms based on games which are, using the ‘neat trick’, shown to be an axiomatisation for the class.

The main purpose of this paper is to prove a general theorem that includes the relevant results of [15, 24] as special cases, and is also applicable in a wide variety of other situations. The strategy is to first formalise the concept of a ‘separation property’ in a way that allows the necessary results to go through, while also being intuitive enough to be useful in practice. This is done in Section 2. In particular, the basic definition of a *separation subclass* is made. The sense in which separation subclasses can be, for example, *essentially countable*, or *essentially recursively enumerable*, is also explained.

We formalise the concept of a separation subclass using infinitary monadic second-order logic. We show that if \mathcal{A} is a class of structures and \mathcal{B} is a subclass of \mathcal{A} that is elementary relative to \mathcal{A} , then \mathcal{B} can always be described as separation subclass of \mathcal{A} (Proposition 2.7). More interestingly, we show that every separation subclass of an elementary class has a first-order axiomatisation relative to the superclass (Theorem 2.14). Thus the classes of separation subclasses and elementary subclasses of an elementary class coincide. However, the important difference is that descriptions as separation subclasses can often be much easier to find than first-order axiomatisations. Moreover, as we shall see, provided the superclass is elementary, we can use a description of a subclass as an essentially recursively enumerable separation subclass to automate the construction of an explicit first-order axiomatisation.

In Section 3 we describe a class of games played between two players, \forall and \exists . The key result is that, if \mathcal{B} is an essentially countable separation subclass of \mathcal{A} , then given $A \in \mathcal{B}$, the player \exists has a strategy for never losing in every relevant game. Conversely, if $A \in \mathcal{A}$ is countable, then \exists having such strategies implies that $A \in \mathcal{B}$ (Proposition 3.1).

Section 4 formalises the existence of strategies for \exists in first-order logic. The main result, which is stated as Corollary 4.6, is that an essentially recursively enumerable separation subclass \mathcal{B} of an elementary class \mathcal{A} always has a recursive first-order axiomatisation relative to \mathcal{A} , which we can generate systematically by examining the relevant class of games. Moreover, we present simple sufficient conditions for the axiomatisation produced to be universal.

In Section 5 we collect together some previous results to make explicit the connections between the various kinds of separation subclasses and the various ways a class can be elementary relative to its superclass (Proposition 5.2). We also

make some simple observations regarding decision problems and complexity (Propositions 5.4 and 5.6).

Finally, in Section 6 we present some applications of the general theory we have developed. First we show how the work in [24] on disjoint union partial algebras fits into the framework of separation subclasses, and how this automatically proves some of the results of that paper (Section 6.1). Following this we consider graph colourings. In particular, in Section 6.2, from the fact that the class of N -colourable graphs has a simple description as a separation subclass of the class of all graphs, we are able to find easy proofs of several model theoretic results relating to these structures. We present new proofs of the known results that, for all $N \geq 2$, the class of N -colourable graphs has a universal Horn axiomatisation, but is not finitely axiomatisable, and also that, when $N \geq 3$, the class of graphs with chromatic number N is not elementary. The proofs follow directly from the general results on separation subclasses. In this sense, N -colourable graphs provide a good example of a class where a characterisation as a separation subclass is obvious, but where results relating to first-order axiomatisability are not so obvious.

In Sections 6.3 and 6.4 we describe the classes of graphs with N -clique covers, and harmonious N -colourings, respectively, as separation subclasses. Thus, as an immediate consequence, we can show that both classes have recursive universal axiomatisations relative to the class of all graphs. Moreover, our method proves that the class of graphs with harmonious N -colourings is actually finitely axiomatisable for each $N \geq 1$.

§2. Separation subclasses. We adopt the convention that indexing sets are denoted by capital letters, and arbitrary indices taken from these sets use the corresponding lowercase letters. When dealing with a set $\{x_1, \dots, x_N\}$, we will use x_n to denote an arbitrary element from this set. For variable symbols, we adopt the convention that, e.g., \vec{x}_N denotes the set $\{x_1, \dots, x_N\}$. Given a first-order formula ϕ and a set of variable symbols \vec{x}_N , it is common to write something like $\phi(\vec{x}_N)$ to denote that the free variables of ϕ are precisely \vec{x}_N . In this paper we use a relaxed version of this convention. Here we will write, e.g., $\phi(\vec{x}_N)$ to denote that the free variables of ϕ are from among \vec{x}_N , but are not necessarily all of them. We will write things like $\forall \vec{x}_N \phi(\vec{x}_N)$ to stand for $\forall x_1 \dots x_N \phi(\vec{x}_N)$. Given a set of variables X we may also write, e.g., $\phi(X)$ to denote that the free variables of ϕ are from X .

We will also use the $\bar{}$ notation in a different but closely related way as follows. If, e.g., b_1, \dots, b_N are not necessarily distinct elements of some \mathcal{L} -structure \mathcal{B} , we will write $\phi(\bar{b}_N)$ to mean $\phi(\vec{x}_N)$ where each free x_n is interpreted as b_n for all $n \in \{1, \dots, N\}$. We will sometimes write something like $\bar{b}_N \in \mathcal{B}$ to denote a sequence b_1, \dots, b_N of elements of a structure \mathcal{B} .

Before diving further into the technicalities, we will try to build up some intuition for what is to be done. The situation to be captured is as follows. We have some kind of a structure, e.g., a poset P , and we want to say that given some elements of this structure collectively meeting some condition that is definable in first-order logic, e.g., $p, q \in P$ with $p \not\leq q$, there are subsets of the original structure which can be specified to either contain or not contain the elements in question. For example, given $p \not\leq q \in P$ we may want to demand that there is a prime filter of P containing p and not q . Moreover, these sets have to satisfy first-order closure

conditions. For example, prime filters must be closed upwards, among other things. It is convenient to represent these sets using new monadic predicate symbols. We will ultimately want to existentially quantify over these predicate symbols, so we consider them to be second-order variables. If C is the monadic predicate standing for the prime filter in our poset example, then upward closure can be expressed using the sentence

$$\forall x \forall y ((x \leq y \wedge C(x)) \rightarrow C(y)). \tag{\dagger}$$

The idea is that the relationship between elements, in this case picked out by the variables x and y , puts some constraint on their containment or otherwise in the set, in this case picked out by the second-order variable C . It will be useful to rewrite this so that the purely first-order part, i.e., the part not involving C , is the antecedent of an implication, with the part involving C being the consequent. In other words, to clearly reflect the fact that the first-order relationship between the elements places constraints on the placements of the variables in the sets represented by the second-order variables. These constraints on set containment can be expressed by Boolean combinations of clauses demanding that particular elements either are or are not included in a particular set. From (\dagger) , for example, we can obtain the equivalent formula

$$\forall x \forall y (x \leq y \rightarrow (\neg C(x) \vee C(y))). \tag{\ddagger}$$

What we call a *closure rule*, and define explicitly in Definition 2.1 shortly, is just a conjunction of constraints of this form for some given sets. For example, we describe the closure properties of prime filters by conjoining (\ddagger) with an infinite number of other sentences of similar form (see Example 2.4 below for the details).

In Definition 2.2 we will formally introduce what we call a *separation rule*. The intuition here is that a separation rule expresses that if a set of elements from the structure meet some first-order definable condition, then there exist some subsets of the structure such that 1) there are constraints on the containment of the original elements in these sets, and 2) these sets satisfy a closure rule. Given a poset P , for example, if given a monadic predicate C we use $C[P]$ to denote $\{p \in P : C(p)\}$, the idea that given $p \not\leq q \in P$ there is a prime filter containing p and not q can be expressed as a separation rule as follows:

$$\forall p, q (p \not\leq q \rightarrow \exists C (C(p) \wedge \neg C(q) \wedge \text{‘closure rule expressing that } C[P] \text{ is a prime filter’})).$$

DEFINITION 2.1. Let \mathcal{L} be a first-order signature, let $1 \leq K \in \omega$, and let C_1, \dots, C_K be unary predicate symbols not appearing in \mathcal{L} . Define $\mathcal{L}_{C_K}^+ = \mathcal{L} \cup \{C_1, \dots, C_K\}$.

A \vec{C}_K -closure rule is a conjunction $\bigwedge_I \tau_i$ for some $I \neq \emptyset$, where for each $i \in I$, the formula τ_i is an $\mathcal{L}_{C_K}^+$ -sentence of form

$$\forall \vec{y}_{M_i} (\gamma(\vec{y}_{M_i}) \rightarrow \psi(\vec{y}_{M_i})),$$

where γ is a first-order \mathcal{L} -formula with free variables taken from $\vec{y}_{M_i} = \{y_1, \dots, y_{M_i}\}$, and ψ is a quantifier-free first-order $\mathcal{L}_{C_K}^+$ -formula whose free variables are also taken from $\{y_1, \dots, y_{M_i}\}$. Note that I may be infinite.

DEFINITION 2.2. A *separation rule* for a first-order signature \mathcal{L} is anything falling into either of the following two categories:

- (1) \mathcal{L} -sentences.
- (2) Monadic second-order sentences of form

$$\forall \vec{x}_N \left(\mu(\vec{x}_N) \rightarrow \exists \vec{C}_K (\eta(\vec{x}_N) \wedge \tau) \right),$$

where $1 \leq K < \omega$, where μ is an \mathcal{L} -formula with free variables taken from $\vec{x}_N = \{x_1, \dots, x_N\}$, where η is a quantifier-free $\mathcal{L}_{\vec{C}_K}^+$ -formula whose free variables are also taken from \vec{x}_N , and where τ is either a \vec{C}_K -closure rule or the tautology \top .

So, in the case where $\tau \neq \top$, a separation rule has form

$$\forall \vec{x}_N \left(\mu(\vec{x}_N) \rightarrow \exists \vec{C}_K (\eta(\vec{x}_N) \wedge \bigwedge_I \forall \vec{y}_{M_i} (\gamma_i(\vec{y}_{M_i}) \rightarrow \psi_i(\vec{y}_{M_i}))) \right)$$

for some set $I \neq \emptyset$ (which may be infinite).

The *order* of a separation rule of type (1) is said to be zero, and the order of a separation rule of type (2) is the value of K used in its definition.

A separation rule σ is said to be *finite* if it either has order zero, or if the order is positive with $\sigma = \forall \vec{x}_N (\mu(\vec{x}_N) \rightarrow \exists \vec{C}_K (\eta(\vec{x}_N) \wedge \bigwedge_I \tau_i))$ and I is finite. A separation rule that is not finite is said to be *infinite*. Similarly, σ is *countable* if it is either finite or I is countable. Finally, σ is said to be *recursively enumerable* (r.e.) if it is either finite or there is an algorithm for listing the formulas τ_i . In other words, while infinite separation rules cannot be written out in full, if they are r.e. they can at least be approximated to arbitrary precision by including more of the infinite conjunction $\bigwedge_I \tau_i$.

A set Σ of separation rules is called a *separation scheme*. Informally, a separation scheme is just a set of constraints on a structure that can be expressed as separation rules. A separation scheme Σ is said to be *essentially finite* if Σ is finite and each $\sigma \in \Sigma$ is also finite. Similarly, Σ is *essentially countable* if it is countable and each $\sigma \in \Sigma$ is countable. A separation scheme Σ is said to be *essentially recursively enumerable* if $\Sigma = \{\sigma_n : n \in \omega\}$, and σ_n is r.e. for each $n \in \omega$.

DEFINITION 2.3 ($\Sigma^{>0}$). If Σ is a separation scheme, we use $\Sigma^{>0}$ to denote the subset of Σ containing all the separation rules of order strictly greater than zero (i.e., all those of type (2)).

EXAMPLE 2.4. As mentioned previously, there are several non-equivalent ways to generalise the concept of a prime filter from lattices to posets. For us, a ‘prime filter’ of a poset P is a subset Γ of P that is closed upwards, closed under finite existing meets (greatest lower bounds), and has the ‘primality’ property that if the join (least upper bound) of a finite set S of elements of P is defined in P and is in Γ , then $\Gamma \cap S \neq \emptyset$. We will in future refer to a subset of a poset satisfying these closure properties as an ω -filter, to avoid ambiguity. We want to phrase the condition that, given a poset P , if $p \not\leq q \in P$, then there is an ω -filter of P containing p but not q as a separation rule. We proceed as follows.

Let $\mathcal{L} = \{\leq\}$ be the signature of ordered sets. For each $M \geq 1$ let $\vec{y}_M = \{y_1, \dots, y_M\}$ be a set of variable symbols, and, with a new variable symbol z , define $j_M(\vec{y}_M, z)$ and $m_M(\vec{y}_M, z)$ to be the universal \mathcal{L} -formulas stating that z is the least upper bound (join) and greatest lower bound (meet) of the elements of \vec{y}_M respectively.

We define a closure rule to capture the closure properties of ω -filters. To do this we introduce a unary predicate C meant to represent an ω -filter. First we define a clause τ_0 meant to capture upward closure:

$$\tau_0 = \forall yz (y \leq z \rightarrow (C(y) \rightarrow C(z))).$$

Now we define a clause τ_i for each $1 \leq i < \omega$ meant to capture closure under finite meets and the ‘primality’ property. We will use even values of i to capture closure under meets of the various finite cardinalities, and we will use odd values to capture the ‘primality’ properties. This division is purely an accounting device, but it is convenient.

If $i = 2M$ for some $M \in \omega$, then

$$\tau_i = \forall \vec{y}_M z \left(m_M(\vec{y}_M, z) \rightarrow \left(\bigwedge_{m=1}^M C(y_m) \rightarrow C(z) \right) \right),$$

and, if $i = 2M - 1$ for some $M \in \omega$, then

$$\tau_i = \forall \vec{y}_M z \left(j_M(\vec{y}_M, z) \rightarrow \left(C(z) \rightarrow \bigvee_{m=1}^M C(y_m) \right) \right).$$

The closure rule for ω -filters is then the conjunction $\tau = \bigwedge_{i \in \omega} \tau_i$. We can now define our separation rule σ as follows:

$$\sigma = \forall pq \left(p \not\leq q \rightarrow \exists C \left(C(p) \wedge \neg C(q) \wedge \bigwedge_{i \in \omega} \tau_i \right) \right).$$

Then σ is easily seen to be an r.e. separation rule, and a poset P satisfies σ if and only if whenever $p \not\leq q \in P$ there is an ω -filter of P containing p but not q , as required. Note that σ is not finite as we need τ_i for all $i \in \omega$.

DEFINITION 2.5. Let \mathcal{L} be a first-order signature, let \mathcal{A} be a class of \mathcal{L} -structures, and let \mathcal{B} be a subclass of \mathcal{A} . Then \mathcal{B} is a *separation subclass* of \mathcal{A} if there is a separation scheme Σ such that $\mathcal{B} = \{A \in \mathcal{A} : A \models \Sigma\}$. Here \models is defined using the standard semantics for second-order logic. A separation subclass is said to be *essentially r.e./countable/finite* when it can be defined using a separation scheme with the corresponding property. If \mathcal{A} is the class of all \mathcal{L} -structures then we say \mathcal{B} is a *separation class*.

EXAMPLE 2.6. We say a poset P is *representable* if there is a set X and an order embedding $h : P \rightarrow \wp(X)$ such that h preserves finite meets and joins from P whenever they exist (here $\wp(X)$ is considered as a lattice with operations \cup and \cap). It is easy to prove that a poset P is representable if and only if whenever $p \not\leq q \in P$ there is an ω -filter of P containing p and not q (see, for example, [15, Theorem 2.4]). Thus, building on Example 2.4, we see that the class of representable posets is an essentially r.e. separation subclass of the class of posets, using the separation scheme

$\Sigma = \{\sigma\}$. Note that Σ is essentially r.e. but not essentially finite, as while it contains only a single separation rule, this separation rule is infinite.

Generalising, given any $2 \leq \alpha, \beta \leq \omega$ we say a poset P is (α, β) -representable if there is a set X and an order embedding $h : P \rightarrow \wp(X)$ such that h preserves meets of cardinality strictly less than α , and joins of cardinality strictly less than β . Adapting the previous argument we can show the class of (α, β) -representable posets is an essentially r.e. separation subclass of the class of all posets, and is essentially finite when $\alpha, \beta < \omega$.

As may be expected given the definitions, the machinery of separation subclasses is not weaker than the machinery of first-order logic when it comes to specifying subclasses of classes of \mathcal{L} -structures. We make this precise in the following proposition. More surprisingly, it turns out that is not stronger either. This is the result of Theorem 2.14.

PROPOSITION 2.7. *If \mathcal{A} is a class of \mathcal{L} -structures, and if $\mathcal{B} \subseteq \mathcal{A}$ is elementary relative to \mathcal{A} , then \mathcal{B} is a separation subclass of \mathcal{A} . Moreover, if the axiomatisation of \mathcal{B} relative to \mathcal{A} is finite/countable/r.e., then \mathcal{B} is an essentially finite/countable/r.e. separation subclass of \mathcal{A} .*

PROOF. Since separation rules of order zero are just \mathcal{L} -sentences, this is an immediate consequence of Definition 2.5. ⊢

It will be useful to slightly generalise the familiar notion of a pseudoelementary class.

DEFINITION 2.8. Let \mathcal{A} be a class of \mathcal{L} -structures, and let $\mathcal{B} \subseteq \mathcal{A}$. Then \mathcal{B} is *pseudoelementary relative to \mathcal{A}* if there is an extension \mathcal{L}' of \mathcal{L} , and an \mathcal{L}' -theory T such that

$$\mathcal{B} = \{A \in \mathcal{A} : \text{we can interpret the additional symbols of } \mathcal{L}' \text{ so that } A \models T\}.$$

If \mathcal{A} is the class of all \mathcal{L} -structures, then being pseudoelementary relative to \mathcal{A} is the same as being pseudoelementary as it is usually defined. We say a pseudoelementary class \mathcal{B} is *essentially finite/countable/r.e.* relative to \mathcal{A} if \mathcal{L}' and T are both finite/countable/r.e. Classes that are essentially finite as pseudoelementary classes relative to the class of all \mathcal{L} -structures are often referred to as being *basic pseudoelementary*.

LEMMA 2.9. *If \mathcal{B} is a separation subclass of \mathcal{A} , then \mathcal{B} is pseudoelementary relative to \mathcal{A} . Moreover, if \mathcal{B} is essentially finite/countable/r.e. as a separation subclass of \mathcal{A} , then \mathcal{B} is essentially finite/countable/r.e. pseudoelementary relative to \mathcal{A} .*

PROOF. Let \mathcal{A} be a class of \mathcal{L} -structures, let Σ be a separation scheme defining \mathcal{B} relative to \mathcal{A} , and let $\sigma \in \Sigma$ have order K for some $K > 0$ (as there is nothing to do in the case where $K = 0$). So

$$\sigma = \forall \vec{x}_N \left(\mu_\sigma(\vec{x}_N) \rightarrow \exists \vec{C}_K (\eta_\sigma(\vec{x}_N) \wedge \tau_\sigma) \right),$$

where we either have $\tau_\sigma = \top$ or

$$\tau_\sigma = \bigwedge_I \tau_\sigma^i.$$

Moreover, assuming $\tau_\sigma \neq \top$, for each $i \in I$, the formula τ_σ^i is given by

$$\tau_\sigma^i = \forall \vec{y}_{M_i} (\gamma_\sigma^i(\vec{y}_{M_i}) \rightarrow \psi_\sigma^i(\vec{y}_{M_i})).$$

Expand \mathcal{L} to a new signature \mathcal{L}'_σ by adding new $(n + 1)$ -ary predicate symbols R_1, \dots, R_K . Define

$$t_\sigma = \forall \vec{x}_N (\mu_\sigma(\vec{x}_N) \rightarrow \hat{\eta}_\sigma(\vec{x}_N)),$$

where $\hat{\eta}_\sigma$ is η_σ but with every occurrence of $C_k(-)$ replaced by $R_k(\vec{x}_N, -)$, for all $k \in \{1, \dots, K\}$.

Now, assuming $\tau_\sigma \neq \top$, let $i \in I$, and define

$$t_\sigma^i = \forall \vec{x}_N (\mu_\sigma(\vec{x}_N) \rightarrow \forall \vec{y}_{M_i} (\gamma_\sigma^i(\vec{y}_{M_i}) \rightarrow \hat{\psi}_\sigma^i(\vec{y}_{M_i}))),$$

where $\hat{\psi}_\sigma^i$ is defined by replacing occurrences of $C_k(-)$ in ψ_σ^i with $R_k(\vec{x}_N, -)$ for all $k \in \{1, \dots, K\}$.

For $\sigma \in \Sigma^{>0}$, define $T_\sigma = \{t_\sigma\} \cup \{t_\sigma^i : i \in I\}$. If the order of σ is 0 then σ is already an \mathcal{L} -sentence, so we define $T_\sigma = \{\sigma\}$ in this case. Define

$$T = \bigcup_{\sigma \in \Sigma} T_\sigma.$$

Then T is a theory for the expanded signature $\mathcal{L}' = \bigcup_{\sigma \in \Sigma} \mathcal{L}'_\sigma$. We assume that if $\sigma_1 \neq \sigma_2$ then the extra symbols added to \mathcal{L}'_{σ_1} and \mathcal{L}'_{σ_2} are all distinct. Define

$$\mathcal{B}' = \{A \in \mathcal{A} : \text{we can interpret the additional symbols of } \mathcal{L}' \text{ so that } A \models T\}.$$

Let $B \in \mathcal{B}'$, and let $\sigma \in \Sigma^{>0}$ have order K . Then we can interpret the additional symbols of \mathcal{L}' in B so that $B \models T_\sigma$. In particular, if $\vec{b}_N \in B$ is such that $B \models \mu_\sigma(\vec{b}_N)$, then $B \models \hat{\eta}_\sigma(\vec{b}_N)$, and, assuming that $\tau_\sigma \neq \top$ and given $i \in I$, we also have $B \models \forall \vec{y}_{M_i} (\gamma_\sigma^i(\vec{y}_{M_i}) \rightarrow \hat{\psi}_\sigma^i(\vec{y}_{M_i}))$. So, if $\mathcal{L}'_\sigma = \mathcal{L} \cup \{R_1, \dots, R_K\}$, then for each $k \in \{1, \dots, K\}$, whenever $B \models \mu_\sigma(\vec{b}_N)$ we can interpret C_k by

$$C_k(a) \iff R_k(\vec{b}_N, a),$$

and so a routine argument reveals that $B \models \sigma$. There is nothing to do for the case where σ has order 0, and so it follows that $\mathcal{B}' \subseteq \mathcal{B}$.

Conversely, if $B \in \mathcal{B}$ then we can make B into an \mathcal{L}' -structure by interpreting the new relations as follows. If R is one such new relation, then it is associated with a unary predicate symbol C appearing in some separation rule

$$\forall \vec{x}_N (\mu(\vec{x}_N) \rightarrow \exists \vec{C}_K (\eta(\vec{x}_N) \wedge \tau)).$$

Let $\vec{b}_N \in B$ and suppose $B \models \mu(\vec{b}_N)$. Then there is an associated instantiation of C in B which we denote $C_{\vec{b}_N}$. Now, define the interpretation of R in B using

$$R = \{(\vec{b}_N, a) : B \models \mu(\vec{b}_N) \wedge C_{\vec{b}_N}(a)\}.$$

Then another routine argument reveals that $B \models T$, and thus $\mathcal{B} \subseteq \mathcal{B}'$. So $\mathcal{B} = \mathcal{B}'$, and \mathcal{B} is pseudoelementary relative to \mathcal{A} as required.

If \mathcal{B} is essentially finite/countable/r.e. as a separation subclass of \mathcal{A} , then that \mathcal{B} is essentially finite/countable pseudoelementary/r.e. relative to \mathcal{A} follows immediately from the construction of \mathcal{L}' and T . \dashv

COROLLARY 2.10. *If \mathcal{B} is a separation class then \mathcal{B} is pseudoelementary.*

PROOF. This follows immediately from Lemma 2.9 and the definition of separation classes (Definition 2.5). \dashv

Converses to Lemma 2.9 and Corollary 2.10 do not hold in general. To see this, note that we shall show that separation classes are elementary (Theorem 2.14), while pseudoelementary classes may not be.

The following lemma is a mild generalisation of the well known fact that pseudoelementary classes are closed under ultraproducts.

LEMMA 2.11. *If \mathcal{A} is closed under ultraproducts and \mathcal{B} is pseudoelementary relative to \mathcal{A} , then \mathcal{B} is closed under ultraproducts.*

PROOF. Suppose T is an \mathcal{L}' -theory making \mathcal{B} pseudoelementary relative to \mathcal{A} . Let I be an indexing set and for each $i \in I$ let $B_i \in \mathcal{B}$. Let $\prod_U B_i$ be an ultraproduct. For every i we can define an \mathcal{L}' -structure on B_i , which we denote B'_i , such that $B'_i \models T$. Then $\prod_U B'_i \models T$, by Łoś' theorem [32], and as $\prod_U B_i \in \mathcal{A}$ it follows that $\prod_U B_i \in \mathcal{B}$. \dashv

The aim now is to show that separation subclasses of elementary classes are elementary. In view of Lemmas 2.9 and 2.11 it will be sufficient to prove they are closed under taking elementary substructures and appeal to [20, Theorem 2.13]. This is done by the following pair of technical lemmas.

LEMMA 2.12. *Let \mathcal{L} be a first-order signature, let A be an \mathcal{L} -structure and let B be an elementary substructure of A . Let a_1, \dots, a_N be elements of B , and let S_1, \dots, S_K be unary predicate symbols not appearing in \mathcal{L} . Define $\mathcal{L}^+ = \mathcal{L} \cup \{S_1, \dots, S_K\}$. Let $\eta(\vec{z}_N)$ be a quantifier-free first-order \mathcal{L}^+ -formula with free variables from \vec{z}_N . For each $k \in \{1, \dots, K\}$ let $X_k \subseteq A$, and use these sets to make A into an \mathcal{L}^+ -structure by interpreting S_k as X_k for each k . Similarly, make B into an \mathcal{L}^+ -structure by interpreting S_k as $X_k \cap B$ for each k . Then $A \models \eta(\vec{a}_N) \iff B \models \eta(\vec{a}_N)$.*

PROOF. We proceed by structural induction on η . Let $\vec{a}_N \in B$. If η is a pure \mathcal{L} -formula, i.e., if it involves none of the additional predicates, then that $A \models \eta(\vec{a}_N)$ if and only if $B \models \eta(\vec{a}_N)$ follows immediately from the assumption that B is an elementary substructure of A . So the non-trivial base cases are the atomic formulas of form $S_k(t(\vec{a}_N))$ where $t(\vec{x}_N)$ is an \mathcal{L} -term. As B is a substructure of A we have $t(\vec{a}_N) \in B$, and so

$$A \models S_k^A(t(\vec{a}_N)) \iff t(\vec{a}_N) \in X_k \iff t(\vec{a}_N) \in X_k \cap B \iff B \models S_k^B(t(\vec{a}_N)).$$

The inductive step is routine because η is quantifier-free, and so it suffices to deal with \neg and \wedge . \dashv

LEMMA 2.13. *If \mathcal{B} is a separation subclass of \mathcal{A} and \mathcal{A} is closed under taking elementary substructures, then \mathcal{B} is also closed under taking elementary substructures.*

PROOF. Let \mathcal{B} be a separation subclass of \mathcal{A} , let Σ be a separation scheme defining \mathcal{B} relative to \mathcal{A} , and suppose \mathcal{A} is closed under taking elementary substructures. Let $B \in \mathcal{B}$, and let B' be an elementary substructure of B . Then $B' \in \mathcal{A}$, as \mathcal{A} is closed under taking elementary substructures. We aim to show that $B' \models \Sigma$, and thus that $B' \in \mathcal{B}$. So, given arbitrary $\sigma \in \Sigma$ we must show $B' \models \sigma$. If the order of σ is 0 then this follows immediately from the assumption that B' is an elementary substructure of B (recall Definition 2.2). So suppose

$$\sigma = \forall \vec{x}_N \left(\mu(\vec{x}_N) \rightarrow \exists \vec{C}_K \left(\eta(\vec{x}_N) \wedge \bigwedge_I \forall \vec{y}_{M_i} (\gamma_i(\vec{y}_{M_i}) \rightarrow \psi_i(\vec{y}_{M_i})) \right) \right),$$

for some $I \neq \emptyset$, as the case where $\tau = \top$ is similar but easier. Let $a_1, \dots, a_N \in B'$, and suppose $B' \models \mu(\vec{a}_N)$. Then, as B' is an elementary substructure of B we must have $B \models \mu(\vec{a}_N)$, and thus

$$B \models \exists \vec{C}_K \left(\eta(\vec{a}_N) \wedge \bigwedge_I \forall \vec{y}_{M_i} (\gamma_i(\vec{y}_{M_i}) \rightarrow \psi_i(\vec{y}_{M_i})) \right).$$

This is equivalent to saying that we can extend \mathcal{L} with new unary predicate symbols C_1, \dots, C_K to a signature \mathcal{L}^+ , and make B into an \mathcal{L}^+ -structure in such a way that $B \models \eta(\vec{a}_N) \wedge \bigwedge_I \forall \vec{y}_{M_i} (\gamma_i(\vec{y}_{M_i}) \rightarrow \psi_i(\vec{y}_{M_i}))$ when this is treated as an \mathcal{L}^+ -formula in the obvious way. We treat B' as an \mathcal{L}^+ structure by interpreting the new predicates as the restrictions of their interpretations in B . We aim to use Lemma 2.12.

First of all, we have $B' \models \eta(\vec{a}_N)$ by immediate application of Lemma 2.12. Now, let $i \in I$, let $b_1, \dots, b_{M_i} \in B'$, and suppose $B' \models \gamma_i(\vec{b}_{M_i})$. Then, as B' is an elementary substructure of B , we also have $B \models \gamma_i(\vec{b}_{M_i})$, as γ_i is an \mathcal{L} -formula, and thus also $B \models \psi_i(\vec{b}_{M_i})$, as $B \models \sigma$. So, again by Lemma 2.12, we have $B' \models \psi_i(\vec{b}_{M_i})$. This is true for arbitrary choices of $b_1, \dots, b_{M_i} \in B'$, so it follows that $B' \models \forall \vec{y}_{M_i} (\gamma_i(\vec{y}_{M_i}) \rightarrow \psi_i(\vec{y}_{M_i}))$ as required. This is true for all $i \in I$, and so $B' \models \sigma$. This is true for all $\sigma \in \Sigma$, and so $B' \models \Sigma$. Since $B' \in \mathcal{A}$ and $B' \models \Sigma$ it follows that $B' \in \mathcal{B}$ as required. \dashv

THEOREM 2.14. *If \mathcal{A} is elementary and \mathcal{B} is a separation subclass of \mathcal{A} then \mathcal{B} is also elementary.*

PROOF. \mathcal{B} is closed under taking ultraproducts, by Lemmas 2.9 and 2.11, and is also closed under elementary substructures, by Lemma 2.13, so the result follows from [20, Theorem 2.13]. \dashv

Theorem 2.14 is not constructive, but we will later exploit the fact that we know that separation subclasses of elementary classes can be axiomatised to produce explicit axiomatisations.

§3. The separation game. We will define games played between two players, Abelard (\forall) and Eloise (\exists). The point here is that the existence of winning strategies in these games relates to a structure’s membership in separation subclasses, in a sense to be made precise in Proposition 3.1 below, and can also be captured in first-order logic, as we explain in Section 4. So the games provide a means to translate the second-order separation rules defining separation subclasses into first-order logic.

Our games are played over a fixed \mathcal{L} -structure in rounds numbered by the natural numbers starting with zero. In each round, \forall plays first, then \exists must respond. If a

player has no legal move to make when required to play, then that player loses the game immediately, and the game does not continue. If one player loses, then the other player necessarily wins. We say that \forall has an r -strategy if he can play in a way that guarantees he wins no later than round r . We say \exists has an r -strategy if she can play in a way that guarantees that \forall will not win till at least the $(r + 1)$ th round, either by not losing, or by winning herself prior to that point. We say that \exists has an ω -strategy if she can play in such a way that she can either win or survive indefinitely, however \forall plays.

We now define the rules of our games more precisely. Let \mathcal{L} be a first-order signature, and let \mathcal{A} be a class of \mathcal{L} -structures. Let Σ be a separation scheme for \mathcal{L} , and let $\sigma = \forall \vec{x}_N (\mu(\vec{x}_N) \rightarrow \exists \vec{C}_K (\eta(\vec{x}_N) \wedge \tau)) \in \Sigma^{>0}$ (recall that $\Sigma^{>0}$ is the subset of Σ containing the separation rules of positive order). Let $A \in \mathcal{A}$, and for each $k \in \{1, \dots, K\}$ let $S_k, \bar{S}_k \subseteq A$. We define the (A, σ) -game with starting position $(S_1, \dots, S_K, \bar{S}_1, \dots, \bar{S}_K)$. The idea is that, for each k , S_k will contain elements definitely specified by the monadic predicate C_k , and \bar{S}_k will denote a set of elements that are definitely in its complement. Over the course of the game \exists is forced to decide whether elements of A are, or are not, contained in S_k . Note that at any given point in the game, \bar{S}_k will generally be a strict subset of the complement of S_k , as there will usually be elements that \exists has not yet been forced to make a decision about. If \exists cannot make a move that does not violate the conditions to be defined below, then she loses the game. Formally, the game is played as follows:

- In round 0, \forall chooses $a_1, \dots, a_N \in A$ such that $A \models \mu(\vec{a}_N)$. In response, \exists must decide, for each $n \in \{1, \dots, N\}$ and $k \in \{1, \dots, K\}$, whether $a_n \in S_k$. If yes then a_n is added to S_k . If no then a_n is added to \bar{S}_k . \exists must choose in such a way that:

(1) $A \models \eta(\vec{a}_N)$, where η is treated as a formula for signature

$$\mathcal{L}^+ = \mathcal{L} \cup \{C_1, \dots, C_K\},$$

and C_k is interpreted as S_k for all k (where S_k includes any elements newly added by \exists).

(2) $S_k \cap \bar{S}_k = \emptyset$ for all k .

- In round r for $r > 0$, \forall must play a move of form (τ_i, \vec{b}_{M_i}) , where

$$\tau_i = \forall \vec{y}_{M_i} (\gamma(\vec{y}_{M_i}) \rightarrow \psi(\vec{y}_{M_i}))$$

is part of the conjunction $\tau = \bigwedge_I \tau_i$, and $b_1, \dots, b_{M_i} \in A$ such that $A \models \gamma(\vec{b}_{M_i})$. If $\tau = \top$ then \forall cannot do this, and so he loses.

\exists must respond by deciding, for each $m \in \{1, \dots, M_i\}$ and $k \in \{1, \dots, K\}$, whether $b_m \in S_k$. If yes then b_m is added to S_k , and if no then b_m is added to \bar{S}_k . \exists must choose in such a way that:

(1) $A \models \psi(\vec{b}_{M_i})$, where ψ is treated as an \mathcal{L}^+ -formula, and C_k is interpreted as S_k for all k (where S_k includes new elements added by \exists).

(2) $S_k \cap \bar{S}_k = \emptyset$ for all k .

We sometimes refer to the (A, σ) -game with starting position $S_k = \bar{S}_k = \emptyset$ for all $k \in \{1, \dots, K\}$ as the *simple* (A, σ) -game. Note that we only define these games for

σ of positive order, as in the order zero case σ is just a first-order sentence, so there is no need to ‘translate’ it into first-order logic.

PROPOSITION 3.1. *Let \mathcal{A} be a class of \mathcal{L} -structures, let \mathcal{B} be a separation subclass of \mathcal{A} defined by the separation scheme Σ , and let $A \in \mathcal{A}$. Then:*

- (1) *If $A \in \mathcal{B}$, then \exists has an ω -strategy for the simple (A, σ) -game for all $\sigma \in \Sigma^{>0}$.*
- (2) *If \mathcal{B} is essentially countable and A is countable, then the converse is true. That is, if $A \models \sigma$ for all $\sigma \in \Sigma \setminus \Sigma^{>0}$, and if \exists has an ω -strategy for the simple (A, σ) -game for all $\sigma \in \Sigma^{>0}$, then $A \in \mathcal{B}$.*

PROOF. For 1), if $A \in \mathcal{B}$ then $A \models \sigma$, so, given $\vec{a}_N \in A$ with $A \models \mu(\vec{a}_N)$, there are monadic predicates C_1, \dots, C_K such that $A \models \eta(\vec{a}_N) \wedge \tau$. In this case \exists can guarantee to never lose by assigning an element $b \in A$ to S_k if $A \models C_k(b)$, and to \bar{S}_k otherwise, whenever she is forced to make a choice.

For 2), suppose that A is countable, and that \exists has an ω -strategy for the simple (A, σ) -game for every $\sigma \in \Sigma^{>0}$. Let

$$\sigma = \forall \vec{x}_N \left(\mu(\vec{x}_N) \rightarrow \exists \vec{C}_K (\eta(\vec{x}_N) \wedge \tau) \right) \in \Sigma^{>0},$$

let $a_1, \dots, a_N \in A$ with $A \models \mu(\vec{a}_N)$, and suppose $A \not\models \mu(\vec{a}_N) \rightarrow \exists \vec{C}_K (\eta(\vec{a}_N) \wedge \tau)$. Then it follows that $A \not\models \exists \vec{C}_K (\eta(\vec{a}_N) \wedge \tau)$.

If $\tau = \top$, then \exists has no response to the opening move \vec{a}_N by \forall , contradicting the assumption that she has an ω -strategy in the game. So we assume that $\tau = \bigwedge_I \tau_i$ for some non-empty I . Since Σ is essentially countable and A is countable, we can order the moves (τ_i, \vec{b}_M) that \forall could potentially make using the natural numbers. Suppose \forall plays according to the strategy whereby in round 0 he plays \vec{a}_N , and in every subsequent round he plays the lowest ranked legal move that he has not yet played.

Consider the sets S_k for $k \in \{1, \dots, K\}$ constructed by \exists as she follows her ω -strategy against \forall . In other words, each S_k is the union of the corresponding sets from the individual rounds of the game. By the rules governing the opening round of play, and the assumption that \exists is playing according to an ω -strategy, we must have $A \models \eta(\vec{a}_N)$, if C_k is interpreted as S_k for all k . Thus, if $A \not\models \exists \vec{C}_K (\eta(\vec{a}_N) \wedge \tau)$ there must be some $i \in I$ with $\tau_i = \forall \vec{y}_{M_i} \left(\gamma(\vec{y}_{M_i}) \rightarrow \psi(\vec{y}_{M_i}) \right)$ such that $A \not\models \tau_i$, where C_k is interpreted as S_k for all k .

It follows that there must be $b_1, \dots, b_{M_i} \in A$ with $A \models \gamma(\vec{b}_{M_i})$ and $A \not\models \psi(\vec{b}_{M_i})$, again interpreting C_k as S_k for all k . But this corresponds to a legal move by \forall , so he must have played it at some point, as his strategy implies that he eventually plays every move that becomes available after the first round. Thus we must have $A \models \psi(\vec{b}_{M_i})$ when C_k is interpreted as S_k after all, as \exists is following an ω -strategy. This would be a contradiction. Thus we must have $A \models \exists \vec{C}_K (\eta(\vec{a}_N) \wedge \tau)$. Since this is true for every choice of \vec{a}_N such that $A \models \mu(\vec{a}_N)$, we have $A \models \sigma$, and since this argument holds for all $\sigma \in \Sigma^{>0}$, and we have assumed that $A \models \sigma$ for all $\sigma \in \Sigma \setminus \Sigma^{>0}$, it follows that $A \in \mathcal{B}$ as required. \dashv

Note that round 0 is conceptually distinct from the subsequent rounds. We define the *reduced (A, σ) -game with starting position $(S_1, \dots, S_K, \bar{S}_1, \dots, \bar{S}_K)$* to be the (A, σ) -game with the same starting position, but omitting round 0. For convenience

we keep the same labeling for rounds as in the normal game, so the reduced game starts with round 1, not round 0. The concept of an r -strategy for $r \geq 1$ carries over without modification for both players.

§4. Generating recursive axiomatisations. Given a separation scheme Σ , the next step is to find a set of first-order axioms equivalent to \exists having an ω -strategy in the simple (A, σ) -game for all $\sigma \in \Sigma^{>0}$. We must assume that Σ is at least essentially recursively enumerable for the main result (Theorem 4.5) to hold. We also assume for convenience that for every $\sigma \in \Sigma^{>0}$ the associated τ is of form $\bigwedge_{i \in \omega} \tau_i$. Recall that in practice, τ may be \top , or a conjunction of only finitely many formulas. We could avoid this assumption by dividing several of the definitions and proofs to come into ‘finite’ and ‘infinite’ cases, but we trust instead that the necessary alterations for the finite cases will become clear once the infinite case is understood.

Writing down these axioms will involve some quite intricate notational constructions, and we will benefit greatly later from taking the time now to prove some technical results. Recall that for distinct variable symbols x_1, \dots, x_N we use the notation $\vec{x}_N = \{x_1, \dots, x_N\}$. If v is a valuation in the model theoretic sense, we will often write, e.g., $v[\vec{x}_N]$ to stand for $\{v(x_1), \dots, v(x_N)\}$. Similarly, if Z is a set of variables we will use $v[Z]$ to denote $\{v(z) : z \in Z\}$.

First we want to formalise in first-order logic the statement that \exists can survive a certain finite number of rounds in an (A, σ) -game from a given starting position. We will do this by recursion. More explicitly, the statement that \exists has an $(r + 1)$ -strategy will be formed by writing a formula to the effect that, whatever move \forall makes, \exists can distribute the elements picked out by \forall in such a way that she has an r -strategy in the (A, σ) -game with the starting position that results from her choice.

To make this work, the states of the predicates C_k for $k \in \{1, \dots, K\}$ associated with the separation rule σ at a given point in the game will be captured using free variables introduced for this purpose and an assignment v of variables to elements of A . We will do this with free variables grouped into sets Z_1, \dots, Z_K . The idea is that for $k \in \{1, \dots, K\}$, the set $v[Z_k]$ captures the elements \exists has assigned to C_k . Lemma 4.1 below explains how we can translate formulas involving the predicates C_1, \dots, C_K into formulas where they are captured by free variables in this way.

LEMMA 4.1. *Let $\mathcal{L}^+ = \mathcal{L} \cup \{C_1, \dots, C_K\}$, where each C_k is a unary predicate symbol not appearing in \mathcal{L} , let Y be a finite set of variable symbols, and let $\psi(Y)$ be a quantifier-free \mathcal{L}^+ -formula whose free variables are taken from Y . For each $k \in \{1, \dots, K\}$ let Z_k be a finite set of variable symbols. Let $\vec{Z}_K = (Z_1, \dots, Z_K)$. Then we can define a quantifier-free \mathcal{L} -formula $\psi_{\vec{Z}_K}$ whose free variables are from the set $Y \cup \bigcup_{k=1}^K Z_k$, such that, whenever A is an \mathcal{L} -structure and v is an assignment, we have*

$$A, v \models_{\mathcal{L}} \psi_{\vec{Z}_K}(Y \cup \bigcup_{k=1}^K Z_k) \iff A, v \models_{\mathcal{L}^+} \psi(Y),$$

where A is treated as an \mathcal{L}^+ -structure by interpreting C_k so that

$$A, v \models C_k(x) \iff v(x) \in v[Z_k],$$

for all $k \in \{1, \dots, K\}$.

PROOF. We use induction on the construction of ψ . If ψ is an atomic formula, then either:

- (1) $\psi(Y) = R(t_1(Y), \dots, t_N(Y))$, where R is some N -ary relation symbol from \mathcal{L} and t_n is an \mathcal{L} -term for each $n \in \{1, \dots, N\}$,
- (2) $\psi(Y) = t_1(Y) \approx t_2(Y)$ where t_1 and t_2 are \mathcal{L} -terms, or
- (3) $\psi(Y) = C_k(t(Y))$, where t is an \mathcal{L} -term and $k \in \{1, \dots, K\}$.

In cases (1) and (2), the interpretation of the additional predicates of \mathcal{L}^+ isn't relevant, so we can define

$$\psi_{\bar{Z}_K}(Y) = \psi(Y).$$

In the third case, define

$$\psi_{\bar{Z}_K}(Y \cup Z_k) = \left(\bigvee_{z \in Z_k} t(Y) \approx z \right).$$

Then

$$\begin{aligned} & A, v \models_{\mathcal{L}} \psi_{\bar{Z}_K}(Y \cup Z_k) \\ \iff & v(t(Y)) = v(z) \text{ for some } z \in Z_k \\ \iff & A, v \models_{\mathcal{L}^+} \psi(Y), \text{ where } C_k \text{ is interpreted as } v[Z_k] \text{ as described.} \end{aligned}$$

For the inductive step, consider first $\neg\psi$ such that $\psi_{\bar{Z}_K}$ is known to exist for ψ . In this case we can just use $\neg\psi_{\bar{Z}_K}$, as

$$\begin{aligned} A, v \models_{\mathcal{L}} \neg\psi_{\bar{Z}_K} & \iff A, v \not\models_{\mathcal{L}} \psi_{\bar{Z}_K} \\ & \iff A, v \not\models_{\mathcal{L}^+} \psi \\ & \iff A, v \models_{\mathcal{L}^+} \neg\psi. \end{aligned}$$

Consider next $\psi^1 \vee \psi^2$, such that appropriate $\psi^1_{\bar{Z}_K}$ and $\psi^2_{\bar{Z}_K}$ exist. We use $\psi^1_{\bar{Z}_K} \vee \psi^2_{\bar{Z}_K}$, because

$$\begin{aligned} A, v \models_{\mathcal{L}} \psi^1_{\bar{Z}_K} \vee \psi^2_{\bar{Z}_K} & \iff A, v \models_{\mathcal{L}} \psi^1_{\bar{Z}_K} \text{ or } A, v \models_{\mathcal{L}} \psi^2_{\bar{Z}_K} \\ & \iff A, v \models_{\mathcal{L}^+} \psi^1 \text{ or } A, v \models_{\mathcal{L}^+} \psi^2 \\ & \iff A, v \models_{\mathcal{L}^+} \psi^1 \vee \psi^2. \end{aligned}$$

Since ψ is quantifier-free, we are done. ⊣

As explained above, the state of each C_k during the (A, σ) -game will be captured by $v[Z_k]$ for some set Z_k of variable symbols and assignment v . During the game, \exists must assign elements both to C_k and its complement. For each k , we will use a set \bar{Z}_k of variable symbols for this purpose. Explicitly, $v[\bar{Z}_k]$ captures the elements that have been assigned to the complement of C_k . In each round of the game, \exists is presented with new elements of A that she must assign to C_k sets or their complements. These new elements will be captured using v and a new set of variable symbols, e.g., as $v[Y]$.

As discussed previously, a key idea for us is to define for each $r \geq 1$ formulas to the effect of 'given a starting position $(v[Z_1], \dots, v[Z_K], v[\bar{Z}_1], \dots, v[\bar{Z}_K])$ in an

(A, σ) -game, whatever move \forall makes involving elements $v[Y]$, \exists can for each k assign each element of $v[Y]$ to either $v[Z_k]$ or $v[\bar{Z}_k]$ to obtain $v[Z'_k]$ and $v[\bar{Z}'_k]$ in such a way that she has an $(r - 1)$ -strategy in the (A, σ) -game with starting position $(v[Z'_1], \dots, v[Z'_K], v[\bar{Z}'_1], \dots, v[\bar{Z}'_K])$. As mentioned above, this will be done using recursion, so the formula stating that \exists has an $(r - 1)$ -strategy in the (A, σ) -game with starting position $(v[Z'_1], \dots, v[Z'_K], v[\bar{Z}'_1], \dots, v[\bar{Z}'_K])$ must involve the variables of Y , but now cast in specified roles as members of the sets $Z'_1, \dots, Z'_K, \bar{Z}'_1, \dots, \bar{Z}'_K$ (as these sets of variables are used to track the members of C_1, \dots, C_K).

The next definition sets up a notation for this process of adding new variables to sets. The situation to be described is that we have sets of variables $Z_1, \dots, Z_K, \bar{Z}_1, \dots, \bar{Z}_K$, and another set of variables Y . For each $y \in Y$ and for each $1 \leq k \leq K$ we want to add y to either Z_k or to \bar{Z}_k . What we do for a given $y \in Y$ and $k \in \{1, \dots, K\}$ is controlled by a function $f : Y \times \{1, \dots, K\} \rightarrow \{0, 1\}$. Explicitly, if $f(y, k) = 1$ then we add y to Z_k , otherwise we add y to \bar{Z}_k . The functions Δ_K and $\Delta_{\bar{K}}$ in Definition 4.2 below formalise this. The input is the data of the sets Z_1, \dots, Z_K, Y (for Δ_K), or $Z_1, \dots, Z_K, \bar{Z}_1, \dots, \bar{Z}_K, Y$ (for $\Delta_{\bar{K}}$), and the ‘control’ function f , and the output is either Z_1, \dots, Z_K (for Δ_K), or $Z_1, \dots, Z_K, \bar{Z}_1, \dots, \bar{Z}_K, Y$ (for $\Delta_{\bar{K}}$) after the elements of Y have been added as just described. We define two functions because sometimes we care about the C_k sets and their complements, and other times just the sets themselves.

DEFINITION 4.2 ($\Delta_K, \Delta_{\bar{K}}, F_Y^K$). Given a set Y and $1 \leq K < \omega$, let $f : Y \times \{1, \dots, K\} \rightarrow \{0, 1\}$. For each $k \in \{1, \dots, K\}$ let Z_k and \bar{Z}_k be sets. We use the shorthand $\vec{Z}_K = (Z_1, \dots, Z_K)$, and $\vec{\bar{Z}}_K = (\bar{Z}_1, \dots, \bar{Z}_K)$. Define

$$\Delta_K(\vec{Z}_K, f) = (Z'_1, \dots, Z'_K),$$

where for $k \in \{1, \dots, K\}$ we have

$$Z'_k = Z_k \cup \{y \in Y : f(y, k) = 1\}.$$

Similarly, define

$$\Delta_{\bar{K}}(\vec{Z}_K, \vec{\bar{Z}}_K, f) = (Z'_1, \dots, Z'_K, \bar{Z}'_1, \dots, \bar{Z}'_K),$$

where Z'_k is as above for all $k \in \{1, \dots, K\}$, and

$$\bar{Z}'_k = \bar{Z}_k \cup \{y \in Y : f(y, k) = 0\}.$$

We will use F_Y^K to denote the set of functions from $Y \times \{1, \dots, K\}$ to $\{0, 1\}$.

We now define \mathcal{L} -formulas as follows, noting the assumptions made about Σ stated at the start of this section. We assume we are working with a countably infinite pool of variable symbols.

- For each $1 \leq K < \omega$, and for each $\vec{Z}_K = (Z_1, \dots, Z_K)$ and $\vec{\bar{Z}}_K = (\bar{Z}_1, \dots, \bar{Z}_K)$ such that Z_k and \bar{Z}_k are finite sets of variables for all $k \in \{1, \dots, K\}$, define

$$D_{(\vec{Z}_K, \vec{\bar{Z}}_K)}$$

to be a quantifier-free \mathcal{L} -formula with free variables exactly $\bigcup_{k=1}^K (Z_k \cup \bar{Z}_k)$ such that

$$A, v \models D_{(\bar{Z}_K, \bar{Z}_K)} \iff v[Z_k] \cap v[\bar{Z}_k] = \emptyset \text{ for all } k \in \{1, \dots, K\}.$$

- For each $1 \leq K < \omega$, for each $i \in \omega$, for each $\sigma \in \Sigma^{>0}$, and for each (\bar{Z}_K, \bar{Z}_K) , define

$$\alpha_{(\bar{Z}_K, \bar{Z}_K)0i}^\sigma = D_{(\bar{Z}_K, \bar{Z}_K)}.$$

- For each

$$\sigma = \forall \bar{x}_N \left(\mu(\bar{x}_N) \rightarrow \exists \bar{C}_K (\eta(\bar{x}_N) \wedge \bigwedge_{i \in \omega} \forall \bar{y}_{M_i} (\gamma^i(\bar{y}_{M_i}) \rightarrow \psi^i(\bar{y}_{M_i}))) \right) \in \Sigma^{>0},$$

for each $i \in \omega$, for each $1 \leq r < \omega$, and for each (\bar{Z}_K, \bar{Z}_K) , recursively define

$$\alpha_{(\bar{Z}_K, \bar{Z}_K)ri}^\sigma = \bigwedge_{j \leq i} \forall \bar{y}_{M_j} \left(\gamma^j(\bar{y}_{M_j}) \rightarrow \bigvee_{f \in F_{\bar{y}_{M_j}}^K} (\psi_{\Delta_K(\bar{Z}_K, f)}^j \wedge \alpha_{\Delta_K(\bar{Z}_K, f)(r-1)i}^\sigma) \right),$$

where $\psi_{\Delta_K(\bar{Z}_K, f)}^j$ is constructed from ψ^j as in Lemma 4.1. In other words,

$$A, v \models \psi_{\Delta_K(\bar{Z}_K, f)}^j(\bar{y}_{M_j} \cup \bigcup_{k=1}^K Z'_k) \iff A, v \models_{\mathcal{L}^+} \psi^j(\bar{y}_{M_j}),$$

where each C_k is interpreted as $v[Z'_k]$, and Z'_k is constructed from Z_k and \bar{y}_{M_j} according to f (recall Definition 4.2). What these formulas are intended to capture is the idea that \exists can respond to all moves involving τ_j for $j \leq i$ played by \forall , and moreover can do so in such a way that she will continue to be able to respond successfully for at least r rounds. This will be made precise in Lemma 4.3. Note that although it is not apparent from the notation, we are assuming that every new occurrence of \bar{y}_{M_j} in the construction of these formulas involves only fresh variable symbols. If we allow variable symbols to be repeated then it turns out we do not properly capture the concept of ‘adding elements to Z_k ’, which is what the Δ operations are supposed to be for. This is explained in the proof of the following lemma.

LEMMA 4.3. *Let $1 \leq K < \omega$, let $A \in \mathcal{A}$, let \mathcal{B} be a separation subclass of \mathcal{A} defined by the essentially r.e. separation scheme Σ , let $1 \leq K, < \omega$, and let $\sigma \in \Sigma$ be a separation rule of order K . Then for all finite $Z_1, \dots, Z_K, \bar{Z}_1, \dots, \bar{Z}_K$, for all assignments v , for all $i \in I$, and for all $1 \leq r < \omega$, the following are equivalent:*

- (1) $A, v \models \alpha_{(\bar{Z}_K, \bar{Z}_K)ri}^\sigma$.
- (2) \exists has an r -strategy in the reduced (A, σ) -game with starting position

$$(v[Z_1], \dots, v[Z_K], v[\bar{Z}_1], \dots, v[\bar{Z}_K]),$$

where \forall can only play moves involving τ_j when $j \leq i$.

PROOF. We use induction on r . For the base case ($r = 1$), let $i \in \omega$ and suppose first that $A, v \models \alpha^\sigma_{(\vec{Z}_K, \vec{\bar{Z}}_K)1i}$. Let v' be an assignment agreeing with v about everything except, possibly, \vec{y}_{M_j} , and suppose $A, v' \models \gamma^j(\vec{y}_{M_j})$ for some $j \leq i$. Then there is $f \in F^K_{\vec{y}_{M_j}}$ such that $A, v' \models \psi^j_{\vec{Z}'_K} \wedge D_{(\vec{Z}'_K, \vec{\bar{Z}}'_K)}$ where $(Z'_1, \dots, Z'_K) = \Delta_K(Z_1, \dots, Z_K, f)$ and $(Z'_1, \dots, Z'_K, \vec{Z}'_1, \dots, \vec{Z}'_K) = \Delta_{\vec{K}}(\vec{Z}_K, \vec{\bar{Z}}_K, f)$. That is, Z'_k is the new value of Z_k as controlled by f , and similar for \vec{Z}'_k , for each $k \in \{1, \dots, K\}$ (recall Definition 4.2). Because the variables of \vec{y}_{M_j} do not appear in any Z_k (by the assumption mentioned after the definition of $\alpha^\sigma_{(\vec{Z}_K, \vec{\bar{Z}}_K)ri}$), we have $v[Z_k] = v'[Z_k]$ for all $k \in \{1, \dots, K\}$, and similar for \vec{Z}_k .

So, suppose \forall chooses $\vec{a}_{M_j} \in A$ and plays the move (τ_j, \vec{a}_{M_j}) . We define v' to be v except that $v'(y_m) = a_m$ for all $m \in \{1, \dots, M_j\}$. As discussed above, there is f such that $A, v' \models \psi^j_{\Delta_K(\vec{Z}_K, f)} \wedge D_{\Delta_{\vec{K}}(\vec{Z}_K, \vec{\bar{Z}}_K, f)}$. So, appealing to Lemma 4.1 applied to $\psi^j_{\Delta_K(\vec{Z}_K, f)}$, there is a way, described by f , to assign each a_m to either $v[Z_k]$ or $v[\vec{Z}_k]$ for each $k \in \{1, \dots, K\}$ such that $A \models_{\mathcal{L}^+} \psi^j$ if we interpret each C_k as the modified $v[Z_k]$. Moreover, from the definition of D , we see that each modified $v[Z_k]$ is disjoint from the modified $v[\vec{Z}_k]$, and so \exists can survive the first round of the reduced (A, σ) -game with starting position $(v[Z_1], \dots, v[Z_K], v[\vec{Z}_1], \dots, v[\vec{Z}_K])$, so long as \forall starts with a move involving τ_j for some $j \leq i$. This proves that (1) \implies (2) for $r = 1$.

For the converse, suppose (2) holds and that $A, v' \models \gamma^j(\vec{y}_{M_j})$ for some $j \leq i$ with v' agreeing with v about everything except, possibly, \vec{y}_{M_j} . Then \exists 's strategy tells us how to find $f \in F^K_{\vec{y}_{M_j}}$ appropriately. That is, if \exists 's strategy assigns $v'(y_m)$ to C_k , then f should assign y_m to Z_k (formally, $f(y_m, k) = 1$), and if not it should assign y_m to \vec{Z}_k (formally, $f(y_m, k) = 0$). With this f , if Z'_k and \vec{Z}'_k are, respectively, the modified Z_k and \vec{Z}_k for all $k \in \{1, \dots, K\}$, we have $v'[Z'_k] \cap v'[\vec{Z}'_k] = \emptyset$, and $A, v' \models_{\mathcal{L}^+} \psi^j(\vec{y}_{M_j})$ if we interpret C_k as $v'[Z'_k]$ for all $k \in \{1, \dots, K\}$. From Lemma 4.1 it follows that $A, v' \models \psi^j_{\Delta_K(\vec{Z}_K, f)} \wedge D_{\Delta_{\vec{K}}(\vec{Z}_K, \vec{\bar{Z}}_K, f)}$. So (2) \implies (1).

For the inductive step, let $1 < R < \omega$ and suppose the claim is true for all $1 \leq r < R$. Then, by the inductive hypothesis, and appealing to similar reasoning as used for the base case, $A, v \models \alpha^\sigma_{(\vec{Z}_K, \vec{\bar{Z}}_K)Ri}$ if and only if, whatever move involving τ_j for $j \leq i$ \forall plays, \exists can respond in such a way that she has an $(R - 1)$ -strategy in the game whose starting position corresponds to her response. But this is the same as saying that \exists has an R -strategy as claimed. \dashv

The formulas $\alpha^\sigma_{(\vec{Z}_K, \vec{\bar{Z}}_K)ri}$ we have just defined let us deal with reduced (A, σ) -games, but we need more to handle the full games. With that in mind we now define formulas where round zero is included. These will include the α formulas as subformulas.

- Let $\sigma \in \Sigma^{>0}$, and for all $i \in \omega$, for all finite $(\vec{Z}_K, \vec{\bar{Z}}_K)$, and for all $0 \leq r < \omega$, define

$$\beta^\sigma_{(\vec{Z}_K, \vec{\bar{Z}}_K)ri} = \forall \vec{x}_N (\mu(\vec{x}_N) \rightarrow \bigvee_{f \in F^K_{\vec{x}_N}} (\eta_{\Delta_K(\vec{Z}_K, f)} \wedge \alpha^\sigma_{\Delta_{\vec{K}}(\vec{Z}_K, \vec{\bar{Z}}_K, f)ri})).$$

Here $\eta_{\Delta_K(\vec{Z}_K, f)}$ is constructed from η as in Lemma 4.1. We assume that the variables in \vec{x}_N do not appear in any subformula not involving μ . The formula $\beta_{(\vec{Z}_K, \vec{Z}_K)ri}^\sigma$ is intended to say that, whatever move \forall makes in round zero of the (A, σ) -game with starting position (\vec{Z}_K, \vec{Z}_K) , \exists can respond in such a way that she survives this opening round, and can survive in the resulting reduced game for at least r rounds, so long as \forall only plays moves involving τ_j for $j \leq i$. Lemma 4.4 makes this precise.

LEMMA 4.4. *Let $1 \leq K < \omega$, let $A \in \mathcal{A}$, let \mathcal{B} be a separation subclass of \mathcal{A} defined by the essentially r.e. separation scheme Σ , and let $\sigma \in \Sigma$ be a separation rule of order K . Then for all (\vec{Z}_K, \vec{Z}_K) , for all assignments v , for all $i \in I$, and for all $r \in \omega$, the following are equivalent:*

- (1) $A, v \models \beta_{(\vec{Z}_K, \vec{Z}_K)ri}^\sigma$.
- (2) \exists has an r -strategy in the (A, σ) -game with starting position

$$(v[Z_1], \dots, v[Z_K], v[\vec{Z}_1], \dots, v[\vec{Z}_K]),$$

where \forall can only play moves involving τ_j when $j \leq i$.

PROOF. (1) is the statement that whenever v' is an assignment agreeing with v about everything except, possibly, \vec{x}_N , if $A, v' \models \mu(\vec{x})$ then there is a way \exists can assign the variables of \vec{x}_N to \vec{Z}_K, \vec{Z}_K so that $A, v' \models \eta_{\vec{Z}_K}$ and $A, v' \models \alpha_{(\vec{Z}'_K, \vec{Z}'_K)ri}^\sigma$ for the resulting values $Z'_1, \dots, Z'_K, \vec{Z}'_1, \dots, \vec{Z}'_K$.

So, suppose \forall plays \vec{a}_N as an opening move. Define v' to be like v except that $v'(x_n) = a_n$ for all $n \in \{1, \dots, N\}$. By the assumption that \vec{x}_N involves only variables not appearing in subformulas not involving μ , we have $v[Z_k] = v'[Z_k]$ for all k . Then $A, v' \models \mu(\vec{x}_N)$, and by the preceding paragraph, there is a way \exists can assign variables from \vec{x}_N so that $A, v' \models \eta_{\vec{Z}'_K}$ and $A, v' \models \alpha_{(\vec{Z}'_K, \vec{Z}'_K)ri}^\sigma$. So, for her response, \exists adds a_n to $v[Z_k]$ if x_n is added to Z_k , and otherwise adds it to $v[\vec{Z}_k]$. Since $A, v' \models \eta_{\vec{Z}'_K}$, it follows from Lemma 4.1 that \exists can survive the opening round, and since $A, v' \models \alpha_{(\vec{Z}'_K, \vec{Z}'_K)ri}^\sigma$, it follows from Lemma 4.3 that she has an r -strategy in the resulting reduced game. Thus (1) \implies (2).

Conversely, if \exists has an r -strategy in the game, if v' is an assignment agreeing with v about everything except possibly \vec{x}_N and $A, v' \models \mu(\vec{x}_N)$, then this strategy implies the existence of a suitable f . Explicitly, $f(x_n, k) = 1$ if $v'(x_n)$ is added to $v[Z_k]$, and $f(x_n, k) = 0$ otherwise. Thus (2) \implies (1) as required. \dashv

The formulas $\beta_{(\vec{Z}_K, \vec{Z}_K)ri}^\sigma$ defined above say that \exists has an r strategy in a constrained (A, σ) -game with a specified starting position. What we want for our main result are formulas stating that this is true for the ‘simple’ starting position. That is, where the sets all start off empty. We also want to cover separation rules of order zero. With is in mind, we proceed as follows.

- For all $\sigma \in \Sigma$, for all $i \in \omega$, and for all $r \in \omega$, define

$$\hat{\beta}_{ri}^\sigma = \begin{cases} \beta_{(\vec{0}_K, \vec{0}_K)ri}^\sigma & \text{if } \sigma \in \Sigma^{>0} \text{ has order } K, \\ \sigma & \text{otherwise.} \end{cases}$$

Here $\vec{0}_K$ stands for \vec{Z}_K where $Z_k = \emptyset$ for all k . So, given any $\sigma \in \Sigma^{>0}$, and given any $r, i \in \omega$, the first-order sentence $\hat{\beta}_{ri}^\sigma$ states that \exists has an r -strategy in the simple (A, σ) -game in which \forall can only play moves involving τ_j for $j \leq i$. Moreover, if Σ is essentially r.e., the set $\{\hat{\beta}_{ri}^\sigma : \sigma \in \Sigma, i, r \in \omega\}$ can be made recursively enumerable.

This brings us to the following theorem, which is the main result of this section, and is a considerable generalisation of [24, Theorem 4.5] and [15, Theorem 5.6]. Nevertheless, the key ingredients of the proofs are essentially the same.

THEOREM 4.5. *Let \mathcal{A} be an elementary class of \mathcal{L} -structures, let $A \in \mathcal{A}$, and let \mathcal{B} be a separation subclass of \mathcal{A} defined by the essentially r.e. separation scheme Σ . Then*

$$A \in \mathcal{B} \iff A \models \hat{\beta}_{ri}^\sigma \text{ for all } \sigma \in \Sigma \text{ and for all } r, i \in \omega.$$

PROOF. If $A \in \mathcal{B}$ then, for all $\sigma = \forall \vec{x}_N (\mu(\vec{x}_N) \rightarrow \exists \vec{C}_K (\eta(\vec{x}) \wedge \tau)) \in \Sigma^{>0}$, we can use the predicates C_k for $k \in \{1, \dots, K\}$ to guide the strategy of \exists in the appropriate games. By Proposition 3.1, \exists has an ω -strategy in every simple (A, σ) -game for $\sigma \in \Sigma^{>0}$, and thus $A \models \beta_{(\vec{0}_K, \vec{0}_K)ri}^\sigma$ for all $\sigma \in \Sigma^{>0}$ and for all $r, i \in \omega$, by Lemma 4.4.

If $\sigma \in \Sigma \setminus \Sigma^{>0}$ then $\sigma = \hat{\beta}_{ri}^\sigma$, and so it follows immediately that $A \models \hat{\beta}_{ri}^\sigma$ for all σ, r, i .

Conversely, suppose first that A is countable and that $A \notin \mathcal{B}$. Then, either $A \not\models \sigma$ for some $\sigma \in \Sigma \setminus \Sigma^{>0}$, or, by Proposition 3.1, there is $\sigma \in \Sigma^{>0}$ such that \exists does not have an ω -strategy in the simple (A, σ) -game. In the former case we immediately have $A \not\models \hat{\beta}_{ri}^\sigma$, just by definition of $\hat{\beta}_{ri}^\sigma$, so we consider the latter. It follows from König’s Tree Lemma [30] that some game tree for the simple (A, σ) -game is finite (otherwise \exists would have a strategy defined using an infinite branch). There are only a finite number of \forall moves in this game tree, and so, if $i \in \omega$ is the largest index of a τ_i used in a move by \forall in this tree, we have $A \not\models \hat{\beta}_{ri}^\sigma$ for some $r \in \omega$, by Lemma 4.4.

Now, suppose A is uncountable, and suppose also that $A \models \hat{\beta}_{ri}^\sigma$ for all $\sigma \in \Sigma$ and for all $r, i \in \omega$. Then, by the downward Löwenheim–Skolem Theorem, A has a countable elementary substructure, A' , and, as $A' \models \hat{\beta}_{ri}^\sigma$ for all σ, i, r , it follows from our proof of the countable case that $A' \in \mathcal{B}$. Moreover, by Theorem 2.14, \mathcal{B} is elementary, so A' is a model of the elementary theory defining \mathcal{B} . But A and A' are elementarily equivalent, so A is also a model of this theory, and thus $A \in \mathcal{B}$ as claimed. \dashv

COROLLARY 4.6. *Let \mathcal{A} be an elementary class of \mathcal{L} -structures, let $A \in \mathcal{A}$, and let \mathcal{B} be an essentially r.e. separation subclass of \mathcal{A} defined by the separation scheme Σ . Then \mathcal{B} has a recursive axiomatisation relative to \mathcal{A} .*

Moreover, if for every $\sigma = \forall \vec{x}_N (\mu(\vec{x}_N) \rightarrow \exists \vec{C}_K (\eta(\vec{x}_N) \wedge \bigwedge_I \tau_i)) \in \Sigma^{>0}$, the prenex normal form of μ contains no universal quantifiers, and, in addition, for every conjunct $\tau_i = \forall \vec{y}_{M_i} (\gamma_i(\vec{y}_{M_i}) \rightarrow \psi_i(\vec{y}_{M_i}))$ of τ , the prenex normal form of γ_i contains no universal quantifiers, then, so long as the prenex normal form of every $\sigma \in \Sigma \setminus \Sigma^{>0}$ is universal, there is a recursive universal axiomatisation of \mathcal{B} relative to \mathcal{A} .

PROOF. Since Σ is essentially r.e. the set $\mathcal{T} = \{\hat{\beta}_{ri}^\sigma : \sigma \in \Sigma, r, i \in \omega\}$ can be assumed to be recursively enumerable. By Theorem 4.5, we know \mathcal{T} axiomatises \mathcal{B} relative to \mathcal{A} , and by Craig’s trick, any class with an r.e. axiomatisation relative to a superclass also has a recursive axiomatisation relative to that superclass (see, for example, [25, Exercise 6.3.1]).

Now, let $\sigma = \forall \vec{x}_N (\mu(\vec{x}_N) \rightarrow \exists \vec{C}_K (\eta(\vec{x}_N) \wedge \bigwedge_I \tau_i)) \in \Sigma^{>0}$, and let $\hat{\beta}_{ri}^\sigma \in \mathcal{T}$ be one of the generated \mathcal{L} -sentences axiomatising \mathcal{B} relative to \mathcal{A} . That is,

$$\hat{\beta}_{ri}^\sigma = \forall \vec{x}_N (\mu(\vec{x}_N) \rightarrow \bigvee_{f \in F_{\vec{x}_N}^K} (\eta_{\Delta_K(\vec{0}_K, f)} \wedge \alpha_{\Delta_{\vec{K}}(\vec{0}_K, \vec{0}_K, f)ri}^\sigma)).$$

Suppose $\hat{\beta}_{ri}^\sigma$ is not logically equivalent to a universal \mathcal{L} -sentence. Then, in particular the prenex normal form of $\hat{\beta}_{ri}^\sigma$ contains an existential quantifier. Note that μ is the antecedent of an implication, so if the prenex normal form of μ contains no universal quantifiers, then, as η is quantifier-free, this implies there must be $f \in F_{\vec{x}_N}^K$ such that the prenex normal form of $\alpha_{\Delta_{\vec{K}}(\vec{0}_K, \vec{0}_K, f)ri}^\sigma$ contains an existential quantifier. This requires that $r \geq 1$, as when $r = 0$ this formula is just $D_{\Delta_{\vec{K}}(\vec{0}_K, \vec{0}_K, f)}$.

Now, by definition, for all \vec{Z}_k and \vec{Z}'_K we have

$$\alpha_{(\vec{Z}_K, \vec{Z}'_K)ri}^\sigma = \bigwedge_{j \leq i} \forall \vec{y}_{M_j} (\gamma_j(\vec{y}_{M_j}) \rightarrow \bigvee_{f \in F_{\vec{y}_{M_j}}^K} (\psi^j_{\Delta_K(\vec{Z}_K, \vec{Z}'_K, f)} \wedge \alpha_{\Delta_{\vec{K}}(\vec{Z}_K, \vec{Z}'_K, f)(r-1)i}^\sigma)).$$

If for each i the prenex normal form of γ_i contains no universal quantifiers, then, as ψ^i is quantifier-free for all i , for the prenex normal form of $\alpha_{(\vec{Z}_K, \vec{Z}'_K)ri}^\sigma$ to contain an existential quantifier it is necessary that the prenex normal form of $\alpha_{\Delta_{\vec{K}}(\vec{Z}_K, \vec{Z}'_K, f)(r-1)i}^\sigma$ contains an existential quantifier for some $f \in F_{\vec{y}_{M_j}}^K$. But then the same argument applies to $\alpha_{(\vec{Z}_K, \vec{Z}'_K)(r-1)i}^\sigma$, and thus we conclude by induction that $\alpha_{(\vec{Z}'_K, \vec{Z}'_K)0i}^\sigma$ contains an existential quantifier for some \vec{Z}'_K and \vec{Z}'_K , but this is impossible, as it is quantifier-free by definition.

This proves the claim, because it follows that provided the conditions are met, we can obtain a universal axiomatisation by putting every formula $\hat{\beta}_{ri}^\sigma$ into prenex normal form. \dashv

The following lemma articulates an essentially trivial but useful observation.

LEMMA 4.7. *Let \mathcal{B} be a separation subclass of \mathcal{A} defined using the essentially finite separation scheme Σ . Suppose there are natural numbers r' and i' such that, for all $\sigma \in \Sigma^{>0}$ and for all $A \in \mathcal{A}$, the following statement holds :*

- *If \exists has an r' -strategy in the (A, σ) -game where \forall uses moves with index at most i' , then she has an ω -strategy in the usual (A, σ) -game.*

Then the axiomatisations produced in the proof of Corollary 4.6 are equivalent to finite subsets of themselves.

PROOF. By Lemma 4.4, given $\sigma \in \Sigma^{>0}$, \exists having an r -strategy in the (A, σ) -game bounded by i for $A \in \mathcal{A}$ is equivalent to saying that $A \models \hat{\beta}_{ri}^\sigma$. So if r' and i' exist as claimed we have $\hat{\beta}_{r'i'}^\sigma \models \hat{\beta}_{ri}^\sigma$ for all r, i . Thus by Theorem 4.5 we have $A \in \mathcal{B} \iff A \models \hat{\beta}_{r'i'}^\sigma$ for all $\sigma \in \Sigma$. The result follows as Σ is finite. \dashv

EXAMPLE 4.8. Returning to Example 2.6, by Corollary 4.6 we see that the class of representable posets has a recursive axiomatisation (as was proved in [15]). However, the universal quantifiers in the j and m formulas mean that the axiomatisation produced is not universal. Indeed, the class of representable posets has no universal axiomatisation, as it is not closed under substructures (see [15, Corollary 2.9]).

We also note the following alternative approach to constructing a recursive axiomatisation for \mathcal{B} relative to \mathcal{A} when \mathcal{B} is an essentially r.e. separation subclass of \mathcal{A} and \mathcal{A} is elementary. [23, Chapter 9] provides a method for generating a recursive first-order axiomatisation for the elementary closure of any pseudoelementary class whose defining theory in the extended language is recursive. Since in the situation we are describing \mathcal{B} is elementary (by Theorem 2.14), the elementary closure is just the class itself, and, since we have an axiomatisation of \mathcal{B} as an essentially recursive pseudoelementary class by Lemma 2.9, this method can be applied to find a recursive axiomatisation for \mathcal{B} . This method also produces a universal axiomatisation when \mathcal{B} is pseudouniversal, in the sense of [23, Definition 9.1].

Also of interest is the result presented as [23, Theorem 9.14], where it is attributed to Mal'cev and Tarski. According to this theorem, every pseudoelementary class that is closed under ultraroots is elementary, and, moreover, if it is also closed under substructures it is universal. If the pseudoelementary theory is r.e. then so too will be the elementary, or universal, axiomatisations. Appropriate sets of axioms are defined, but not made explicit. The reader is directed to the discussion following [23, Corollary 9.15] for some comments on this.

A notable advantage of the recursive axiomatisation generated in the proof of Theorem 4.5 is that, as it has an explicit connection to \exists 's ability to survive in certain combinatorial games, it can give us some insight into the question of whether an essentially r.e. separation subclass \mathcal{B} of an elementary class \mathcal{A} is finitely axiomatisable relative to \mathcal{A} . To understand how this works, let $T = \{\psi_0, \psi_1, \psi_2, \dots\}$ be the recursive axiomatisation obtained from Theorem 4.5. Then, if \mathcal{B} is finitely axiomatisable relative to \mathcal{A} , there must be some $K \in \omega$ with $A \models \bigwedge_{k=0}^K \psi_k \implies A \models \psi_j$ for all $j \in \omega$, for all $A \in \mathcal{A}$. So, to prove that no such finite axiomatisation exists, it suffices to construct, for each $K \in \omega$, a structure $A_K \in \mathcal{A}$ such that $A_K \models \bigwedge_{k=0}^K \psi_k$, but $A_K \not\models \psi_{K+1}$.

Translating this back into the setting of games, for an essentially finite separation subclass the idea is to construct objects $A_r \in \mathcal{A}$ such that \exists has r -strategies for all simple (A, σ) -games, but not an $(r + 1)$ -strategy for at least one such game. The non-finite case is similar, but we must consider the maximum indices of allowed \forall moves, and we also have to take the possibly infinite number of separation rules of order zero into account. Of course, the substance of any such proof is to be found in the constructions themselves, but this can be a useful approach, where it applies. For example, this method is essentially the engine of the proofs of the titular result of [14], and the results of [24, Section 5], though the work in these examples is




FIGURE 1. Class inclusions for separation subclasses.

phrased in terms of ultraproducts. Note that the argument as described here has an advantage over the originals as reasoning about properties of the ultraproduct is not required. We present a simple application of this technique in Section 6.2.

§5. Expressive power and decision problems. To begin this section we organise our results on the expressive power of the formalism of separation subclasses vis-à-vis first-order logic.

LEMMA 5.1. *There is a basic elementary class \mathcal{A} , and an essentially finite separation subclass \mathcal{B} of \mathcal{A} , such that \mathcal{B} is not finitely axiomatisable.*

PROOF. We have shown that the class of (α, β) -representable posets is an essentially finite separation subclass of the class of posets whenever $2 \leq \alpha, \beta < \omega$ (see Example 2.6), and this class is also known to not be finitely axiomatisable for $\alpha, \beta \geq 3$ [14]. \dashv

PROPOSITION 5.2. *Let \mathcal{A} be an elementary class and make the following definitions:*

- $\mathbf{S}_{\mathcal{A}}$ is the class of separation subclasses of \mathcal{A} .
- $\mathbf{RS}_{\mathcal{A}}$ is the class of essentially r.e. separation subclasses of \mathcal{A} .
- $\mathbf{FS}_{\mathcal{A}}$ is the class of essentially finite separation subclasses of \mathcal{A} .
- $\mathbf{E}_{\mathcal{A}}$ is the class of elementary subclasses of \mathcal{A} .
- $\mathbf{RE}_{\mathcal{A}}$ is the class of subclasses of \mathcal{A} with recursive axiomatisations relative to \mathcal{A} .
- $\mathbf{FE}_{\mathcal{A}}$ is the class of subclasses of \mathcal{A} that are finitely axiomatisable relative to \mathcal{A} .

Then Figure 1 represents the class inclusions that always hold (with arrows from subclass to superclass). In cases where there is no arrow there are choices of \mathcal{A} for which the inclusion does not hold.

PROOF. The horizontal arrows come straight from the definitions, and the lack of backward arrows is also straightforward. The downward arrows come from Corollary 4.6 and Theorem 2.14, and the upward arrows come from Proposition 2.7. The lack of an arrow from $\mathbf{FS}_{\mathcal{A}}$ to $\mathbf{FE}_{\mathcal{A}}$ comes from Lemma 5.1. \dashv

Now we present some easy results on the decision problem for separation subclasses.

DEFINITION 5.3 (Subclass decision problem). Given classes \mathcal{A} and \mathcal{B} with $\mathcal{B} \subseteq \mathcal{A}$, the decision problem for \mathcal{B} relative to \mathcal{A} is the question:

“Given a finite $A \in \mathcal{A}$, is $A \in \mathcal{B}$?”




FIGURE 2. Complexity class inclusions.

PROPOSITION 5.4. *Let \mathcal{A} be a class of \mathcal{L} -structures, and let \mathcal{B} be an essentially r.e. separation subclass of \mathcal{A} . Then the complement to the decision problem for \mathcal{B} relative to \mathcal{A} is semidecidable.*

PROOF. If \mathcal{A} is elementary then \mathcal{B} has a recursive axiomatisation relative to \mathcal{A} , by Corollary 4.6, and the result follows immediately.

Suppose now that \mathcal{A} is not elementary, and that \mathcal{B} is defined by the essentially r.e. separation scheme Σ . Our algorithm is as follows. Given finite $A \in \mathcal{A}$, using dovetailing we work through the separation rules of Σ . If σ is finite then it can be checked directly if $A \models \sigma$. If on the other hand

$$\sigma = \forall \vec{x}_N \left(\mu(\vec{x}_N) \rightarrow \exists \vec{C}_K (\eta(\vec{x}_N) \wedge \bigwedge_I \tau_i) \right)$$

is infinite then it cannot be checked directly if $A \models \sigma$, but, as Σ is essentially r.e., we can assume without loss of generality that $I = \omega$, and for each $n \in \omega$ we can define

$$\sigma_n = \forall \vec{x}_N \left(\mu(\vec{x}_N) \rightarrow \exists \vec{C}_K (\eta(\vec{x}_N) \wedge \bigwedge_{i \leq n} \tau_i) \right).$$

Now with dovetailing we can check if $A \models \sigma_n$ for each $n \in \omega$ and each infinite $\sigma \in \Sigma$. If $A \not\models \sigma$ then there is n with $A \not\models \sigma_n$, so if such a σ exists our algorithm will eventually find it. As soon as σ with $A \not\models \sigma$ is found the algorithm terminates, as this shows $A \notin \mathcal{B}$. \dashv

LEMMA 5.5. *If \mathcal{A} is a class of \mathcal{L} -structures and \mathcal{B} is an essentially finite separation subclass of \mathcal{A} , then the decision problem for \mathcal{B} relative to \mathcal{A} is in \mathbf{NP} .*

PROOF. By Lemma 2.9, an essentially finite separation subclass is essentially finite pseudo-elementary relative to the superclass, and being essentially finite pseudo-elementary is equivalent to being finitely axiomatisable in existential second-order logic. Finally, by Fagin’s Theorem [18], the problem of checking whether a finite structure satisfies an existential second-order sentence is in \mathbf{NP} . \dashv

PROPOSITION 5.6. *Let \mathcal{A} be an elementary class and make the following definitions in addition to those of Proposition 5.2 :*

$\mathbf{P}_{\mathcal{A}}$ is the class of subclasses of \mathcal{A} whose decision problem relative to \mathcal{A} is in \mathbf{P} .

$\mathbf{NP}_{\mathcal{A}}$ is the class of subclasses of \mathcal{A} whose decision problem relative to \mathcal{A} is in \mathbf{NP} .

Then Figure 2 represents the class inclusions that always hold, using the same system as in Figure 1, but with the addition that the existence of a full arrow in either of the places indicated by dotted arrows is equivalent to $\mathbf{P} = \mathbf{NP}$.

PROOF. The arrow from $\mathbf{FS}_{\mathcal{A}}$ to $\mathbf{NP}_{\mathcal{A}}$ comes from Lemma 5.5. To see that there is no arrow from $\mathbf{P}_{\mathcal{A}}$ to $\mathbf{FS}_{\mathcal{A}}$ let \mathcal{A} be, for example, the class of all sets, and let \mathcal{B} be the class of all finite sets. Then the decision problem for \mathcal{B} relative to \mathcal{A} is trivially in \mathbf{P} (as every instance is a yes instance), but finiteness has no first-order characterisation, and thus cannot be formalised as a separation subclass (by Theorem 2.14). That there is no arrow from $\mathbf{NP}_{\mathcal{A}}$ to $\mathbf{FS}_{\mathcal{A}}$ follows immediately.

If there is an arrow from $\mathbf{FS}_{\mathcal{A}}$ to $\mathbf{P}_{\mathcal{A}}$ then, for example, deciding whether a finite poset is $(4, 4)$ -representable is in \mathbf{P} , as the $(4, 4)$ -representable posets are an essentially finite separation subclass of the class of posets (see Example 2.6), and thus $\mathbf{P} = \mathbf{NP}$, as this problem is \mathbf{NP} -complete [39]. Conversely, if $\mathbf{P} = \mathbf{NP}$ then there is an arrow from $\mathbf{NP}_{\mathcal{A}}$ to $\mathbf{P}_{\mathcal{A}}$, and thus an arrow from $\mathbf{FS}_{\mathcal{A}}$ to $\mathbf{P}_{\mathcal{A}}$. \dashv

We have established that every essentially finite separation subclass of an elementary class is recursively axiomatisable relative to the superclass, and also that the converse does not hold in general (see Proposition 5.2). The following result says that the converse still doesn't hold when we restrict to varieties and recursively axiomatisable subvarieties.

PROPOSITION 5.7. *There is a finitely axiomatised variety \mathcal{A} , and a recursively axiomatised subvariety \mathcal{B} of \mathcal{A} such that \mathcal{B} is not an essentially finite separation subclass of \mathcal{A} .*

PROOF. The class \mathbf{RRA} of representable relation algebras is a variety (by [37], or see [23, Theorem 3.37]), and can be recursively axiomatised by equations (see [23, Theorem 8.4]), but the decision problem for \mathbf{RRA} relative to the class of relation algebras, \mathbf{RA} , is not decidable (by [22], or see [23, Theorem 18.13]), and thus cannot be an essentially finite separation subclass (appealing to Lemma 5.5). \dashv

Finally, as in the argument used in the proof of Lemma 5.5, an essentially finite separation subclass \mathcal{B} of a basic elementary class \mathcal{A} can be finitely axiomatised in existential second-order logic. Of course, it follows immediately from Fagin's Theorem and Proposition 5.6 that there are subclasses that are finitely axiomatisable in existential second-order logic relative to their superclasses that cannot be expressed as essentially finite separation subclasses.

§6. Applications. In this section we use the general theory of separation subclasses to get some axiomatisation results in graph theory and theoretical computer science.

6.1. Disjoint union partial algebras. Here we deal with a class of structures introduced in [24].

DEFINITION 6.1. A *partial algebra* is a set equipped with a number of partial operations of fixed arities, and also possibly some constants. In order to accommodate this in first-order logic we think of partial algebras as relational structures, where each n -ary partial operation corresponds to an $(n + 1)$ -ary relation, and for each such relation R we have a sentence

$$\forall x_1 \dots x_n yz ((R(x_1, \dots, x_n, y) \wedge R(x_1, \dots, x_n, z)) \rightarrow y \approx z)$$

expressing that the associated partial function is well defined.

DEFINITION 6.2. A disjoint union partial algebra (DUPA) is a partial algebra with a single ternary relation d (disjoint union). We will usually write $d(x, y, z)$ as $x \dot{\cup} y = z$.

For more on DUPAs and their uses in computer science see [24].

DEFINITION 6.3 (Representable DUPA). A DUPA is representable if it is isomorphic (as a relational structure) to a DUPA whose universe is a set of sets, and whose relation d is defined by

$$d(X, Y, Z) \iff X \cap Y = \emptyset \text{ and } Z = X \cup Y$$

The following is a minor adaptation of [24, Definition 4.1].

DEFINITION 6.4 (Basic sets). If A is a disjoint union partial algebra, and if $\Gamma \subseteq A$, then we say Γ is basic if:

- (1) $a \dot{\cup} b \in \Gamma \implies$ either $a \in \Gamma$ or $b \in \Gamma$.
- (2) If either $a \in \Gamma$ or $b \in \Gamma$, and if $a \dot{\cup} b$ is defined, then $a \dot{\cup} b \in \Gamma$.
- (3) If both $a \in \Gamma$ and $b \in \Gamma$ then $a \dot{\cup} b$ is not defined.

The word ‘basic’ here refers to being part of the ‘base’ of a representation as an algebra of sets.

LEMMA 6.5. If A is a DUPA, then A is representable if and only if:

- (1) For all $a \neq b \in A$, there is a basic $\Gamma \subseteq A$ with either $a \in \Gamma$ and $b \notin \Gamma$, or $b \in \Gamma$ and $a \notin \Gamma$.
- (2) For all $a, b \in A$, if $a \dot{\cup} b$ is undefined then there is a basic $\Gamma \subseteq A$ with $\{a, b\} \subseteq \Gamma$.

PROOF. This is [24, Lemma 4.2]. ⊢

PROPOSITION 6.6. The class of representable DUPAs is an essentially finite separation subclass of the class of all DUPAs.

PROOF. Let $\mathcal{L} = \{d\}$ be the signature of disjoint union partial algebras, and let $\mathcal{L}^+ = \mathcal{L} \cup \{C\}$, where C is a unary predicate symbol. Define the following \mathcal{L}^+ -sentences:

$$\begin{aligned} \tau_0 &= \forall y_1 y_2 y_3 \left(d(y_1, y_2, y_3) \rightarrow (C(y_3) \rightarrow (C(y_1) \vee C(y_2))) \right), \\ \tau_1 &= \forall y_1 y_2 y_3 \left(d(y_1, y_2, y_3) \rightarrow ((C(y_1) \vee C(y_2)) \rightarrow C(y_3)) \right), \\ \tau_2 &= \forall y_1 y_2 y_3 \left(d(y_1, y_2, y_3) \rightarrow (\neg C(y_1) \vee \neg C(y_2)) \right). \end{aligned}$$

Then $\tau = \tau_0 \wedge \tau_1 \wedge \tau_2$ states that the set defined by C is basic, and, moreover, τ is a closure rule as defined in Definition 2.1. Now define

$$\sigma_1 = \forall x_1 x_2 \left(\neg(x_1 \approx x_2) \rightarrow \exists C \left(((C(x_1) \wedge \neg C(x_2)) \vee (C(x_2) \wedge \neg C(x_1))) \wedge \tau \right) \right),$$

and

$$\sigma_2 = \forall x_1 x_2 \left(\neg \exists x_3 d(x_1, x_2, x_3) \rightarrow \exists C (C(x_1) \wedge C(x_2) \wedge \tau) \right).$$

Then σ_1 and σ_2 are separation rules, as defined in Definition 2.2. Moreover, by Lemma 6.5, $\{\sigma_1, \sigma_2\}$ axiomatises the class of representable DUPAs relative to the class of all DUPAs, which is what we are required to prove. ⊢

Having established that the class of representable DUPAs is an essentially finite separation subclass of the class of all DUPAs (which is basic elementary), we can use general results for separation subclasses to easily prove some results that were obtained with more effort in [24]. For example:

COROLLARY 6.7. *The class of representable DUPAs is basic pseudoelementary.*

PROOF. This follows from Proposition 6.6 and Lemma 2.9. ⊢

COROLLARY 6.8. *The class of representable DUPAs has a recursive axiomatisation in first-order logic.*

PROOF. This follows from Proposition 6.6 and Theorem 4.5. ⊢

Note that the appearance of $\neg\exists$ in σ_2 means the recursive axiomatisation generated is not universal. Indeed, by [24, Corollary 3.3] we know that no such universal axiomatisation can exist.

COROLLARY 6.9. *The decision problem for the class of representable DUPAs is in NP.*

PROOF. This follows from Lemma 5.5. ⊢

6.2. *N*-colourable graphs. Here and elsewhere we assume all graphs are undirected and simple. Given $1 \leq N < \omega$, a graph $G = (V, E)$ is *N*-colourable if it is possible to assign to each vertex $v \in V$ one of *N* colours in such a way that no adjacent vertices have the same colour. Equivalently, *G* is *N*-colourable if there is a homomorphism $h : G \rightarrow K_N$ where K_N is the complete graph with *N* vertices. Let \mathcal{G} be the (elementary) class of all graphs, and, given $1 \leq N < \omega$, define \mathcal{G}_N to be the class of *N*-colourable graphs. Note that if *G* is *N*-colourable via $h : G \rightarrow K_N$, and if *H* is any other graph, then the composition of *h* with the projection function, $h \circ \pi_G : G \times H \rightarrow K_N$, is a homomorphism. So, in particular, \mathcal{G}_N is closed under taking direct products for all $1 \leq N < \omega$.

Let $\mathcal{L} = \{E\}$ be the standard signature for graphs (so *E* stands for the binary edge relation). Let $\mathcal{L}^+ = \mathcal{L} \cup \{C_1, \dots, C_N\}$ (here C_n is a monadic predicate intended to pick out the vertices coloured by the *n*th colour), and define

$$\tau_0 = \forall y \left(\top \rightarrow \bigvee_{n=1}^N C_n(y) \right),$$

$$\tau_1 = \forall y \left(\top \rightarrow \bigwedge_{m \neq n} \neg(C_m(y) \wedge C_n(y)) \right),$$

$$\tau_2 = \forall y_1 y_2 \left(E(y_1, y_2) \rightarrow \bigwedge_{n=1}^N \neg(C_n(y_1) \wedge C_n(y_2)) \right),$$

and

$$\sigma = \forall x \left(\top \rightarrow \exists C_1 \dots C_N \left(\top \wedge \tau_0 \wedge \tau_1 \wedge \tau_2 \right) \right).$$




FIGURE 3. A game played on C_5 , with different ‘colours’ being denoted by \diamond and \circ . Here \forall first asks \exists to colour the top vertex, which she does with \diamond (both players making arbitrary choices here). To minimise the maximum length of a chain of uncoloured elements \forall then demands that one of the bottom two vertices be coloured (the left one say). To be consistent with the closest coloured vertex \exists responds by colouring with \diamond . Now \forall demands that either of the vertices between the two coloured ones (moving clockwise) be coloured (the higher one, say). To be consistent with the nearest neighbour \exists must colour it with \circ . Now \forall can force a forbidden colouring in the next round.

Then σ is a separation rule, and if \mathcal{C}_N is the separation subclass of \mathcal{G} defined by $\{\sigma\}$, then \mathcal{C}_N is exactly the class of all N -colourable graphs. Thus we see that \mathcal{C}_N has the various pleasant properties associated with essentially finite separation subclasses of elementary classes. In particular, from Corollary 4.6 we obtain a recursive universal axiomatisation for \mathcal{C}_N as a class of \mathcal{L} -structures. This is not a new result. Indeed, [41, Theorem 1.4] proves that \mathcal{C}_N has a recursive axiomatisation using universal Horn formulas, and that paper attributes to W. Taylor a proof of the same result using the De Bruijn–Erdős theorem for graphs (i.e., that a graph is N -colourable when all its finite subgraphs are) [10].

Now, being universal, \mathcal{C}_N is closed under isomorphisms, substructures, and ultraproducts, and, as the class is also closed under taking direct products, it follows that \mathcal{C}_N is a universal Horn class (see, e.g., [5, Theorem V.2.23]). The universal Horn theory of \mathcal{C}_N must be precisely the universal Horn consequences of our recursive axiomatisation, and so is also a recursively enumerable set, and consequently defines a recursive axiomatisation using Craig’s trick. Thus the *prima facie* stronger result of [41] follows easily from our version, which we got more or less for free from the general theory. Note that such a universal Horn axiomatisation is the best that can be hoped for, as for $N \geq 2$ there can be no finite axiomatisation of \mathcal{C}_N [41, Theorem 1.5]. Of course, \mathcal{C}_1 is just the class of totally disconnected (edgeless) graphs.

Making good on the claims in the comments at the end of Section 4, we can also use our game-generated axioms to find a simple proof that \mathcal{C}_N is not finitely axiomatisable when $N \geq 2$. First, for each $n \geq 1$ consider the cycle graph C_{2n+1} , and consider also the class \mathcal{C}_2 as a separation subclass of \mathcal{G} . Then the number of rounds \exists can guarantee to survive in the simple (C_{2n+1}, σ) game scales linearly with $\log_2 n$. Here \exists ’s strategy is to always colour vertices consistently with their closest neighbour, and \forall ’s best strategy is to make the maximum size of a chain of uncoloured vertices as small as possible each round—see Figure 3 for an illustration. Note that \exists can use her strategy against any strategy used by \forall . The one described is optimal for him in the sense that it wins fastest.

Now, if \mathcal{C}_2 were finitely axiomatisable then a graph would be 2-colourable if and only if \exists could guarantee survival for a fixed finite number of rounds. By choosing n large enough, \exists can find a graph C_{2n+1} where she does have such a strategy, but which is nevertheless not 2-colourable, and this would result in contradiction. Thus the axiomatisation of \mathcal{C}_2 generated by Corollary 4.6 cannot be logically equivalent to a finite subset of itself, and it follows that \mathcal{C}_2 is not finitely axiomatisable. Note that as $C_{2n+1} \in \mathcal{C}_3$ for all n , this argument also shows that \mathcal{C}_2 is not finitely axiomatisable relative to \mathcal{C}_3 .

Generalising, let $N > 2$, and for each $n \geq 1$ define G_n to be the graph obtained by taking the disjoint union of the cycle graph C_{2n+1} and the complete graph K_{N-2} , and adding edges between every vertex of C_{2n+1} and every vertex of K_{N-2} . Then, in the game where \exists attempts to colour G_n using N colours, the choice of colours for K_{N-2} forces her to attempt to colour C_{2n+1} with two colours. We know this is impossible, but the number of rounds she can survive again scales with $\log_2 n$. Here \exists 's strategy is to choose $N - 2$ colours for K_{N-2} , and to use her strategy from the $N = 2$ case for C_{2n+1} with the two remaining colours. Thus \mathcal{C}_N is not finitely axiomatisable for all $N \geq 2$. This provides a proof of [41, Theorem 1.5] that does not use the fact that the class of graphs with chromatic number N is not elementary for all $N \geq 3$ [38, Theorem 6.3]. Note that, combined with the result for $N = 2$, this argument also shows that \mathcal{C}_N is not finitely axiomatisable relative to \mathcal{C}_{N+1} for all $N \geq 2$.

Moreover, let $N \geq 3$ and define χ_N to be the class of graphs with chromatic number N . We can use our results on the lack of a finite axiomatisation for \mathcal{C}_{N-1} relative to \mathcal{C}_N to prove that χ_N is not elementary. To see this, first note that $\chi_N = \mathcal{C}_N \setminus \mathcal{C}_{N-1}$, and so $\mathcal{C}_N = \mathcal{C}_{N-1} \cup \chi_N$, which is a disjoint union. Now, as \mathcal{C}_N and \mathcal{C}_{N-1} are elementary, if χ_N is also elementary then \mathcal{C}_{N-1} will be finitely axiomatisable relative to \mathcal{C}_N , by a variation of the compactness argument that says that if a class and its complement are elementary, then both will be basic elementary. As \mathcal{C}_{N-1} is not finitely axiomatisable relative to \mathcal{C}_N , it follows that χ_N is not elementary. Thus we also obtain an alternative proof of [38, Theorem 6.3] (the original uses Erdős' famous result that for all $m, k \in \omega$ there is a finite graph with chromatic number $\geq m$ and no circuits of length $\leq k$ [17]).

As a final observation, every first-order structure can be embedded into an ultraproduct of its finitely generated substructures (see, e.g., [5, Theorem V.2.14]). Moreover, if a graph G has the property that every finite subgraph is N -colourable, then, as \mathcal{C}_N is elementary, an ultraproduct of these subgraphs must also be N -colourable, by Łoś' theorem. Furthermore, as \mathcal{C}_N is universal, its substructures must also be in \mathcal{C}_N , and so it follows that $G \in \mathcal{C}_N$. Thus, from the axiomatisation of \mathcal{C}_N we also obtain a rather indirect proof of the De Bruijn–Erdős theorem. We must note that much simpler proofs are well known, so this last result is essentially a curiosity.

6.3. Clique covers. Let $N \in \omega$. We say a graph $G = (V, E)$ has an N -clique cover if its vertices can be partitioned into N subsets, each of which is a clique. In other words, if there is a partition V_1, \dots, V_N of V such that the restriction of E to V_n produces a complete graph for all $n \in \{1, \dots, N\}$. Note that a graph G has an N -clique covering if and only if the complement graph $\bar{G} = (V, \bar{E})$ is N -colourable. As we can define \bar{E} as $\neg E$, the results of Section 6.2 apply here, with the following

exception. The class of undirected simple graphs with an N -clique cover is *not* closed under taking direct products for any N . To see this, consider the product of the totally disconnected graph with N vertices with itself. So this class does not have a universal Horn axiomatisation (by [5, Theorem V.2.23] again), though it does have a recursive universal axiomatisation.

6.4. Harmonious colourings. The concept of a harmonious colouring for a graph was introduced in [19] and defined in its current form in [26]. Given $N \in \omega$, we say a graph has a harmonious N -colouring if it has an N -colouring in which each pair of colours can be used to colour a pair of adjacent vertices at most once. That is, if we use $c(v)$ to denote the colour of a vertex, and if v_1, v_2, v_3, v_4 are vertices with $c(v_1) = c(v_3)$ and $c(v_2) = c(v_4)$, then either $v_1 = v_3$ and $v_2 = v_4$, or there cannot be edges between both $\{v_1, v_2\}$ and $\{v_3, v_4\}$. Define \mathcal{L}^+ , τ_0 , τ_1 , and τ_2 as in Section 6.2. In addition, we define a sentence τ_3 expressing that the vertices of two distinct edges cannot be coloured by the same pattern of colours. This sentence says that if $E(v_1, v_2)$ and $E(v_3, v_4)$ (i.e., the edges exist), and if $v_1 \neq v_3$ or $v_2 \neq v_4$ (i.e., the edges are distinct), then either v_1 is not the same colour as v_3 , or v_2 is not the same colour as v_4 .

$$\tau_3 = \forall y_1 y_2 y_3 y_4 \left(\neg((y_1 \approx y_3) \wedge (y_2 \approx y_4)) \wedge E(y_1, y_2) \wedge E(y_3, y_4) \right) \\ \rightarrow \bigwedge_{m,n=1}^N \neg(C_m(y_1) \wedge C_n(y_2) \wedge C_m(y_3) \wedge C_n(y_4)).$$

Now define

$$\sigma = \forall x \left(\top \rightarrow \exists C_1 \dots C_n (\top \wedge \tau_0 \wedge \tau_1 \wedge \tau_2 \wedge \tau_3) \right).$$

Then σ defines the class of graphs with harmonious N -colourings as an essentially finite separation subclass of \mathcal{G} . It again follows from Corollary 4.6 that this class has a universal recursive axiomatisation. Note that when $N \geq 2$ the class is not closed under taking direct products. To see this, note that a graph with a harmonious N -colouring can have at most $\binom{N}{2}$ edges, as this is the maximum number of distinct colour pairs, and consider the product of the complete graphs K_N and K_2 . Each component has a harmonious N -colouring, but the product does not, simply because it has too many edges. So the class does not have a universal Horn axiomatisation. When $N = 1$ the graphs must be totally disconnected just to have an N -colouring, which will be trivially harmonious.

From the fact that a graph with a harmonious N -colouring can have at most $\binom{N}{2}$ edges, we can deduce that the axiomatisation produced here is equivalent to a finite one. Assuming \forall plays in an efficient way, in other words, that he forces \exists to define a new coloured pair each round if possible, he will either definitely be able to force a win in round $(\binom{N}{2} + 1)$ at the latest, or he will have run out of useful moves in an earlier round. So, if \exists has an $(\binom{N}{2} + 1)$ -strategy, then she has an ω -strategy. Appealing to Lemma 4.7 proves the claim.

Note that it is proved in [26] that the problem of deciding, when given a graph G and a positive integer N , whether G has a harmonious colouring with N -colours is NP-complete. As it is known that checking whether a first-order sentence is valid

in a finite structure can be done in polynomial time (see [40, Proposition 3.1]), we may wonder whether we have accidentally proved $\mathbf{P} = \mathbf{NP}$. The answer, sadly, is no, because given (G, N) we have to *construct* the appropriate sentence, which depends on N , before we can check it, and we have no reason to believe we can do this in polynomial time. There is an alternative version of the problem where N is regarded as fixed, and so an instance is just a graph G . In this case our argument does indeed show the decision problem to be in \mathbf{P} , but this version of the problem is not \mathbf{NP} -complete.

Acknowledgments. The author would like to thank Robin Hirsch for, among other things, a suggestion that significantly simplified the key definitions, and the observation described in the preceding paragraph. The author would also like to thank the Department of Computer Science at UCL for hosting him while part of this paper was written. In addition, the author thanks the anonymous referee for their useful comments.

REFERENCES

- [1] A. ABIAN, *Boolean rings with isomorphisms preserving suprema and infima*. *Journal of the London Mathematical Society* (2), vol. 3 (1971), pp. 618–620.
- [2] R. BALBES, *A representation theory for prime and implicative semilattices*. *Transactions of the American Mathematical Society*, vol. 136 (1969), pp. 261–267.
- [3] G. BIRKHOFF, *On the combination of subalgebras*. *Proceedings of the Cambridge Philosophical Society*, vol. 29, (1933), pp. 441–464.
- [4] G. BIRKHOFF and O. FRINK, JR., *Representations of lattices by sets*. *Transactions of the American Mathematical Society*, vol. 64 (1948), pp. 299–316.
- [5] S. BURRIS and H. P. SANKAPPANAVAR, *A Course in Universal Algebra*, Graduate Texts in Mathematics, vol. 78, Springer-Verlag, New York and Berlin, 1981.
- [6] C. C. CHANG, *On the representation of α -complete Boolean algebras*. *Transactions of the American Mathematical Society*, vol. 85 (1957), pp. 208–218.
- [7] C. C. CHANG and A. HORN, *On the representation of α -complete lattices*. *Fundamenta Mathematicae*, vol. 51 (1962/1963), pp. 253–258.
- [8] Y. CHENG and P. KEMP, *Representation of posets*. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, vol. 38 (1992), pp. 269–276.
- [9] W. H. CORNISH and R. C. HICKMAN, *Weakly distributive semilattices*. *Acta Mathematica Academiae Scientiarum Hungaricae*, vol. 32 (1978), no. 1–2, pp. 5–16.
- [10] N. G. DE BRUIJN and P. ERDÖS, *A colour problem for infinite graphs and a problem in the theory of relations*. *Proceedings of the Koninklijke Nederlandse Akademie van Wetenschappen. Series A. 54 = Indagationes Mathematicae*, vol. 13 (1951), pp. 369–373.
- [11] R. EGROT, *Representable posets*. *Journal of Applied Microbiology*, vol. 16 (2016), pp. 60–71.
- [12] _____, *Closure operators, frames and neatest representations*. *Bulletin of the Australian Mathematical Society*, vol. 96 (2017), no. 3, pp. 361–373.
- [13] _____, *Non-elementary classes of representable posets*. *Proceedings of the American Mathematical Society*, vol. 145 (2017), no. 11, pp. 4675–4685.
- [14] _____, *No finite axiomatizations for posets embeddable into distributive lattices*. *Annals of Pure and Applied Logic*, vol. 169 (2018), no. 3, pp. 235–242.
- [15] _____, *Recursive axiomatizations for representable posets*. *International Journal of Algebra and Computation*, vol. 29 (2019), no. 4, pp. 699–711.
- [16] R. EGROT and R. HIRSCH, *Completely representable lattices*. *Algebra Universalis*, vol. 67 (2012), pp. 205–217.
- [17] P. ERDÖS, *Graph theory and probability*. *Canadian Journal of Mathematics*, vol. 11 (1959), pp. 34–38.

- [18] R. FAGIN, *Contributions to the model-theory of finite-structures* (ProQuest LLC, Ann Arbor, MI), Ph.D. thesis, University of California, Berkeley, 1973.
- [19] O. FRANK, F. HARARY, and M. PLANTHOLT, *The line-distinguishing chromatic number of a graph*. *Ars Combinatoria*, vol. 14 (1982), pp. 241–252.
- [20] T. FRAYNE, A. C. MOREL, and D. S. SCOTT, *Reduced direct products*. *Fundamenta Mathematicae*, vol. 51 (1962/1963), pp. 195–228.
- [21] R. C. HICKMAN and G. P. MONRO, *Distributive partially ordered sets*. *Fundamenta Mathematicae*, vol. 120 (1984), no. 2, pp. 151–166.
- [22] R. HIRSCH and I. HODKINSON, *Representability is not decidable for finite relation algebras*. *Transactions of the American Mathematical Society*, vol. 353 (2001), no. 4, pp. 1403–1425.
- [23] ———, *Relation Algebras by Games*. North-Holland, Amsterdam, 2002.
- [24] R. HIRSCH and B. MCLEAN, *Disjoint-union partial algebras*. *Logical Methods in Computer Science*, vol. 13 (2017), no. 2, Article no. 10, pp. 1–31.
- [25] W. HODGES, *Model Theory*, Encyclopedia of Mathematics and Its Applications, vol. 42, Cambridge University Press, Cambridge, 1993.
- [26] J. E. HOPCROFT and M. S. KRISHNAMOORTHY, *On the harmonious coloring of graphs*. *SIAM Journal on Algebraic and Discrete Methods*, vol. 4 (1983), no. 3, pp. 306–311.
- [27] K. A. KEARNES, *The class of prime semilattices is not finitely axiomatizable*. *Semigroup Forum*, vol. 55 (1997), no. 1, pp. 133–134.
- [28] H. KEISLER, *Ultraproducts and elementary models*. *Indagationes Mathematicae*, vol. 23 (1961), pp. 477–495.
- [29] P. KEMP, *Representation of partially ordered sets*. *Algebra Universalis*, vol. 30 (1993), pp. 348–351.
- [30] D. KÖNIG, *Sur les correspondances multivoques des ensembles*. *Fundamenta Mathematicae*, vol. 8, (1926), pp. 114–134.
- [31] L. LOOMIS, *On the representation of σ -complete Boolean algebras*. *Bulletin of the American Mathematical Society*, vol. 53 (1947), pp. 757–760.
- [32] J. ŁOŚ, *Quelques remarques, théorèmes et problèmes sur les classes définissables d'algèbres*, *Mathematical Interpretation of Formal Systems*, North-Holland, Amsterdam, 1955, pp. 98–113.
- [33] B. M. SCHEIN, *On the definition of distributive semilattices*. *Algebra Universalis*, vol. 2 (1972), pp. 1–2.
- [34] S. SHELAH, *Every two elementarily equivalent models have isomorphic ultrapowers*. *Israel Journal of Mathematics*, vol. 10 (1971), pp. 224–233.
- [35] R. SIKORSKI, *On the representation of Boolean algebras as fields of sets*. *Fundamenta Mathematicae*, vol. 35 (1948), pp. 247–258.
- [36] M. STONE, *The theory of representations for Boolean algebras*. *Transactions of the American Mathematical Society*, vol. 40 (1936), pp. 37–111.
- [37] A. TARSKI, *Contributions to the theory of models, III*. *Proceedings of the Koninklijke Nederlandse Akademie van Wetenschappen*, vol. 58 (1955), pp. 56–64.
- [38] W. TAYLOR, *Atomic compactness and graph theory*. *Fundamenta Mathematicae*, vol. 65 (1969), pp. 139–145.
- [39] C. J. VAN ALTEN, *Embedding ordered sets into distributive lattices*. *Order*, vol. 33 (2016), no. 3, pp. 419–427.
- [40] M. VARDI, *On the complexity of bounded-variable queries*. *ACM Symposium on Principles of Database Systems*, ACM Press, New York, May 1995, pp. 266–276.
- [41] W. H. WHEELER, *The first order theory of N -colorable graphs*. *Transactions of the American Mathematical Society*, vol. 250 (1979), pp. 289–310.

FACULTY OF INFORMATION AND COMMUNICATION TECHNOLOGY
 MAHIDOL UNIVERSITY
 999 PHUTTHAMONTHON 4 RD, SALAYA
 NAKHON PATHOM 73170, THAILAND
 E-mail: robert.egr@mahidol.ac.th