

MINIMAL FUNCTIONS ON THE RANDOM GRAPH

MANUEL BODIRSKY AND MICHAEL PINSKER

ABSTRACT. We show that there is a system of 14 non-trivial finitary functions on the random graph with the following properties: Any non-trivial function on the random graph generates one of the functions of this system by means of composition with automorphisms and by topological closure, and the system is minimal in the sense that no subset of the system has the same property. The theorem is obtained by proving a Ramsey-type theorem for colorings of tuples in finite powers of the random graph, and by applying this to find regular patterns in the behavior of any function on the random graph. As model-theoretic corollaries of our methods we rederive a theorem of Simon Thomas classifying the first-order closed reducts of the random graph, and prove some refinements of this theorem; also, we obtain a classification of the maximal reducts closed under primitive positive definitions, and prove that all reducts of the random graph are model-complete.

1. INTRODUCTION

1.1. The random graph. The *random graph* (also called the *Rado graph*) is the countably infinite graph $G = (V; E)$ defined uniquely up to isomorphism by the *extension property*: for all finite disjoint subsets U, U' of the countably infinite vertex set V there exists a vertex $v \in V \setminus (U \cup U')$ such that v is in G adjacent to all vertices in U and to no vertex in U' . Alternatively, G is the unique countable graph which is *universal* in the sense that it contains all finite graphs as induced subgraphs, and *homogeneous* in the sense that any isomorphism between finite induced subgraphs of G extends to an automorphism of G . For the many remarkable properties of G and its automorphism group $\text{Aut}(G)$, and various connections to many branches of mathematics, see e.g. [18, 19].

1.2. Minimal functions. We say that a finitary operation $f: V^k \rightarrow V$ *generates* an operation $g: V^l \rightarrow V$ iff g is contained in the topological closure of the set of term functions that can be built from f and the automorphisms of G , where the topology on functions is just the pointwise convergence topology – we refer to Section 2 for a more technical definition.

By this relation of function generation, the functions on the random graph are quasi-ordered with respect to their “generating strength”. The weakest functions in this order are the *trivial functions*, which we define to be those functions that are generated by the identity $\text{id}: V \rightarrow V$; these trivial functions are the self-embeddings of G with possibly additional

Date: December 4, 2012.

2000 Mathematics Subject Classification. Primary 03C10; secondary 05C80; 08A35; 05C55; 03C40.

Key words and phrases. random graph, ω -categoricity, minimal function, Ramsey theory, automorphism, endomorphism, polymorphism, local clone, model-completeness, first-order definition, existential positive definition, primitive positive definition.

The research leading to these results has received funding from the European Research Council under the European Community’s Seventh Framework Programme (FP7/2007-2013 Grant Agreement no. 257039). The second author is grateful for support through Erwin Schrödinger Fellowship J2742-N18 of the Austrian Science Fund and through an APART-fellowship of the Austrian Academy of Sciences.

dummy variables. On the next level are the *minimal functions*: An operation f is called *minimal* iff it is non-trivial, and all non-trivial functions g it generates have at least the arity of f and generate f . Simon Thomas proved in [32] that there are exactly two minimal unary bijective operations on the random graph that do not generate each other. In this paper we generalize this to arbitrary finitary operations, and show that there are exactly 14 minimal operations on G that do not generate each other: these are a constant operation, an operation that maps G injectively to a complete subgraph of G , an operation that maps G injectively to an independent subset of G , the two operations considered by Thomas, and nine binary injective operations.

1.3. Ramsey theory. In our proof, we apply in a systematic way structural Ramsey theory. Any function $f: V \rightarrow V$ induces a coloring of the edges of the random graph by three colors: each edge might either be sent to an edge, to a non-edge, or be collapsed to a single vertex. Similarly, f induces a coloring of the non-edges. If f is not unary, but a function from a power V^k to V , then it induces colorings of pairs of elements of V^k of a fixed type, and so on. We will use a theorem of Nešetřil and Rödl from [26, 27] (and independently by [1]) which states that finite ordered vertex-colored graphs form a Ramsey class in order to prove a Ramsey-type theorem which in turn allows us to find regular patterns in these colorings, for any function f . This makes it feasible to understand the generating process of functions; in particular, all minimal functions turn out to have *canonical* behavior in the sense that seen from the right perspective, the colorings they induce are all constant.

1.4. Universal algebra: Clones. In universal algebra, a *clone* on a set D is a subset of the set \mathcal{O} of all finitary operations on D which is closed under composition of operations and which contains the projections (see e.g. [29, 31]). The set of all clones over D forms a complete lattice with respect to set-theoretical inclusion; see [21] for a survey of results on this lattice for infinite D . In many applications, however, e.g., in theoretical computer science, one does not study the lattice of all clones, but the smaller lattice of those clones which are closed in the pointwise convergence topology on \mathcal{O} . This topology is given by the countable basis of sets of the form

$$\mathcal{O}_A^s := \{f \in \mathcal{O} \mid f|_A = s\},$$

where $A \subseteq D^n$ is finite and $s: A \rightarrow D$ is a function. Clones which are closed subsets of \mathcal{O} in this topology are called *locally closed*, or just *local*. The importance of such clones is reminiscent of that of closed permutation groups (rather than arbitrary permutation groups) for some applications (see e.g. [17]).

The lattice of local clones has been studied in [28], and it turns out to be quite complicated. However, one is often interested in specific parts of this lattice, in particular the interval of all local clones that contain $\text{Aut}(\Gamma)$, for a structure Γ . When Γ is ω -categorical, i.e., every countable model of the first-order theory of Γ is isomorphic to Γ , then it turns out that techniques similar to those for clones over finite domains can be applied for the study of this interval, in particular when classifying its atoms [2, 4].

If we consider the lattice of local clones containing $\text{Aut}(G)$, then it is easy to see that a clone \mathcal{C} is an atom in this lattice iff there exists a minimal operation f on G such that \mathcal{C} is the smallest local clone containing the set $\{f\} \cup \text{Aut}(G)$; we then say that \mathcal{C} is the local clone generated by $\{f\}$ over $\text{Aut}(G)$. Hence, in this paper we determine all atoms of the lattice of local clones containing $\text{Aut}(G)$.

1.5. Groups and monoids. Similarly to clones, the (topologically) closed permutation groups containing $\text{Aut}(G)$ form a complete lattice, with the meet of a set of groups being their intersection; so do the closed transformation monoids containing $\text{Aut}(G)$. By determining the minimal functions on the random graph we find the atoms not only of the lattice of local clones containing $\text{Aut}(G)$, but also of the corresponding group and monoid lattices. In the group case, it turns out that if one continues “climbing up” in the lattice, i.e., if after the atoms of the lattice one determines the next level and so on, one finds the whole lattice as the lattice has only five elements. This was shown already by Thomas in [32], and we will rederive this result. Our methods also allow to follow the same strategy for the other two lattices, but the iteration does not terminate as these lattices have infinite height.

1.6. Model theory: Reducts of the random graph. Results about operations on the random graph G yield model-theoretic results about *reducts* of G , i.e., about relational structures with the same domain as G whose relations have a first-order definition in G ; this is particularly true because G is ω -categorical. In fact, if we consider two reducts equivalent iff they first-order define one another, then the lattice of all reducts, factored by this equivalence, is antiisomorphic to the lattice of closed permutation groups that contain $\text{Aut}(G)$. Similarly, the finer lattice of reducts up to *existential positive interdefinability* corresponds to the lattice of closed transformation monoids that contain $\text{Aut}(G)$. Finally, the lattice of reducts factored by the even finer equivalence of *primitive positive interdefinability* (a first-order formula is *primitive positive* iff it contains no negations, disjunctions, and universal quantifications), corresponds to the lattice of closed clones that contain $\text{Aut}(G)$. Using the latter connection, we obtain a list of the dual atoms in the lattice of reducts up to primitive positive interdefinability; there are 14 such dual atoms, each corresponding to one of the minimal operations mentioned above.

As another application of the techniques in this paper we rederive the full result of Thomas from [32], which is in fact a classification of the reducts of G up to *first-order interdefinability*. We show that the result can be strengthened to obtain a classification of those structures up to *existential interdefinability*. Finally, we show that all reducts Γ of G have a model-complete theory.

1.7. Computational complexity: Constraint satisfaction. Many computational problems in theoretical computer science can be elegantly formalized in the following way. Fix a structure Γ with finite relational signature. Then the *constraint satisfaction problem for Γ* ($\text{CSP}(\Gamma)$) is the problem of deciding whether a given primitive positive sentence is true in Γ . The computational complexity of $\text{CSP}(\Gamma)$ has been determined for all two-element structures in [30], for all three-element structures in [15], and for all structures with a first-order definition in $(\mathbb{Q}; <)$ in [7]. The results of the present paper provide the necessary mathematical techniques for a complexity classification for $\text{CSP}(\Gamma)$ when Γ has a first-order definition in the random graph [10]; such constraint satisfaction problems constitute a generalization of Boolean constraint satisfaction problems to the language of graphs. We have to refer to the introduction of [4], the survey paper [9], or the papers [12] and [11] for a more detailed description of the general connection of reducts of ω -categorical structures with the CSP.

1.8. Structure of this paper. This introduction will be followed by Section 2 in which we present our results on functions on the random graph in full detail. We then discuss these results from a model-theoretic perspective and draw some corollaries in Section 3; this section can be skipped by anyone not interested in the matter. The proof of the results in Section 2

starts with Section 4, where we recall some Ramsey-type theorems and extend them for our purposes. We then apply these theorems to mappings from V to V in order to get hold of such mappings in Section 5. This allows us to determine the minimal unary functions in Section 6. Turning to functions of higher arity in Section 7, we show that minimal higher arity functions are always binary injections. In order to understand binary minimal functions, we develop further Ramsey-theoretic tools in Section 8. Finally, in Section 9, we determine the minimal binary injections, completing our proof.

2. RESULTS

2.1. The minimal functions result. When $f: V^n \rightarrow V$ and $g_1, \dots, g_n: V^m \rightarrow V$ are operations, then the *composition* of f with g_1, \dots, g_n is the operation defined by

$$(x_1, \dots, x_m) \mapsto f(g_1(x_1, \dots, x_m), \dots, g_n(x_1, \dots, x_m)).$$

An operation $f: V^n \rightarrow V$ is called a *projection* iff there exists $1 \leq i \leq n$ such that $f(x_1, \dots, x_n) = x_i$ for all $x_1, \dots, x_n \in V$.

The following definition of the notion *generates* is equivalent to the one in the introduction.

Definition 1. Let $f: V^k \rightarrow V$ and $g: V^l \rightarrow V$. We say that f *generates* g iff for every finite subset S of V^l the restriction of g to S equals the restriction to S of an l -ary operation that can be obtained from f , automorphisms of G , and the projections by a sequence of compositions of operations.

We also give the definition of a *minimal function* in full detail.

Definition 2. Let $f: V^k \rightarrow V$ and $g: V^l \rightarrow V$.

- f and g are *equivalent* iff f generates g and g generates f .
- g is *trivial* iff it is equivalent to the identity function on V .
- f is *minimal* iff it is not trivial, and all non-trivial functions g generated by f have arity at least k and are equivalent to f .

We now define a small number of special operations on G . The random graph contains all countable graphs as induced subgraphs, and in particular, it contains an infinite complete subgraph, denoted by K_ω . It follows from the homogeneity of G that all injective operations from V to V whose image induces K_ω in G generate each other. Let e_E be one such injective operation.

We define $N := \{(x, y) \in V^2 \mid (x, y) \notin E \wedge x \neq y\}$. Pairs $\{x, y\}$ with $(x, y) \in N$ are referred to as *non-edges*. G contains an infinite independent set, denoted by I_ω . Let e_N be an injective operation from V to V whose image induces I_ω in G .

It is clear that the complement graph of G , i.e., the graph on V obtained by flipping all edges and non-edges of G , is isomorphic to G . Again, note that by homogeneity of G all isomorphisms between G and its complement generate each other. Let $-$ be one such isomorphism. In formulas, we will write $-x$ for $-(x)$.

For any finite non-empty subset S of V , if we flip edges and non-edges between S and $V \setminus S$ in G , then the resulting graph is isomorphic to G (it is straightforward to verify the extension property). All such isomorphisms generate each other. For each non-empty finite S , we let i_S be such an isomorphism. We also write sw for $i_{\{0\}}$, where $0 \in V$ is a fixed element for the rest of the paper, and refer to this operation as the *switch*.

To take a break from the definitions, we state the first part of our main theorem, which characterizes the minimal unary functions on G .

Theorem 3. *Any minimal unary function on G is equivalent to exactly one of the following operations:*

- (1) a constant operation;
- (2) e_N ;
- (3) e_E ;
- (4) $-$;
- (5) sw.

We now turn to minimal functions of higher arity.

Definition 4. Let $f: V^2 \rightarrow V$ be a binary injective operation.

The dual f^* of f is defined by $f^*(x, y) = -f(-x, -y)$.

We say that $f: V^2 \rightarrow V$ is

- of type p_1 iff for all $x_1, x_2, y_1, y_2 \in V$ with $x_1 \neq x_2$ and $y_1 \neq y_2$ we have $(f(x_1, y_1), f(x_2, y_2)) \in E$ if and only if $(x_1, x_2) \in E$;
- of type max iff for all $x_1, x_2, y_1, y_2 \in V$ with $x_1 \neq x_2$ and $y_1 \neq y_2$ we have $(f(x_1, y_1), f(x_2, y_2)) \in E$ if and only if $(x_1, x_2) \in E$ or $(y_1, y_2) \in E$;
- balanced in the first argument iff for all $x_1, x_2, y \in V$ with $x_1 \neq x_2$ we have $(f(x_1, y), f(x_2, y)) \in E$ if and only if $(x_1, x_2) \in E$;
- balanced in the second argument iff $(x, y) \mapsto f(y, x)$ is balanced in the first argument;
- E -dominated in the first argument iff for all $x_1, x_2, y \in V$ with $x_1 \neq x_2$ we have that $(f(x_1, y), f(x_2, y)) \in E$;
- E -dominated in the second argument iff $(x, y) \mapsto f(y, x)$ is E -dominated in the first argument.

We can now state our main result.

Theorem 5. *Any minimal function on G is equivalent to one of the unary operations in Theorem 3, or to exactly one of the following operations:*

- (6) a binary injection of type p_1 that is balanced in both arguments;
- (7) a binary injection of type max that is balanced in both arguments;
- (8) a binary injection of type max that is E -dominated in both arguments;
- (9) a binary injection of type p_1 that is E -dominated in both arguments;
- (10) a binary injection of type p_1 that is balanced in the first and E -dominated in the second argument;

or to one of the duals of the last four operations (the dual of the operation in (6) is equivalent to the operation itself).

2.2. Results on groups and monoids. The technique to show this result can be applied several times to unary bijective operations in order to rederive a result by Simon Thomas (Theorem 6 below). A permutation group \mathcal{G} acting on a set D is called *(locally) closed* iff it is closed in the space of all permutations on D equipped with the pointwise convergence topology; equivalently, \mathcal{G} contains all permutations which can be interpolated by elements of \mathcal{G} on arbitrary finite subsets of D . We call the smallest closed group containing a set of permutations \mathcal{F} on V as well as $\text{Aut}(G)$ the *group generated by \mathcal{F}* .

Theorem 6 (from [32]). *The closed permutation groups containing $\text{Aut}(G)$ are precisely the following.*

- (1) $\text{Aut}(G)$;

- (2) *the group generated by $\{-\}$;*
- (3) *the group generated by $\{\text{sw}\}$;*
- (4) *the group generated by $\{-, \text{sw}\}$;*
- (5) *the group of all permutations on V .*

The arguments given in [32] use a Ramsey-theoretic result by Nešetřil [24], namely that the class of all finite graphs excluding finite cliques of a fixed size forms a *Ramsey class* (in the sense of [25]). We also use a Ramsey-theoretic result, shown by Rödl and Nešetřil [26,27] (and independently by [1]), which is different: we need the fact that finite ordered vertex-colored graphs form a Ramsey class.

Similarly to groups and clones, a monoid \mathcal{M} of operations from a set D to D is called (*locally*) *closed* iff it is closed in the space D^D equipped with the pointwise convergence topology. Our proof moreover shows the following statement about closed transformation monoids that contain $\text{Aut}(G)$; this statement also follows from another combinatorial proof of Simon Thomas given in [33] (where he uses the notion of *pseudo-reducts* instead of a formulation in terms of closed monoids).

Theorem 7. *For any closed monoid \mathcal{M} containing $\text{Aut}(G)$, one of the following cases applies.*

- (1) *\mathcal{M} contains a constant operation.*
- (2) *\mathcal{M} contains e_E .*
- (3) *\mathcal{M} contains e_N .*
- (4) *The permutations in \mathcal{M} form a group which is a dense subset of \mathcal{M} in the space V^V .*

3. MODEL-THEORETIC COROLLARIES

We now discuss a model-theoretic interpretation of these results as well as further model-theoretic consequences; this section can be skipped without affecting readability of the rest of the paper.

Since G is homogeneous in a finite language it is ω -categorical (Corollary 6.4.2 of [22]). The reducts of a countable ω -categorical structure Γ are ω -categorical (see e.g. [22]). In particular, this is true for all reducts of G .

We say that two structures Γ and Δ on the same domain are *first-order interdefinable* when Γ is first-order definable in Δ and vice versa. Let $f: D^n \rightarrow D$ be an operation and let $R \subseteq D^m$ be a relation. For tuples $r_1, \dots, r_n \in D^m$ we write $f(r_1, \dots, r_n)$ for the m -tuple that is obtained by applying f to r_1, \dots, r_n componentwise, i.e., for the m -tuple whose i -th component is $f(r_1^i, \dots, r_n^i)$, where r_j^i denotes the i -th component of r_j , for $1 \leq j \leq n$ and $1 \leq i \leq m$. We say that f *preserves* R iff $f(r_1, \dots, r_n) \in R$ whenever $r_1, \dots, r_n \in R$, and that f *violates* R otherwise. The theorem of Engeler, Ryll-Nardzewski, and Svenonius (see e.g. [22, Theorem 6.3.1]) implies that a relation R is first-order definable in a countable ω -categorical structure Δ if and only if R is preserved by all automorphisms of Δ . As a consequence, the reducts of a countable ω -categorical structure Δ are, up to first-order interdefinability, in one-to-one correspondence with the locally closed permutation groups containing $\text{Aut}(\Delta)$. To illustrate this, we restate Theorem 6 by means of this connection.

On the random graph, let $R^{(k)}$ be the k -ary relation that holds on $x_1, \dots, x_k \in V$ iff x_1, \dots, x_k are pairwise distinct, and the number of edges between these k vertices is odd. Note that $R^{(4)}$ is preserved by $-$, $R^{(3)}$ is preserved by sw , and that $R^{(5)}$ is preserved by $-$ and by sw , but not by all permutations of V .

Theorem 8 (Simon Thomas [32]). *Any reduct of G is first-order interdefinable with precisely one of the following structures.*

- (1) $G = (V; E)$;
- (2) $(V; R^{(4)})$;
- (3) $(V; R^{(3)})$;
- (4) $(V; R^{(5)})$;
- (5) $(V; =)$.

For any reduct Γ of G , a case of Theorem 8 applies iff the case with the same number applies for $\text{Aut}(\Gamma)$ in Theorem 6. We will not prove this relational description in this paper; however, given Theorem 6 and the discussion above, verifying the equivalence is merely an exercise.

In the same way as automorphisms of a countable ω -categorical structure Δ can be used to characterize first-order definability in Δ , self-embeddings of Δ (that is, embeddings of Δ into itself) can be used to characterize existential definability, endomorphisms of Δ can be used to characterize existential positive definability, and polymorphisms of Δ (i.e., homomorphisms from a finite power Δ^n to Δ , or simply finitary operations preserving all relations of Δ) can be used to characterize primitive positive definability in Δ .

A first-order formula ϕ is called *existential* iff it is of the form $\exists x_1, \dots, x_k. \psi$, where ψ is quantifier-free. If ψ is even of the form $\psi_1 \wedge \dots \wedge \psi_m$ for atomic formulas ψ_1, \dots, ψ_m , then ϕ is called *primitive positive*. A formula is called *existential positive* iff it is a disjunction of primitive positive formulas. Call two structures Γ and Δ *primitive positive interdefinable* iff every relation in Γ has a definition by a primitive positive formula in Δ and vice versa; we have analogous definitions for existential positive and existential interdefinability. To translate results about operations on G into results about primitive positive definability in reducts of G , the following theorem is central.

Theorem 9 (from [8]). *Let Γ be a countable ω -categorical structure. Then a relation R is primitive positive definable in Γ if and only if R is preserved by all polymorphisms of Γ .*

The operational generating process can be linked to preservation of relations of structures [2, 4, 8].

Proposition 10. *Let $f: V^k \rightarrow V$ and $g: V^l \rightarrow V$ be operations. Then f generates g if and only if every relation with a first-order definition in G that is preserved by g is also preserved by f .*

Recall from Subsection 1.6 that the reducts of G , factored by the equivalence of primitive positive interdefinability, form a complete lattice in which the order is given by primitive positive definability. Using Theorem 9 and Proposition 10, we obtain the following equivalent formulation of Theorem 5. The *dual* of Γ is the structure that consists of the relations $-R := \{(-t_1, \dots, -t_k) \mid (t_1, \dots, t_k) \in R\}$ for all relations R in Γ .

Corollary 11. *Let Γ be a member of a dual atom of the lattice of reducts of G up to primitive positive interdefinability. Then it is primitive positive interdefinable with exactly one of the following 14 structures, namely the structures with all relations that are first-order definable in G and preserved by*

- (1) a constant operation;
- (2) e_N ;

- (3) e_E ;
- (4) $-$;
- (5) sw ;
- (6) a binary operation of type p_1 that is balanced in both arguments;
- (7) a binary operation of type max that is balanced in both arguments;
- (8) a binary operation of type max that is E -dominated in both arguments;
- (9) a binary operation of type p_1 that is E -dominated in both arguments;
- (10) a binary operation of type p_1 that is balanced in the first and E -dominated in the second argument;

or to one of the duals of the last four structures.

Existential positive and existential definability in a countable ω -categorical structure Γ can be described in terms of the endomorphism monoid of Γ .

Proposition 12. *A relation R has an existential positive (existential) definition in a countable ω -categorical structure Γ if and only if R is preserved by the endomorphisms (self-embeddings) of Γ .*

Proof. It is easy to verify that existential positive formulas are preserved by endomorphisms, and existential formulas are preserved by self-embeddings of Γ .

For the other direction, note that the endomorphisms and self-embeddings of Γ contain the automorphisms of Γ , and hence the theorem of Ryll-Nardzewski shows that R has a first-order definition in Γ ; let ϕ be a formula defining R . Suppose for contradiction that R is preserved by all endomorphisms of Γ but has no existential positive definition in Γ . We use the homomorphism preservation theorem (see [22, Section 5.5, Exercise 2]), which states that a first-order formula ϕ is equivalent to an existential positive formula modulo a first-order theory T if and only if ϕ is preserved by all homomorphisms between models of T . Since by assumption ϕ is not equivalent to an existential positive formula in Γ , there are models Γ_1 and Γ_2 of the first-order theory of Γ and a homomorphism h from Γ_1 to Γ_2 that violates ϕ . By the Theorem of Löwenheim-Skolem (see e.g. [22]) the first-order theory of the two-sorted structure $(\Gamma_1, \Gamma_2; h)$ has a countable model $(\Gamma'_1, \Gamma'_2; h')$. Since both Γ'_1 and Γ'_2 must be countably infinite, and because Γ is ω -categorical, we have that Γ'_1 and Γ'_2 are isomorphic to Γ , and h' can be seen as an endomorphism of Γ that violates ϕ ; a contradiction.

The argument for existential definitions and self-embeddings is similar, but instead of the homomorphism preservation theorem we use the Theorem of Łos-Tarski which states that a first-order formula ϕ is equivalent to an existential formula modulo a first-order theory T if and only if ϕ is preserved by all embeddings between models of T (see e.g. [22, Corollary 5.4.5]). \square

Using Proposition 12, we obtain an interesting and perhaps surprising consequence of Theorem 7. A theory T is called *model-complete* iff every embedding between models of T is elementary, i.e., preserves all first-order formulas. It is well-known that a theory T is model-complete if and only if every first-order formula is modulo T equivalent to an existential formula (see [22, Theorem 7.3.1]). A structure is said to be model-complete iff its first-order theory is model-complete. From the definition of model-completeness and ω -categoricity it is easy to see that a countable ω -categorical structure Γ is model-complete iff all self-embeddings of Γ preserve all first-order formulas. We write $\text{Emb}(\Gamma)$ for the monoid of all self-embeddings of Γ .

Lemma 13. *A countable ω -categorical structure Γ is model-complete if and only if $\text{Aut}(\Gamma)$ is dense in $\text{Emb}(\Gamma)$.*

Proof. First assume that all self-embeddings of Γ are in the topological closure of $\text{Aut}(\Gamma)$. Let ϕ be a first-order formula. By the equivalent characterization of model-completeness mentioned above it suffices to show that ϕ is equivalent to an existential formula. Since ϕ is preserved by automorphisms of Γ , it is also preserved by self-embeddings of Γ . Then Proposition 12 implies that ϕ is equivalent to an existential formula.

Conversely, suppose that all first-order formulas are equivalent to an existential formula in Γ . Since existential formulas are preserved by self-embeddings of Γ , also the first-order formulas are preserved by self-embeddings of Γ . Then the theorem of Engeler, Ryll-Nardzewski, and Svenonius shows that every relation that is preserved by all automorphisms of Γ is also preserved by the self-embeddings of Γ . Now if there were a self-embedding e not in the closure of $\text{Aut}(\Gamma)$, then there would be a finite tuple t in Γ such that $e(t) \neq \alpha(t)$ for all $\alpha \in \text{Aut}(\Gamma)$. Let $R := \{\alpha(t) \mid \alpha \in \text{Aut}(\Gamma)\}$. Then R is preserved by all automorphisms of Γ but not by e , a contradiction. \square

It follows from a result in [7, Proposition 19] (based on a proof of a result by Cameron [16] from [23]) that all reducts of the linear order of the rationals $(\mathbb{Q}; <)$ are model-complete. We now see that the same is true for the random graph. Recall that the homogeneity of G implies that it has *quantifier-elimination*: every first-order formula is in G equivalent to a quantifier-free first-order formula.

Corollary 14. *All reducts of the random graph are model-complete.*

Proof. Let Γ be a reduct. We apply Theorem 7 to $\text{Emb}(\Gamma)$. If Case (4) of the theorem holds, then we are done by Lemma 13. Note that $\text{Emb}(\Gamma)$ cannot contain a constant operation as all its operations are injective. So suppose that $\text{Emb}(\Gamma)$ contains e_N (the argument for e_E is analogous). Let R be any relation of Γ , and ϕ_R be its defining quantifier-free formula. Let ψ_R be the formula obtained by replacing all occurrences of E by *false*; so ψ_R is a formula over the empty language. Then a tuple a satisfies ϕ_R in G iff $e_N(a)$ satisfies ϕ_R in G (because e_N is an embedding) iff $e_N(a)$ satisfies ψ_R in G (as there are no edges on $e_N(a)$) iff $e_N(a)$ satisfies ψ_R in the substructure induced by $e_N[V]$ (since ψ_R does not contain any quantifiers). Thus, Γ is isomorphic to the structure on $e_N[V]$ which has the relations defined by the formulas ψ_R ; hence, Γ is isomorphic to a structure with a first-order definition over the empty language. This structure has, of course, all injections as self-embeddings, and all permutations as automorphisms, and hence is model-complete by Lemma 13; thus, the same is true for Γ . \square

Although G has quantifier-elimination, the same is not true for its reducts. For example, any two 2-element substructures of the structure

$$\Gamma = (V; \{(x, y, z) \mid (x, y) \in E \wedge (y, z) \in N\})$$

are isomorphic. But since there is a first-order definition of G in Γ , an isomorphism between a 2-element substructure with an edge and a 2-element substructure without an edge cannot be extended to an automorphism of Γ . However, our results imply that a structure Γ with a first-order definition in the random graph is homogeneous when Γ is expanded by all relations with an existential definition in Γ .

As another application, we refine Theorem 8 by giving a finer (at least in theory) classification of the reducts of the random graph.

Corollary 15. *Up to existential interdefinability, the random graph has exactly five reducts.*

Proof. In the same way as in the proof of Corollary 14, we can use Theorem 7 to show that either the self-embeddings of a reduct Γ are generated by the automorphisms, and Γ is existentially interdefinable with one of the structures described in Theorem 6; or otherwise Γ has an existential definition in $(V; =)$, which is again one of the five cases from Theorem 6. \square

The endomorphism monoid $\text{End}(G)$ of the random graph has been studied in [13, 14, 20]. By Proposition 12, studying closed transformation monoids containing $\text{Aut}(G)$ is equivalent to studying reducts of G up to existential positive interdefinability. A complete classification of all locally closed transformation monoids that contain all permutations of V , and hence of the reducts of $(V; =)$ up to existential positive interdefinability, has been given in [4]; there is only a countable number of such monoids. The results of the present paper are far from providing a full classification of the locally closed transformation monoids that contain $\text{Aut}(G)$ – this is left for future investigation.

4. RAMSEY-THEORETIC PRELIMINARIES

We recall some Ramsey-type theorems and extend these theorems for our purposes. This will allow us to find patterns in colorings of edges and non-edges of graphs and of graphs equipped with additional structure.

We start by recalling a theorem on ordered structures due to Nešetřil and Rödl [26] of which we will make heavy use. Let $\tau = \tau' \cup \{<\}$ be a relational signature, and let $\mathcal{C}(\tau)$ be the class of all finite τ -structures \mathcal{S} where $<$ denotes a linear order on the domain of \mathcal{S} . For τ -structures \mathcal{A}, \mathcal{B} , let $\binom{\mathcal{A}}{\mathcal{B}}$ be the set of all substructures of \mathcal{A} that are isomorphic to \mathcal{B} (we also refer to members of $\binom{\mathcal{A}}{\mathcal{B}}$ as *copies of \mathcal{B} in \mathcal{A}*). For a finite number $k \geq 1$, a k -coloring of the copies of \mathcal{B} in \mathcal{A} is simply a mapping χ from $\binom{\mathcal{A}}{\mathcal{B}}$ into a set of size k .

Definition 16. For $\mathcal{S}, \mathcal{H}, \mathcal{P} \in \mathcal{C}(\tau)$ and $k \geq 1$, we write $\mathcal{S} \rightarrow (\mathcal{H})_k^{\mathcal{P}}$ iff for every k -coloring χ of the copies of \mathcal{P} in \mathcal{S} there exists a copy \mathcal{H}' of \mathcal{H} in \mathcal{S} such that all copies of \mathcal{P} in \mathcal{H}' have the same color under χ .

Theorem 17 (from [1, 26, 27]). *The class $\mathcal{C}(\tau)$ of all finite relational ordered τ -structures is a Ramsey class, i.e., for all $\mathcal{H}, \mathcal{P} \in \mathcal{C}(\tau)$ and $k \geq 1$ there exists $\mathcal{S} \in \mathcal{C}(\tau)$ such that $\mathcal{S} \rightarrow (\mathcal{H})_k^{\mathcal{P}}$.*

Corollary 18. *For every finite graph \mathcal{H} and for all colorings χ_E and χ_N of the edges and the non-edges of the random graph G , respectively, by finitely many colors, there exists an isomorphic copy of \mathcal{H} in G on which both colorings are constant.*

Proof. Let k be the number of colors used altogether by χ_E and χ_N . Let $<$ be any total order on the domain of \mathcal{H} , and denote the structure obtained from \mathcal{H} by adding the order $<$ to the signature by $\bar{\mathcal{H}}$. Consider the complete graph \mathcal{K}_2 on two vertices, and order its two vertices anyhow to arrive at a structure $\bar{\mathcal{K}}_2$. Then the coloring χ_E of the edges of \mathcal{H} can be viewed as a coloring of the copies of $\bar{\mathcal{K}}_2$ in $\bar{\mathcal{H}}$. Let $\bar{\mathcal{S}}$ with $\bar{\mathcal{S}} \rightarrow (\bar{\mathcal{H}})_{k^2}^{\bar{\mathcal{K}}_2}$ be provided by the preceding theorem, and let \mathcal{S} be $\bar{\mathcal{S}}$ without the order. Then \mathcal{S} is a graph with the property that whenever we color its edges with k colors, then there is a copy of \mathcal{H} in \mathcal{S} all of whose edges have the same color. Now we repeat the argument for the non-edges, starting from \mathcal{S} instead of \mathcal{H} . We then arrive at a graph \mathcal{T} with the property that whenever we color its edges and non-edges by k colors, then there is a copy \mathcal{H}' of \mathcal{H} in \mathcal{T} such that all edges of \mathcal{H}' have the same color, and such that non-edges of \mathcal{H}' have the same color. \mathcal{T} has a copy in G , proving the claim. \square

We will not only need to color edges of graphs, but also of graphs equipped with additional structure.

Definition 19.

- An n -partitioned graph is a structure $\mathcal{U} = (U; F, U_1, \dots, U_n)$, where $(U; F)$ is a graph and each U_i is a subset of U such that the U_i form a partition of U .
- An n -constant graph is a structure $\mathcal{U} = (U; F, u_1, \dots, u_n)$, where $\mathcal{U} = (U; F)$ is a graph, and $u_i \in U$ are distinct.

Observe that n -constant graphs are not relational structures; therefore, in order to apply Theorem 17, we have to make them relational: To every n -constant graph $\mathcal{U} = (U; F, u_1, \dots, u_n)$ we can assign an $n+2^n$ -partitioned graph $\tilde{\mathcal{U}} = (U; F, \{u_1\}, \dots, \{u_n\}, U_1, \dots, U_{2^n})$ in which the u_i belong to singleton sets, and in which for every possible relative position (edge or non-edge) to the u_i we have a set U_j of all elements in $U \setminus \{u_1, \dots, u_n\}$ having this position. (In the language of model theory, every of the $n+2^n$ sets corresponds to a maximal quantifier-free 1-type over the structure \mathcal{U} .) We call the parts U_i the *proper parts* of $\tilde{\mathcal{U}}$.

Definition 20. Let Γ be a structure and $a^1, \dots, a^m \in \Gamma$. We write $\text{tp}_{\text{qf}}(a^1, \dots, a^m)$ for the set of quantifier-free formulas satisfied by the tuple (a^1, \dots, a^m) in Γ , and refer to this set as the *type* of (a^1, \dots, a^m) in Γ .

Definition 21. Let Γ be a structure and let $m \geq 1$. A coloring χ of the m -element subsets of Γ is called *canonical* iff for all tuples (a^1, \dots, a^m) and (b^1, \dots, b^m) enumerating m -element subsets of Γ , if $\text{tp}_{\text{qf}}(a^1, \dots, a^m) = \text{tp}_{\text{qf}}(b^1, \dots, b^m)$, then they induce subsets of equal color under χ .

In this section, we will consider colorings of the two-element subsets of graphs, n -partitioned graphs and n -constant graphs. For disjoint subsets S_1 and S_2 of any such structure, we will say that a coloring is canonical on S_1 iff it satisfies the definition of canonicity for subsets of S_1 ; moreover, we will say that a coloring is canonical between S_1 and S_2 iff it satisfies the definition of canonicity for subsets of $S_1 \cup S_2$ which have precisely one element in S_1 and one element in S_2 .

Lemma 22 (The n -partitioned graph Ramsey lemma). *Let $n, k \geq 1$. For any finite n -partitioned graph $\mathcal{U} = (U; F, U_1, \dots, U_n)$ there exists a finite n -partitioned graph $\mathcal{Q} = (Q; D, Q_1, \dots, Q_n)$ with the property that for all colorings of the two-element subsets of Q with k colors, there exists a copy of \mathcal{U} in \mathcal{Q} on which the coloring is canonical.*

Proof. We show the lemma for $n = 2$; the generalization to larger n is straightforward. For $n = 2$, we apply Theorem 17 six times: Once for the edges in U_1 , once for the edges in U_2 , once for the edges between U_1 and U_2 , and then the same for all three kinds of non-edges.

In general, we would have to apply the theorem 2 ($n + \binom{n}{2}$) times: Once for the edges of each part U_i , once for the edges between any two distinct parts U_i, U_j , and then the same for all non-edges on and between parts.

So assume $n = 2$. We exhibit the idea in detail for the edges between U_1 and U_2 . Let \prec be any total order on U with the property that $u_1 \prec u_2$ for all $u_1 \in U_1, u_2 \in U_2$. Consider the 2-partitioned graph $\mathcal{L}^1 = (\{a, b\}; \{(a, b), (b, a)\}, \{a\}, \{b\})$ and order its vertices by setting $a \prec b$; so \mathcal{L}^1 consists of two adjacent vertices which are ordered somehow, and which lie in different parts. By Theorem 17, there exists an ordered partitioned graph $\mathcal{Q}^1 = (Q^1; D^1, Q_1^1, Q_2^1, \prec)$ such that $\mathcal{Q}^1 \rightarrow (\mathcal{U})_k^{\mathcal{L}^1}$.

Now, if we change the order on \mathcal{Q}^1 in such a way that $r \prec s$ for all $r \in Q_1^1$ and all $s \in Q_2^1$ and such that the order within the parts Q_1^1, Q_2^1 remains unaltered, then the statement $\mathcal{Q}^1 \rightarrow (\mathcal{U})_k^{\mathcal{L}^1}$ still holds: For, given a coloring of the copies of \mathcal{L}^1 with respect to the new ordering, we obtain a coloring of (possibly fewer) copies of \mathcal{L}^1 with respect to the old ordering. There, we obtain a copy \mathcal{U}' of \mathcal{U} such that all copies of \mathcal{L}^1 in \mathcal{U}' have the same color. But in this copy, by the choice of the order on \mathcal{U} , we have that $r \prec s$ for all $r \in U_1'$ and all $s \in U_2'$. Therefore, this copy is also a substructure of \mathcal{Q}^1 with respect to the new ordering.

Since we can change the ordering on \mathcal{Q}^1 in the way described above, the colorings of the copies of \mathcal{L}^1 are just colorings of those pairs $\{r, s\}$, with $r \in Q_1^1$ and $s \in Q_2^1$, which are edges.

Now we repeat the process with the structure $\mathcal{L}^2 = (\{a, b\}; \{(a, b), (b, a)\}, \{a, b\}, \emptyset)$, ordered again by setting $a \prec b$, starting with \mathcal{Q}^1 . We then obtain a structure \mathcal{Q}^2 ; this step takes care of the edges which lie within U_1 . After that we proceed with $\mathcal{L}^3 = (\{a, b\}; \{(a, b), (b, a)\}, \emptyset, \{a, b\})$, thereby taking care of the edges within U_2 . We then apply Theorem 17 three more times with the structures $\mathcal{L}^4 = (\{a, b\}; \emptyset, \{a\}, \{b\})$, $\mathcal{L}^5 = (\{a, b\}; \emptyset, \{a, b\}, \emptyset)$, and $\mathcal{L}^6 = (\{a, b\}; \emptyset, \emptyset, \{a, b\})$, in order to ensure homogeneous non-edges. \square

We now arrive at the goal of this section, namely the following lemma, which we are going to apply to operations on the random graph numerous times in the sections to come.

Lemma 23 (The n -constant graph Ramsey lemma). *Let $n, k \geq 1$. For any finite n -constant graph $\mathcal{U} = (U; F, u_1, \dots, u_n)$ there exists a finite n -constant graph $\mathcal{Q} = (Q; D, q_1, \dots, q_n)$ with the property that for all colorings of the two-element subsets of Q with k colors, there exists a copy of \mathcal{U} in \mathcal{Q} on which the coloring is canonical.*

Proof. Let $\tilde{\mathcal{U}} := (U; F, \{u_1\}, \dots, \{u_n\}, U_1, \dots, U_{2^n})$ be the partitioned graph associated with \mathcal{U} . We would like to use the partitioned graph Ramsey lemma (Lemma 22) in order to obtain \mathcal{Q} ; but we want the singleton sets $\{u_i\}$ of the partition to remain singletons, which is not guaranteed by that lemma.

So consider the 2^n -partitioned graph $\mathcal{R} := (U \setminus \{u_1, \dots, u_n\}; F, U_1, \dots, U_{2^n})$, and apply the partitioned graph Ramsey lemma to this graph to obtain a partitioned graph \mathcal{R}^0 .

Equip \mathcal{R}^0 with any linear order. Now consider the ordered 2^n -partitioned graph \mathcal{L}^1 which has just one vertex, and whose first part contains this single vertex. Apply Theorem 17 in order to obtain an ordered partitioned graph \mathcal{R}^1 such that $\mathcal{R}^1 \rightarrow (\mathcal{R}^0)_{k^n}^{\mathcal{L}^1}$.

Next, consider the ordered 2^n -partitioned graph \mathcal{L}^2 which has just one vertex, and whose second part contains this single vertex. Apply Theorem 17 in order to obtain an ordered partitioned graph \mathcal{R}^2 such that $\mathcal{R}^2 \rightarrow (\mathcal{R}^1)_{k^n}^{\mathcal{L}^2}$.

Repeat this procedure with the ordered 2^n -partitioned graphs $\mathcal{L}^3, \dots, \mathcal{L}^{2^n}$; \mathcal{L}^i has its single vertex in its i -th part. We end up with an ordered partitioned graph \mathcal{R}^{2^n} . We now forget its order and denote the resulting structure by $\mathcal{T} = (T; C, T_1, \dots, T_{2^n})$.

\mathcal{T} has the following property: Whenever we color its vertices with k^n colors, then we find a copy of \mathcal{R}^0 in \mathcal{T} such that the coloring is constant on each part of this copy. Hence, it has the property that if we color its two-element subsets and its vertices with k and k^n colors, respectively, then we find in it a copy of \mathcal{R} on which the first coloring is canonical, and such that the color of the vertices depends only on the part the vertex lies in.

Now consider the structure $\mathcal{S} := (T \cup \{u_1, \dots, u_n\}; B, \{u_1\}, \dots, \{u_n\}, T_1, \dots, T_{2^n})$, where B consists of the edges of \mathcal{T} , plus edges connecting the u_i with the vertices of some parts T_i , depending on whether u_i was in \mathcal{U} connected to the vertices in U_i or not. Clearly, \mathcal{S} is the partitioned graph of the n -constant graph $\mathcal{Q} := (T \cup \{u_1, \dots, u_n\}; B, u_1, \dots, u_n)$. We

claim that \mathcal{Q} has the property we want to prove. Assume that we color the two-element subsets of $T \cup \{u_1, \dots, u_n\}$ with k colors. We must find a copy of \mathcal{U} in \mathcal{Q} on which the coloring is canonical. Divide the coloring into two colorings, namely the coloring restricted to two-element subsets of T , and the coloring of two-element subsets which contain at least one element u_i outside T . The color of the sets $\{u_i, u_j\}$ completely outside T is irrelevant for what we want to prove, so forget about these.

Now the coloring of those sets which have exactly one element outside T can be encoded in a coloring of the vertices of T : Each vertex is given one of k^n colors, depending on the colors of its edges leading to u_1, \dots, u_n . So we have encoded the original coloring into a coloring of two-elements subsets of T and a coloring of the vertices of T . With our observation above, this proves the lemma. \square

5. FINDING STRUCTURE IN MAPPINGS ON THE RANDOM GRAPH

In this section we show how to use the Ramsey-theoretic results from the last section in our context. That is, we will use those results in order to find regular patterns in the behavior of unary functions from V to V .

Definition 24. Let τ be any signature and let $\mathcal{C}(\tau)$ be a class of finite τ -structures. We say that a property P holds for arbitrarily large elements of $\mathcal{C}(\tau)$ iff for any $\mathcal{F} \in \mathcal{C}(\tau)$ there exists $\mathcal{H} \in \mathcal{C}(\tau)$ such that \mathcal{F} embeds into \mathcal{H} and $P(\mathcal{H})$ holds. We say that P holds for all sufficiently large elements of $\mathcal{C}(\tau)$ iff there exists $\mathcal{F} \in \mathcal{C}(\tau)$ such that P holds for \mathcal{H} whenever \mathcal{F} embeds into \mathcal{H} .

Our classes $\mathcal{C}(\tau)$ will be closed under induced substructures; moreover, our properties P will be *hereditary*, i.e., if $P(\mathcal{H})$ holds, then P also holds for all substructures of \mathcal{H} . The definition then says that P holds for arbitrarily large elements of $\mathcal{C}(\tau)$ iff for any $\mathcal{F} \in \mathcal{C}(\tau)$ there is $\mathcal{F}' \in \mathcal{C}(\tau)$ isomorphic to \mathcal{F} such that $P(\mathcal{F}')$ holds.

In our situation, $\mathcal{C}(\tau)$ will also have the *joint embedding property (JEP)*, i.e., for any two structures in $\mathcal{C}(\tau)$ there exists a structure in $\mathcal{C}(\tau)$ that embeds both structures. We then have that if P holds for all sufficiently large elements of $\mathcal{C}(\tau)$, then it holds for arbitrarily large elements of $\mathcal{C}(\tau)$. Observe also that under (JEP), if arbitrarily large structures in $\mathcal{C}(\tau)$ have one of finitely many properties, then one of those properties holds for arbitrarily large elements of $\mathcal{C}(\tau)$.

Definition 25. Let $e, f: V \rightarrow V$. We say that e behaves as f on $F \subseteq V$ iff there is an automorphism α of G such that $f(x) = \alpha(e(x))$ for all $x \in F$. We say that e interpolates f modulo automorphisms iff for every finite $F \subseteq V$ there is an automorphism β of G such that $e(\beta(x))$ behaves as f on F ; so this is the case iff there exist automorphisms α, β such that $\alpha(e(\beta(x))) = f(x)$ for all $x \in F$.

Note that if e interpolates f modulo automorphisms, then it also generates f .

Definition 26. Let \mathcal{U} be a graph, an n -partitioned graph, or an n -constant graph. Any function $f: \mathcal{U} \rightarrow \mathcal{U}$ induces a coloring of the two-element subsets of \mathcal{U} as follows: the color of a set $\{x, y\}$ is the type of $(f(x), f(y))$ with respect to the graph relation of \mathcal{U} . We say that f is *canonical* iff the coloring it induces is canonical.

Proposition 27. Let $e: V \rightarrow V$ be a mapping on the random graph G . Then e is canonical on arbitrarily large subgraphs of G , and interpolates either the identity, e_E , e_N , a constant function, or – modulo automorphisms.

Proof. We show that arbitrarily large finite subgraphs of G have the property that e behaves on them like one of the operations of the proposition. Since there are finitely many operations to choose from, e then behaves like one fixed operation p from the list on arbitrarily large finite subgraphs of G . By the homogeneity of G , we can freely move finite graphs around by automorphisms, proving that e interpolates p .

So let \mathcal{F} be any finite graph; we have to find a copy \mathcal{F}' of \mathcal{F} in G such that e behaves like one of the mentioned operations on this copy.

We color all pairs $\{x, y\}$ of distinct vertices of G

- by 1 if $e(x) = e(y)$,
- by 2 if $(e(x), e(y)) \in E$,
- by 3 if $(e(x), e(y)) \in N$.

By Corollary 18 there exists a copy \mathcal{F}' of \mathcal{F} in G such that all edges and all non-edges of \mathcal{F}' have the same color χ_E and χ_N , respectively. If $(\chi_E, \chi_N) = (1, 1)$, then e behaves like the constant function on \mathcal{F}' . If $(\chi_E, \chi_N) = (2, 3)$, then it behaves like the identity, and if $(\chi_E, \chi_N) = (3, 2)$, then e behaves like $-$. If $(\chi_E, \chi_N) = (2, 2)$ or $(\chi_E, \chi_N) = (3, 3)$, then e behaves like e_E or e_N , respectively. Finally, it is easy to see that $(\chi_E, \chi_N) = (1, q)$ or $(\chi_E, \chi_N) = (q, 1)$, where $q \in \{2, 3\}$, is impossible if \mathcal{F} contains the two three-element graphs with one and two edges, respectively. \square

Definition 28. Let τ be a signature and let T be a theory in this language. We call a τ -structure \aleph_0 -universal for T iff it satisfies T and embeds all finite models of T .

Lemma 29 (The n -partite graph interpolation lemma). *Let $\mathcal{U} = (U; C, U_1, \dots, U_n)$ be an \aleph_0 -universal partitioned graph, and let $f: U \rightarrow U$. Then every finite partitioned graph has a copy in \mathcal{U} on which f is canonical.*

Proof. This is immediate from the n -partitioned graph Ramsey lemma (Lemma 22): Just like in the proof of Proposition 27, we color the edges and non-edges of \mathcal{U} according to what f does to them. \square

Lemma 30 (The n -constant graph interpolation lemma). *Let $\mathcal{U} = (U; C, u_1, \dots, u_n)$ be an \aleph_0 -universal n -constant graph, and let $f: U \rightarrow U$. Then every finite n -constant graph has a copy in \mathcal{U} on which f is canonical.*

Proof. This is immediate from the n -constant graph Ramsey lemma (Lemma 23). \square

6. UNARY FUNCTIONS

We now have the tools to settle the unary case: In this section, we will prove Theorem 3 which characterizes the unary minimal functions, Theorem 6 which lists the five closed supergroups of $\text{Aut}(G)$, and Theorem 7 which states that any closed monoid containing $\text{Aut}(G)$ either is generated by the group of its permutations, or contains e_E , e_N , or a constant function. We start by applying Lemma 30 to prove

Lemma 31. *Let $e: V \rightarrow V$ be so that it preserves N but not E . Then e generates e_N . Dually, if $e: V \rightarrow V$ is so that it preserves E but not N , then e generates e_E .*

Proof. We prove that for every finite subset F of V , e generates an operation which behaves like e_N on F . We first claim that there are adjacent vertices $a, b \in V$ such that $(e(a), e(b)) \in N$. Since e does not preserve E , there exist u, v with $(u, v) \in E$ such that $(e(u), e(v)) \notin E$. If

$(e(u), e(v)) \in N$, then we are done. If $e(u) = e(v)$, then choose w such that $(w, u) \in E$ and $(w, v) \in N$. We have $(e(w), e(u)) = (e(w), e(v)) \in N$, so u, w prove the claim.

Now, $\mathcal{U} := (V; E, a, b)$ is an \aleph_0 -universal 2-constant graph. Therefore, by Lemma 30, e is canonical on arbitrarily large substructures of \mathcal{U} . Since e preserves N , it is easy to see that if e is canonical on a 2-constant graph which is sufficiently large, then e must be injective; for example, if e is canonical on a graph which contains the three-element graph with two edges, then e cannot collapse any edges of that graph. Hence, e is canonical and injective on arbitrarily large 2-constant subgraphs of \mathcal{U} . Since e preserves N , we have that for arbitrarily large substructures of \mathcal{U} , it behaves like the identity or like e_N on and between the parts of these structures; in particular, it does not turn any non-edges into edges on and between the parts of these structures. Hence, for any finite 2-constant subgraph of \mathcal{U} , by applying an automorphism of G and then e , we can delete the edge between the two constants without turning any non-edge of that 2-constant graph into an edge. But that means that starting from any finite graph, we can delete all edges by repeating this process, choosing any edge we want to get rid of in each step. This proves the lemma. \square

Lemma 32. *If $e: V \rightarrow V$ preserves neither E nor N and is not injective, then e generates a constant operation.*

Proof. Since e is not injective, it collapses without loss of generality an edge (otherwise dualize). Since e violates N , it either collapses a non-edge or sends some non-edge to an edge, which, with the help of an appropriate automorphism, can be collapsed by another application of e . Thus e generates operations g, h which collapse an edge and a non-edge, respectively. Having this, one sees that e generates a constant function on each finite subset F of V , by shifting F around with automorphisms and applying g and h to collapse all points in F to a single vertex. \square

The following proposition already identifies the five minimal functions of Theorem 3.

Proposition 33. *Let Γ be a reduct of G . Then one of the following cases applies.*

- (1) Γ has a constant endomorphism.
- (2) Γ has e_E as an endomorphism.
- (3) Γ has e_N as an endomorphism.
- (4) Γ has $-$ as an automorphism.
- (5) Γ has sw as an automorphism.
- (6) $\text{Aut}(G)$ is dense in $\text{End}(\Gamma)$.

Proof. If Γ has an endomorphism e which preserves E but not N or N but not E , then we can refer to Lemma 31. If all of its endomorphisms preserve both N and E , then $\text{Aut}(G)$ is dense in $\text{End}(\Gamma)$. We thus assume henceforth that Γ has an endomorphism e which violates both E and N .

If e is not injective, then it generates a constant operation, by Lemma 32. So suppose that e is injective. Fix $(x, y) \in E$ such that $(e(x), e(y)) \in N$.

By Proposition 27, e is canonical on arbitrarily large finite subgraphs of G . If e interpolates $-$, e_E , or e_N modulo automorphisms, then we are done. So assume this is not the case, i.e., there is a finite graph \mathcal{F}_0 with the property that on all copies of \mathcal{F}_0 in G , e does not behave like any of these operations. Observe that e then behaves like the identity on arbitrarily large subgraphs of G . Moreover, this assumption implies that if a finite subgraph \mathcal{F} of G is sufficiently large (i.e., if it embeds \mathcal{F}_0), and e is canonical on \mathcal{F} , then e behaves like the identity on \mathcal{F} .

We now make a series of observations which rule out bad behavior of e between subsets of the random graph, and which follow from our assumptions of the preceding paragraph; the easily verifiable details are left to the reader.

- If e behaves like $-$ between the parts of arbitrarily large finite 2-partitioned subgraphs of G , then it generates sw .
- If e behaves like e_N between the parts of arbitrarily large finite 2-partitioned subgraphs of G , then it generates e_N .
- If e behaves like e_E between the parts of arbitrarily large finite 2-partitioned subgraphs of G , then it generates e_E .

We assume therefore that for sufficiently large finite 2-partitioned subgraphs of G , if e is canonical on such a graph, then e behaves like the identity on and between the parts.

Now observe that $\mathcal{Q} := (V; E, x, y)$ is an \aleph_0 -universal 2-constant graph. Let $\mathcal{F} = (F; D, f_1, f_2)$ be any finite 2-constant graph. By the n -constant interpolation lemma (Lemma 30), there is a copy \mathcal{F}' of \mathcal{F} in \mathcal{Q} on which e is canonical. By our assumption above, if only \mathcal{F} is sufficiently large, then being canonical on a proper part F'_i of the 6-partitioned graph $\tilde{\mathcal{F}}' = (F'; E, \{x\}, \{y\}, F'_1, \dots, F'_4)$ corresponding to \mathcal{F}' means behaving like the identity thereon, and being canonical between proper parts means behaving like the identity between these parts. Therefore, all 2-constant graphs \mathcal{F} have a copy $\mathcal{F}' = (F'; E, x, y)$ in \mathcal{Q} such that e behaves like the identity on and between all of the parts F'_i, F'_j of the corresponding partitioned graph $\tilde{\mathcal{F}}' = (F'; E, \{x\}, \{y\}, F'_1, \dots, F'_4)$.

Of a two-constant graph \mathcal{F} , consider the reduct $\mathcal{H} = (F; D, f_1)$. This reduct has a copy \mathcal{H}' in $\mathcal{Q}^x = (V; E, x)$ on which e is canonical. The corresponding partitioned graph has two parts H'_1, H'_2 , and x is adjacent to, say, all vertices in H'_1 and to none in H'_2 . Since e is canonical on \mathcal{H}' , either it preserves all edges between x and H'_1 , or it turns all these edges into non-edges. Similarly with the non-edges between x and H'_2 . If all edges are deleted and all non-edges kept for arbitrarily large \mathcal{H} , then e generates e_N . If all edges are deleted and all non-edges edged for arbitrarily large \mathcal{H} , then e interpolates sw modulo automorphisms. If all edges are kept and all non-edges edged for arbitrarily large \mathcal{H} , then e generates e_E . So we assume that if only \mathcal{H} is sufficiently large, then all edges and non-edges are kept by e on those copies of \mathcal{H} on which e is canonical.

We use the same argument with the reduct $(F; D, f_2)$ and $\mathcal{Q}^y = (V; E, y)$, and arrive at the conclusion that if the two-constant graph \mathcal{F} is sufficiently large, then on every copy of \mathcal{F} in \mathcal{Q} on which e is canonical, the edges and non-edges leading from x and y to the other vertices of the copy are kept.

Combining this with what we have established before, we conclude that if only \mathcal{F} is sufficiently large, and \mathcal{F}' is a copy of \mathcal{F} in \mathcal{Q} on which e is canonical, then e behaves like the identity on \mathcal{F}' except between x and y , where it deletes the edge. Hence, for any finite \mathcal{F} we can find a copy in \mathcal{Q} on which e behaves that way. But this implies that starting from any finite graph $\mathcal{S} := (F; D)$, we can pick any edge in \mathcal{S} , say between vertices f_1, f_2 , and then find a copy of $\mathcal{F} := (F; D, f_1, f_2)$ in \mathcal{Q} such that e deletes exactly that edge from the copy without changing the rest. Hence, by shifting finite graphs around with automorphisms, we can delete a single edge from an arbitrary finite subgraph of G without changing the rest of the graph. Applying this successively, we can remove all edges from arbitrary finite graphs, proving that e generates e_N . \square

Now Theorem 3 follows: let f be a minimal function and Γ the reduct whose endomorphism monoid is generated by f . We apply Proposition 33 to Γ . Observe that Case (6) of that proposition cannot hold for Γ , since its endomorphism f is non-trivial, and hence not generated $\text{Aut}(G)$. Thus, Γ contains one of the functions of the other cases, meaning that f is equivalent to one of those functions. This finishes the proof.

Proposition 34. *Let Γ be a reduct of G , and suppose Γ is preserved by $-$, but not by e_N, e_E , or a constant operation. Then the endomorphisms of Γ are generated by $\{-\} \cup \text{Aut}(G)$, or Γ is preserved by sw .*

Proof. Suppose the endomorphisms of Γ are not generated by $\{-\} \cup \text{Aut}(G)$. Then, by Proposition 10, there is a relation R invariant under $\{-\} \cup \text{Aut}(G)$ and an endomorphism e of Γ which violates R ; that is, there exists a tuple $a := (a_1, \dots, a_n) \in R$ such that $e(a) = (e(a_1), \dots, e(a_n)) \notin R$.

Since R is definable in G , e violates either an edge or a non-edge. Hence, as in the proof of Proposition 33, the assumption that e does not generate e_N, e_E , or a constant operation implies that e is injective.

Let $\mathcal{F} = (F; D, f_1, \dots, f_n)$ be any finite n -constant graph. By the n -constant interpolation lemma (Lemma 30), there is a copy \mathcal{F}' of \mathcal{F} in the \aleph_0 -universal n -constant graph $\mathcal{Q} := (V; E, a_1, \dots, a_n)$ such that e is canonical on this copy.

We now make a series of observations on the behavior of e on and between subsets of V where it is canonical.

- Since by assumption, e does not interpolate e_E, e_N , or a constant operation modulo automorphisms, it behaves like $-$ or the identity on sufficiently large finite subgraphs of G where it is canonical.
- Suppose that for arbitrarily large finite 2-partitioned subgraphs of G , e behaves like the identity on the parts and like $-$ between the parts. Then e generates sw .
- Suppose that for arbitrarily large finite 2-partitioned subgraphs of G , e behaves like the identity on the parts and like e_N (like e_E) between the parts. Then e generates e_N (e_E).
- Suppose that for arbitrarily large finite 2-partitioned subgraphs of G , e behaves like $-$ on the parts and like the identity / e_N / e_E between the parts. Then e and $-$ together generate sw / e_E / e_N . This is because we can apply the preceding two observations to $-e$.
- Suppose that for arbitrarily large finite 2-partitioned subgraphs of G on which e is canonical, e behaves like $-$ on one part and like the identity on the other part. Then e and $-$ together generate e_N .

To see the last assertion for the case where e behaves like the identity between the parts, select an edge within one of the parts that is mapped to a non-edge. For arbitrary finite $A \subseteq V$ we can now use the operation e to get rid of one edge in the graph induced by A in G and preserve all other edges, and so eventually generate an operation that behaves like e_N on A . For the case where e behaves like $-$ between the parts, we can apply the same argument to $-e$. If e behaves like e_N between the parts, then we can all the more delete edges. If it behaves like e_E between the parts, then $-e$ behaves like e_N and we are back in the preceding case.

Summarizing our observations, we can assume that for an arbitrary finite n -constant graph \mathcal{F} there is a copy of \mathcal{F} in \mathcal{Q} such that e behaves like the identity on and between all proper

parts F'_i, F'_j of the corresponding partitioned graph, or like $-$ on and between all of its parts. If only the second case holds for arbitrarily large n -constant graphs \mathcal{F} , then we simply proceed our argument with $-e$ instead of e . We can do that since also $-e(a) \notin R$. Otherwise, picking an automorphism α of G such that $\alpha(-(-x)) = x$ for all $x \in V$, we would have $\alpha(-(-e(a))) = e(a) \in R$, contrary to our choice of a . Thus we assume that for arbitrary finite n -constant graphs \mathcal{F} there is a copy of \mathcal{F} in \mathcal{Q} such that e behaves like the identity on and between all proper parts of that copy.

As in the proof of Proposition 33, we may assume that if a copy $\mathcal{F}' = (F'; E, a_1, \dots, a_n)$ of \mathcal{F} in \mathcal{Q} is large enough and e is canonical on \mathcal{F}' and behaves like the identity on and between all proper parts F'_i, F'_j of the corresponding n -partitioned graph $\tilde{\mathcal{F}}'$, then it leaves the edges and non-edges between the a_i and the vertices in $F' \setminus \{a_1, \dots, a_n\}$ unaltered. It follows that for arbitrary finite n -constant graphs \mathcal{F} there is a copy of \mathcal{F} in \mathcal{Q} such that the only edges or non-edges changed by e on this copy are those between the a_i .

Finally, note that since R is definable in the random graph and $e(a) \notin R$, e destroys at least one edge or one non-edge on $\{a_1, \dots, a_n\}$. Without loss of generality, say that a_1, a_2 are adjacent but their values under e are not. We have shown that for arbitrarily large 2-constant graphs \mathcal{H} , there is a copy of \mathcal{H} in $(V; E, a_1, a_2)$ such that e behaves like the identity on this copy, except for the edge between a_1 and a_2 , which is destroyed. This clearly implies that e generates e_N . \square

Proposition 35. *Let Γ be a reduct of G , and suppose Γ is preserved by sw , but not by e_N, e_E , or a constant operation. Then the endomorphisms of Γ are generated by $\{\text{sw}\} \cup \text{Aut}(G)$, or Γ is preserved by $-$.*

Proof. The proof is very similar to the proof of the preceding proposition. This time we know that unless the endomorphisms are generated by $\{\text{sw}\} \cup \text{Aut}(G)$, there exists an endomorphism e that violates a relation R which is preserved by $\{\text{sw}\} \cup \text{Aut}(G)$. Fix a tuple a as before.

As in the preceding proof, we may assume that e is injective. If e interpolates $-$ modulo automorphisms, we are done. Suppose therefore that if e is canonical on a finite partitioned graph which is sufficiently large, then it must behave like the identity on its parts.

If e behaves like e_N (e_E) between the parts of arbitrarily large finite 2-partitioned subgraphs of G , then it generates e_N (e_E). Thus we may assume that it behaves like the identity or $-$ between such parts.

Suppose that for arbitrarily large finite 3-partitioned subgraphs $\mathcal{F} = (F; E, F_1, F_2, F_3)$ of G on which e is canonical, e behaves like the $-$ between exactly two of the parts, say between F_1, F_2 , and like the identity between F_2, F_3 and F_1, F_3 . Then e is easily seen to generate both e_N and e_E . Indeed, if we want to delete¹ any edge from a finite graph, then we can view the vertices of the edge as two parts of a 3-partitioned graph, where the third part contains all the other vertices. If e behaves like $-$ between the two vertices whose edge we want to delete, and like the identity on and between the other parts, what happens is exactly that the edge is deleted.

If for arbitrarily large finite 3-partitioned subgraphs \mathcal{F} of G on which e is canonical, e behaves like $-$ between, say, F_1, F_2 and F_1, F_3 , and like the identity between F_2, F_3 , then by applying a suitable switch operation i_A to e we are back in the preceding case. Note here that there is an automorphism α of G such that $i_A(\alpha(i_A(x))) = x$ for all $x \in V$. Therefore, $i_A(e(a)) \notin R$; for otherwise, we would have $i_A(\alpha(i_A(e(a)))) = e(a) \in R$, a contradiction.

¹For the purposes of the proof, we identify ourselves with the endomorphism monoid.

The latter argument works also if e behaves like $-$ between all three parts. Summarizing, we may assume that if e is canonical on a finite n -partitioned graph which is sufficiently large, where $n \geq 3$, then it behaves like the identity on and between all of the parts.

As for n -constant graphs on which e is canonical, e might flip edges and non-edges between some parts and the constants. However, this situation can easily be repaired by a single application of sw .

Finally, observe that at least one edge or one non-edge on a_1, \dots, a_n is destroyed, and that we therefore can generate either e_N or e_E . \square

Proposition 36. *Let Γ be a reduct of G , and suppose Γ is preserved by sw and by $-$, but not by e_N, e_E , or a constant operation. Then the endomorphisms of Γ are generated by $\{-, \text{sw}\} \cup \text{Aut}(G)$, or Γ is preserved by all permutations.*

Proof. The argument goes as in the preceding two propositions; we leave the details to the reader. \square

Theorem 7 now is a direct consequence of Proposition 33, and Propositions 34, 35, 36: If a reduct Γ of G does not have e_E, e_N , or a constant operation as an endomorphism, and if its endomorphisms are not generated by $\text{Aut}(G)$, then Proposition 33 implies that it has either $-$ or sw as an endomorphism. Since $\text{Aut}(\Gamma)$ contains $\text{Aut}(G)$, once Γ has $-$ or sw as an endomorphism, it also has its inverse as an endomorphism; thus it has $-$ or sw as an automorphism. But then by the preceding three propositions, either $\text{End}(\Gamma)$ is generated by $\text{Aut}(\Gamma)$, or Γ is preserved by all permutations. The latter case, however, is impossible, as this would imply that e_E and e_N are among its endomorphisms, which we excluded already.

Observe also how Thomas' classification of closed permutation groups containing $\text{Aut}(G)$ (Theorem 6) follows from our results: If a closed group properly contains $\text{Aut}(G)$, then it contains $-$ or sw , by Proposition 33. If it contains $-$ but is not generated by $-$, then it contains sw by Proposition 34. Similarly, if it contains sw but is not generated by sw , then it contains $-$ by Proposition 35. If it contains both $-$ and sw , but is not generated by these operations, then it must already contain all permutations (Proposition 36).

7. PRODUCING BINARY INJECTIONS

Having found the minimal unary operations, we now turn to operations of higher arity. The goal of this section is proving Theorem 38, which (together with Lemma 39) implies that all minimal functions are at most binary.

Definition 37. We say that an operation $f: V^k \rightarrow V$ is *essentially unary* iff there exists a unary function $g: V \rightarrow V$ and $1 \leq i \leq k$ such that $f(x_1, \dots, x_k) = g(x_i)$ for all $x_1, \dots, x_k \in V$. If f is not essentially unary, we call it *essential*.

Theorem 38. *Let f be an essential operation on the random graph that preserves E and N . Then f generates a binary injection.*

Lemma 39. *Minimal essential operations on the random graph must preserve E and N .*

Proof. Suppose an essential function $f: V^n \rightarrow V$ does not preserve E (the argument for N is dual). Then there exist tuples $(x_1, \dots, x_n), (y_1, \dots, y_n)$ such that $(x_i, y_i) \in E$ for all $1 \leq i \leq n$ and such that $f(x_1, \dots, x_n)$ and $f(y_1, \dots, y_n)$ are not connected by an edge. Fix $(u, v) \in E$, and choose automorphisms α_i such that $\alpha_i(u) = x_i$ and $\alpha_i(v) = y_i$, for all $1 \leq i \leq n$. The function $g(x) := f(\alpha_1(x), \dots, \alpha_n(x))$ is unary and violates E ; hence it is non-trivial. But being unary, it cannot generate the essential function f , proving that f is not minimal. \square

The following lemma allows us to work with binary operations; its proof is very similar to the proof of a corresponding lemma in [6].

Lemma 40. *Let $f: V^k \rightarrow V$ be an essential operation. Then f generates a binary essential operation.*

Proof. Assume without loss of generality that f depends all of its arguments and is at least ternary. In particular, there are a_1, \dots, a_k and a'_1 such that $f(a_1, \dots, a_k) \neq f(a'_1, a_2, \dots, a_k)$. We distinguish two cases.

Case 1. There are b_1, \dots, b_k such that $(b_i, a_i) \in E$ for $2 \leq i \leq k$ and $f(b_1, a_2, \dots, a_k) \neq f(b_1, \dots, b_k)$. By the homogeneity of G we can find automorphisms $\alpha_3, \dots, \alpha_k$ such that $\alpha_i(a_2) = a_i$ and $\alpha_i(b_2) = b_i$. Using these automorphisms we define

$$g(x, y) := f(x, y, \alpha_3(y), \dots, \alpha_k(y)),$$

which clearly depends on both arguments.

Case 2. For all b_1, \dots, b_k , if $(a_i, b_i) \in E$ for $2 \leq i \leq k$, then $f(b_1, a_2, \dots, a_k) = f(b_1, b_2, \dots, b_k)$. Since f depends on its second coordinate, there are c_1, \dots, c_k and c'_2 such that

$$f(c_1, c_2, c_3, \dots, c_k) \neq f(c_1, c'_2, c_3, \dots, c_k).$$

Then $f(c_1, a_2, \dots, a_k)$ can be equal to either $f(c_1, c_2, c_3, \dots, c_k)$, or to $f(c_1, c'_2, c_3, \dots, c_k)$, but not to both. We assume without loss of generality that $f(c_1, a_2, \dots, a_k) \neq f(c_1, c_2, c_3, \dots, c_k)$. From the extension property of the random graph we see that we can choose d_2, \dots, d_k such that $(d_i, a_i) \in E$ and $(d_i, c_i) \in E$ for $2 \leq i \leq k$. Since G is homogeneous there are automorphisms $\alpha_3, \dots, \alpha_k$ of G such that $\alpha_i(c_2) = c_i$ and $\alpha_i(d_2) = d_i$. We claim that the operation g defined by

$$g(x, y) := f(x, y, \alpha_3(y), \dots, \alpha_k(y))$$

depends on both arguments. Indeed, we know that $g(a_1, d_2) = f(a_1, d_2, \dots, d_k) = f(a_1, \dots, a_k)$, and that $f(a'_1, d_2) = f(a'_1, d_2, \dots, d_k) = f(a'_1, a_2, \dots, a_k)$. By the choice of the values a_1, \dots, a_k and a'_1 these two values are distinct, and we have that g depends on the first argument. For the second argument, note that $g(c_1, d_2) = f(c_1, d_2, \dots, d_k) = f(c_1, a_2, \dots, a_k)$ and that $g(c_1, c_2) = f(c_1, c_2, \dots, c_k)$. Because $f(c_1, a_2, \dots, a_k)$ and $f(c_1, c_2, \dots, c_k)$ are distinct, we have that g also depends on the second argument. \square

We are left with the task of producing a binary injection from a binary essential function preserving E and N . This will be prepared in the general Lemma 42.

Definition 41. A relation $R \subseteq X^k$ is called *intersection-closed* iff for all $(u_1, \dots, u_k), (v_1, \dots, v_k) \in R$ there is a tuple $(w_1, \dots, w_k) \in R$ such that for all $1 \leq i, j \leq k$ we have $w_i \neq w_j$ whenever $u_i \neq u_j$ or $v_i \neq v_j$.

Lemma 42. *Let Γ be a countable ω -categorical structure where \neq is primitive positive definable. Then the following are equivalent.*

- (1) *If ϕ is a primitive positive formula such that both $\phi \wedge x \neq y$ and $\phi \wedge u \neq v$ are satisfiable over Γ , then $\phi \wedge x \neq y \wedge u \neq v$ is satisfiable over Γ as well.*
- (2) *Every finite induced substructure of Γ^2 admits an injective homomorphism into Γ .*
- (3) *Γ is preserved by a binary injective operation.*
- (4) *All primitive positive definable relations in Γ are intersection-closed.*

We would like to remark that the first item in Lemma 42 is inspired from joint work of the alphabetically first author with Peter Jonsson and Timo von Oertzen in [5].

Proof. Throughout the proof, let e_1, e_2, \dots be an enumeration of the domain D of Γ . If f is a binary injective polymorphism of Γ , then clearly every relation in Γ is intersection-closed, so (3) implies (4). The implication from (4) to (1) is straightforward as well.

We now show the implication from (1) to (2). Let S be a finite induced substructure of Γ^2 . Without loss of generality we can assume that S is induced in Γ^2 by a set of the form $\{e_1, \dots, e_n\}^2$, for sufficiently large n . Consider the formula ϕ whose variables x_1, \dots, x_{n^2} are the elements of S ,

$$x_1 := (e_1, e_1), \dots, x_n := (e_1, e_n), \dots, x_{n^2-n+1} := (e_n, e_1), \dots, x_{n^2} := (e_n, e_n),$$

and which is the conjunction over all literals $R((e_{i_1}, e_{j_1}), \dots, (e_{i_k}, e_{j_k}))$ such that $R(e_{i_1}, \dots, e_{i_k})$ and $R(e_{j_1}, \dots, e_{j_k})$ hold in Γ . So ϕ states precisely which relations hold in S .

Using induction over the number of inequalities, we will now show that for any conjunction $\sigma := \bigwedge_{1 \leq k \leq m} x_{i_k} \neq x_{j_k}$ with the property that $i_k \neq j_k$ for all $1 \leq k \leq m$, the formula $\phi \wedge \sigma$ is satisfiable over Γ . This implies that there exists an n^2 -tuple t in Γ with pairwise distinct entries which satisfies ϕ ; the assignment that sends every $x_i \in S$ to t_i is an injective homomorphism from S into Γ .

For the induction beginning, let $x_i \neq x_j$ be any inequality. Let r, s be the n^2 -tuples defined as follows.

$$\begin{aligned} r &:= (e_1, \dots, e_1, e_2, \dots, e_2, \dots, e_n, \dots, e_n) \\ s &:= (e_1, e_2, \dots, e_n, e_1, e_2, \dots, e_n, \dots, e_1, e_2, \dots, e_n). \end{aligned}$$

These two tuples satisfy ϕ , because the projections to the first and second coordinate, respectively, are homomorphisms from S to Γ . Now either r or s satisfies $x_i \neq x_j$, proving that $\phi \wedge x_i \neq x_j$ is satisfiable in Γ .

In the induction step, let a conjunction $\sigma := \bigwedge_{1 \leq k \leq m} x_{i_k} \neq x_{j_k}$ be given, where $m \geq 2$. Set $\sigma' := \bigwedge_{3 \leq k \leq m} x_{i_k} \neq x_{j_k}$, and $\phi' := \phi \wedge \sigma'$. Observe that ϕ' has a primitive positive definition in Γ , as ϕ and \neq have such definitions. By induction hypothesis, both $\phi' \wedge x_{i_1} \neq x_{j_1}$ and $\phi' \wedge x_{i_2} \neq x_{j_2}$ are satisfiable in Γ . But then $\phi' \wedge x_{i_1} \neq x_{j_1} \wedge x_{i_2} \neq x_{j_2} = \phi \wedge \sigma$ is satisfiable over Γ as well by (1), concluding the proof.

The implication from (2) to (3) is by a standard application of König's lemma. This is because the fact that Γ is ω -categorical implies that for every $n \geq 1$, there are only finitely many "behaviors" of functions from $\{e_1, \dots, e_n\}^2$ to Γ , by the theorem of Ryll-Nardzewski.

We give the standard argument for completeness. We say that two homomorphisms f_1, f_2 from the structure induced by a set $\{e_1, \dots, e_l\}^2$ in Γ to Γ are *equivalent* if there is an automorphism h of Γ such that $h(f_1(x, y)) = f_2(x, y)$ for all $x, y \in \{e_1, \dots, e_l\}$. Consider the infinite tree T whose vertices are the equivalence classes of injective homomorphisms from structures induced by a set of the form $\{e_1, \dots, e_l\}^2$ to Γ . There is an arc from one equivalence class of injective homomorphisms to another in T iff there are representatives f_1 and f_2 of the two classes such that the domain of f_1 is $\{e_1, \dots, e_l\}^2$, and the domain of f_2 is $\{e_1, \dots, e_l, e_{l+1}\}^2$, and f_2 is an extension of f_1 .

The theorem of Ryll-Nardzewski implies that every node in T has a finite number of outgoing arcs, since there are only finitely many inequivalent homomorphisms from a set $\{e_1, \dots, e_l\}^2$ to Γ . Since T is infinite by (2), König's lemma asserts the existence of an infinite branch B in T . This infinite branch gives rise to an injective binary polymorphism f of Γ , which is defined inductively as follows. The restriction of f to $\{e_1, \dots, e_n\}$ will be an element from the n -th node of B . To start the induction, pick any function f_1 from the first node of B ; f_1 has domain $\{e_1\}^2$. Set f to equal f_1 on $\{e_1\}^2$. Suppose f is already defined on

e_1, \dots, e_n , for $n \geq 1$. By definition of T , we find representatives f_n and f_{n+1} of the n -th and the $n+1$ -st element of B such that f_n is a restriction of f_{n+1} . The inductive assumption gives us an automorphism h of Γ such that $h(f_n(x, y)) = f(x, y)$ for all $x, y \in \{e_1, \dots, e_n\}$. We set $f(x, y)$ to be $h(f_{n+1}(x, y))$, for all $x, y \in \{e_1, \dots, e_{n+1}\}$. The restriction of f to e_1, \dots, e_{n+1} will therefore be a member of the $n+1$ -st node of B . The operation f defined in this way is indeed an injective homomorphism from Γ^2 to Γ , and we are done. \square

We are now ready to end this section and provide a proof of Theorem 38.

Proof of Theorem 38. Let an essential operation $f: V^k \rightarrow V$ preserving E and N be given. By Lemma 40, f generates a binary essential function; clearly, this function still preserves E and N , so that we may henceforth assume that f is itself binary.

Consider the structure Δ whose relations are the relations that are first-order definable in G and preserved by f . In order to prove that f generates a binary injection, we will refer to Proposition 10 and prove that there is a binary injection preserving Δ .

By its definition, Δ has E and N amongst its relations. We claim that \neq is also among the relations of Δ : This is because $x \neq y$ iff $\exists z (E(x, z) \wedge N(y, z))$, so \neq has a primitive positive definition from E and N , and hence from Δ . Hence, we may apply Lemma 42 to Δ , and in order to show that Δ is preserved by a binary injection, it suffices to show that if ϕ is a primitive positive formula over Δ such that both $\phi \wedge x \neq y$ and $\phi \wedge s \neq t$ are satisfiable over Δ , then $\phi \wedge x \neq y \wedge s \neq t$ is satisfiable over Δ as well.

To this end, let ϕ be a primitive positive formula over the signature of Δ such that

- there is a tuple t_1 that satisfies $\phi \wedge x \neq y$
- there is a tuple t_2 that satisfies $\phi \wedge s \neq t$.

Let a_1, a_2, a_3, a_4 and b_1, b_2, b_3, b_4 be the values for x, y, s, t in t_1 and t_2 , respectively. We have $a_1 \neq a_2$ and $b_3 \neq b_4$. We want to show that $\phi \wedge x \neq y \wedge s \neq t$ is satisfiable over Δ . Thus, if $a_3 \neq a_4$ or $b_1 \neq b_2$, there is nothing to show, and so we assume that $a_3 = a_4$ and $b_1 = b_2$.

We claim that there are automorphisms α, β of G such that in the tuple $t_3 := f(\alpha(t_1), \beta(t_2))$ the value of x is different from the value of y , and the value of s is different from the value of t . Then, since f preserves Δ , the tuple t_3 shows that $\phi \wedge x \neq y \wedge s \neq t$ is satisfiable over Δ , and concludes the proof.

To prove the claim, we will find tuples $c := (c_1, c_2, c_3, c_4)$ and $d := (d_1, d_2, d_3, d_4)$ of the same type as (a_1, a_2, a_3, a_4) and (b_1, b_2, b_3, b_4) , respectively, such that the tuple $e := f(c, d)$ satisfies $e_1 \neq e_2$ and $e_3 \neq e_4$. Then, by the homogeneity of G , we can find automorphisms α and β of G sending a to c and b to d , which suffices for the prove of our claim.

In the sequel, we will assume that $a_1, a_2 \in X$ and $b_3, b_4 \in Y$, where $X, Y \in \{E, N\}$.

Case 1. Suppose first that $a_3 = a_4 \in \{a_1, a_2\}$ and $b_1 = b_2 \in \{b_3, b_4\}$; without loss of generality $a_3 = a_2$ and $b_1 = b_3$.

Case 1.1 There exists $u \in V$ such that for all $p, v \in V$ with $(u, v) \in Y$ we have $f(p, u) = f(p, v)$. Then, because f preserves \neq , we have $f(p, u) \neq f(q, u)$ for all $p \neq q$. Since f is essential there are $p, v \in V$ such that $f(p, u) \neq f(p, v)$: for if $f(p, u) = f(p, v)$ for all $p, v \in V$, then $f(p, v) = f(p, v')$ for all $p, v, v' \in V$, and hence f is not essential. Pick $w \in V$ such that $(w, u), (w, v) \in Y$. Pick moreover $q \in V$ such that $(p, q) \in X$. We have $f(p, v) \neq f(p, u) = f(p, w)$. Moreover, $f(p, w) = f(p, u) \neq f(q, u) = f(q, w)$. Hence, the tuples $c := (q, p, p, p)$ and $d := (w, w, w, v)$ prove the claim.

Case 1.2 For all $u \in V$ there exist $p, v \in V$ with $(u, v) \in Y$ such that $f(p, u) \neq f(p, v)$. Pick $m, n, u \in V$ with $(m, n) \in X$ and $f(m, u) \neq f(n, u)$. Pick $p, v \in V$ such that $(u, v) \in Y$ and $f(p, u) \neq f(p, v)$. If we can pick p in such a way that $(p, m), (p, n) \in X$, then since either $f(m, u) \neq f(p, u)$ or $f(n, u) \neq f(p, u)$ we have that either (m, p, p, p) or (n, p, p, p) proves the claim together with the tuple (u, u, u, v) . So suppose that this is impossible. Then for any $q \in V$ with $(q, m), (q, n) \in X$ we have $f(q, u) = f(q, v) \neq f(p, u)$, so we have that (q, p, p, p) and (u, u, u, v) satisfy the claim.

Case 2. Now suppose that $a_3 = a_4 \in \{a_1, a_2\}$ and $b_1 = b_2 \notin \{b_3, b_4\}$; without loss of generality $a_3 = a_2$. Write $(b_1, b_3) \in Q_3$ and $(b_1, b_4) \in Q_4$, where $Q_3, Q_4 \in \{E, N\}$.

Case 2.1 There exists $u \in V$ such that for all p, v, r with $vr \in Y$, $(u, v) \in Q_3$ and $(u, r) \in Q_4$ we have $f(p, v) = f(p, r)$. Then one easily concludes that for all $p \in V$ and all $v, v' \in V$ with $v, v' \neq u$ we have $f(p, v) = f(p, v')$. This implies that $f(p, v) \neq f(q, v)$ whenever $p \neq q$ and $v \neq u$. Since f is essential, there exist $p, v \in V$ with $(u, v) \in Y$ such that $f(p, u) \neq f(p, v)$. Now pick $w, q \in V$ such that $(w, u) \in Q_3$, $(w, v) \in Q_4$, and $(q, p) \in X$. Then $f(p, w) \neq f(q, w)$, and so the tuples (q, p, p, p) and (w, w, u, v) prove the claim.

Case 2.2 For all u there exist p, v, r with $(v, r) \in Y$, $(u, v) \in Q_3$, $(u, r) \in Q_4$ and $f(p, v) \neq f(p, r)$. Pick m, n, u with $(m, n) \in X$ and $f(m, u) \neq f(n, u)$. Pick $p, v, r \in V$ such that $(v, r) \in Y$, $(u, v) \in Q_3$, $(u, r) \in Q_4$ and $f(p, v) \neq f(p, r)$. If we can pick p in such a way that $(p, m), (p, n) \in X$, then either (m, p, p, p) and (u, u, v, r) or (n, p, p, p) and (u, u, v, r) prove the claim. So suppose that this is impossible. Then for any q with $(q, m), (q, n) \in X$ and all $v, r \in V$ with $(v, r) \in Y$, $(u, v) \in Q_3$, $(u, r) \in Q_4$ we have $f(q, v) = f(q, r)$. This implies that for all such q and all $v, v' \neq u$ we have $f(q, v) = f(q, v')$. Pick w such that $(w, v) \in Q_3$, $(w, r) \in Q_4$. Pick q such that $(q, p) \in X$. We have $f(q, w) \neq f(p, w)$, and so (q, p, p, p) and (w, w, v, r) prove the claim.

Case 3. To finish the proof, suppose that $a_3 = a_4 \notin \{a_1, a_2\}$ and $b_1 = b_2 \notin \{b_3, b_4\}$. Write $(a_3, a_1) \in P_1$, $(a_3, a_2) \in P_2$, $(b_1, b_3) \in Q_3$ and $(b_1, b_4) \in Q_4$, where $P_i, Q_i \in \{E, N\}$.

Case 3.1 There exists u such that for all p, v, r with $(v, r) \in Y$, $(u, v) \in Q_3$ and $(u, r) \in Q_4$ we have $f(p, v) = f(p, r)$. Then one easily concludes that for all $p \in V$ and all $v, v' \in V$ with $v, v' \neq u$ we have $f(p, v) = f(p, v')$. This implies that $f(p, v) \neq f(q, v)$ whenever $p \neq q$ and $v \neq u$. Since f is essential, there exist p, v with $(u, v) \in Y$ such that $f(p, u) \neq f(p, v)$. Now pick w, m, n such that $wu \in Q_3$, $(w, v) \in Q_4$, $(m, n) \in X$, $mp \in P_1$, and $(n, p) \in P_2$. Then the tuples (m, n, p, p) and (w, w, u, v) prove the claim.

Case 3.2 For all u there exist p, v, r with $(v, r) \in Y$, $(u, v) \in Q_3$, $(u, r) \in Q_4$ and $f(p, v) \neq f(p, r)$. Pick m, n, u with $(m, n) \in X$ and $f(m, u) \neq f(n, u)$. Pick p, v, r such that $(v, r) \in Y$, $(u, v) \in Q_3$, $(u, r) \in Q_4$ and $f(p, v) \neq f(p, r)$. If we can pick p in such a way that $(p, m) \in P_1$ and $(p, n) \in P_2$, then (m, n, p, p) and (u, u, v, r) prove the claim, so suppose that this is impossible. Then for any q with $(q, m) \in P_1$ and $(q, n) \in P_2$ and all v, r with $(v, r) \in Y$, $(u, v) \in Q_3$, $(u, r) \in Q_4$ we have $f(q, v) = f(q, r)$. This is easily seen to imply that for all such q and all $v, v' \neq u$ we have $f(q, v) = f(q, v')$. Pick w such that $(w, v) \in Q_3$, $(w, r) \in Q_4$, and $w \neq u$. Pick q, q' such that $(q, q') \in X$, $(q, p) \in P_1$ and $(q', p) \in P_2$. We have $f(q, w) \neq f(q', w)$, and thus (q, q', p, p) and (w, w, v, r) prove the claim. \square

8. THE ORDERED GRAPH PRODUCT RAMSEY LEMMA

In order to find the minimal functions which are not unary, we need to develop the Ramsey-theoretic tools that allow us to find patterns in the behavior of such higher arity functions; this is the purpose of this section.

Definition 43. Let $\Gamma_1, \dots, \Gamma_n$ be structures. For a tuple x in the cartesian product $\Gamma := \Gamma_1 \times \dots \times \Gamma_n$, we write x_i for the i -th coordinate of x . The *type* of a sequence of tuples $a^1, \dots, a^m \in \Gamma$, denoted by $\text{tp}_{\text{qf}}(a^1, \dots, a^m)$, is the cartesian product of the types of (a_i^1, \dots, a_i^m) in Γ_i .

Definition 44. An *ordered graph* is a graph with an additional total order \prec on the vertices. An *ordered graph product* is a cartesian product of ordered graphs.

The following extends the definition of a canonical function on graphs, n -partitioned graphs, and n -constant graphs from Section 5 to functions on (ordered) graph products.

Definition 45. Let F_1, \dots, F_n, Z be (ordered) graphs. Set $F := F_1 \times \dots \times F_n$. An operation $g: F \rightarrow Z$ is *canonical* iff for all $x, y, u, v \in F$ with $\text{tp}_{\text{qf}}(x, y) = \text{tp}_{\text{qf}}(u, v)$ we have $\text{tp}_{\text{qf}}(g(x), g(y)) = \text{tp}_{\text{qf}}(g(u), g(v))$.

Lemma 46 (The ordered graph product Ramsey lemma). *For every finite ordered graph product $F := F_1 \times \dots \times F_n$ there exists a finite ordered graph product $H := H_1 \times \dots \times H_n$ such that for all functions $f: H \rightarrow Z$ to an ordered graph Z there is a copy F' of F in H on which f is canonical.*

Note that Lemma 46 is not true if the graphs are not ordered: let $n = 2$, and let I_2 be the graph which has two vertices and no edges. Set F_1 and F_2 equal to I_2 . Suppose H exists, and order its components H_1, H_2 linearly. Define on the domain of H the following graph Z : Two pairs x, y are connected by an edge iff they are comparable in the product order. Set $f: H \rightarrow Z$ to be the identity function. Now whenever $x_i, y_i \in H_i$ are so that they induce copies F'_i of F_i , for $i = 1, 2$, then f is not canonical on $F'_1 \times F'_2$: for, assume without loss of generality that x_i is smaller than y_i in the order on H_i . Then $\text{tp}_{\text{qf}}((x_1, x_2), (y_1, y_2)) = \text{tp}_{\text{qf}}((x_1, y_2), (y_1, x_2))$, but $\text{tp}_{\text{qf}}(f(x_1, x_2), f(y_1, y_2)) \neq \text{tp}_{\text{qf}}(f(x_1, y_2), f(y_1, x_2))$ as $f(x_1, x_2), f(y_1, y_2)$ are connected by an edge in Z but $f(x_1, y_2), f(y_1, x_2)$ are not.

We remark that as we have seen in Section 5, $n > 1$ is necessary for this problem to appear.

Proof of Lemma 46. We prove the following claim: for every type $\text{tp}_{\text{qf}}(u, v)$ of n -tuples u, v in the ordered graph product F there exists a finite ordered graph product $H := H_1 \times \dots \times H_n$ such that whenever $f: H \rightarrow Z$ is a function, then there is a copy F' of F in H with the property that $\text{tp}_{\text{qf}}(f(x), f(y)) = \text{tp}_{\text{qf}}(f(a), f(b))$ for all $x, y, a, b \in F'$ such that $\text{tp}_{\text{qf}}(x, y) = \text{tp}_{\text{qf}}(a, b) = \text{tp}_{\text{qf}}(u, v)$. Repeated use of the claim for all types of pairs in F then proves the lemma.

In fact, we prove the following more abstract statement, which clearly implies our original claim: for every type $\text{tp}_{\text{qf}}(u, v)$ of two n -tuples in an ordered graph product F and for any $k < \omega$ there exists a finite ordered graph product $H := H_1 \times \dots \times H_n$ such that whenever χ is a coloring of the pairs (x, y) in H with $\text{tp}_{\text{qf}}(x, y) = \text{tp}_{\text{qf}}(u, v)$ with k colors, then there is a copy F' of F in H on which the coloring is constant.

To prove the claim, we use induction over n . The induction beginning $n = 1$ is just a subset of the proof of Corollary 18 (as there, we had to introduce an order for the sake of the proof, and now we are already given ordered graphs). So suppose $n > 1$ and that the claim holds for all $i < n$. Let n -tuples $u, v \in F$ defining the type be given. Set $u' := (u_1, \dots, u_{n-1})$, and define

v' analogously. By induction hypothesis, there is an ordered graph product $H_1 \times \cdots \times H_{n-1}$ such that whenever its pairs (x', y') with $\text{tp}_{\text{qf}}(x', y') = \text{tp}_{\text{qf}}(u', v')$ are colored with k colors, then there is a copy of $F_1 \times \cdots \times F_{n-1}$ in $H_1 \times \cdots \times H_{n-1}$ on which the coloring is constant. Let m be the number of pairs (x', y') in $H_1 \times \cdots \times H_{n-1}$ which have type $\text{tp}_{\text{qf}}(u', v')$. By induction hypothesis, there is an ordered graph $H_{n,1}$ with the property that whenever its pairs (x_n, y_n) with $\text{tp}_{\text{qf}}(x_n, y_n) = \text{tp}_{\text{qf}}(u_n, v_n)$ are colored with k colors, then it contains a monochromatic copy of F_n . Further, there is an ordered graph $H_{n,2}$ with the property that whenever its subsets of this type are colored with k colors, then it contains a monochromatic copy of $H_{n,1}$. Continue constructing ordered graphs like that, arriving at $H_n := H_{n,m}$. We claim that $H := H_1 \times \cdots \times H_n$ has the desired properties. To see this, let a coloring χ of the pairs of H of type $\text{tp}_{\text{qf}}(u, v)$ be given. Let $(x^1, y^1), \dots, (x^m, y^m)$ be an enumeration of all the pairs in $H_1 \times \cdots \times H_{n-1}$ which have type $\text{tp}_{\text{qf}}(u', v')$. For all $1 \leq i \leq m$, define a coloring χ^i of the pairs (p, q) of H_n of type $\text{tp}_{\text{qf}}(u_n, v_n)$ by setting $\chi^i(p, q) := \chi(x^i \smile p, y^i \smile q)$, where $a \smile b$ denotes the concatenation of two tuples a, b . By thinning out H_n m times, we obtain a copy F'_n of F_n in H_n on which each coloring χ^i is constant with color c^i . Now by that construction, all pairs (x^i, y^i) have been assigned a color c^i , the assignment thus being a coloring of all the pairs of type $\text{tp}_{\text{qf}}(u', v')$ in $H_1 \times \cdots \times H_{n-1}$. By the choice of that product, there is a copy $F'_1 \times \cdots \times F'_{n-1}$ of $F_1 \times \cdots \times F_{n-1}$ in $H_1 \times \cdots \times H_{n-1}$ on which that coloring is constant, say with value r . But that means that if $x, y \in F'_1 \times \cdots \times F'_n$ have type $\text{tp}_{\text{qf}}(u, v)$, then $\chi(x, y) = r$, proving our statement. \square

9. MINIMAL BINARY FUNCTIONS

We know from Theorem 38 and Lemma 39 that all essential minimal functions are binary, injective, and preserve both E and N . It is the goal of this section to determine these binary minimal functions.

Let V be equipped with a total order \prec in such a way that $(V; E, \prec)$ is the random ordered graph, i.e., the unique countably infinite homogeneous graph containing all finite ordered graphs (for existence and uniqueness of this structure, see e.g. [22]). The order $(V; \prec)$ is then isomorphic to the order of the rationals \mathbb{Q} .

From now on, until Proposition 54, we see the random graph equipped with this order, in particular when talking about canonical behavior of functions. Note in this context that a function $f: V^k \rightarrow V$ which is canonical with respect to the language of ordered graphs need not be canonical in the language of ordinary graphs; the converse implication does not hold either.

Proposition 47. *Every function $f: V^k \rightarrow V$ is canonical on arbitrarily large finite ordered graph products. In particular, every binary injection generates a binary injection which is canonical with respect to the language of ordered graphs.*

Proof. The first statement is a direct consequence of the ordered graph product Ramsey lemma (Lemma 46). The second statement follows from the fact that there are only finitely many canonical behaviors on every finite ordered graph product, and by local closure. \square

The following is also straightforward to verify.

Proposition 48. *If a function $f: V^k \rightarrow V$ is canonical with respect to some base structure (i.e., the random graph or the ordered random graph), then so are all functions it generates.*

Hence, by Propositions 47 and 48 all minimal binary injections are canonical as functions on the ordered random graph. In the following, we determine those canonical behaviors of binary injections that yield minimal functions.

Definition 49. Let $f: V^2 \rightarrow V$ be injective. If for all $(u_1, u_2), (v_1, v_2) \in V^2$ with $u_1 \prec v_1$ and $u_2 \prec v_2$ we have

- $(f(u_1, u_2), f(v_1, v_2)) \in E$ if and only if $(u_1, v_1) \in E$ and $(u_2, v_2) \in E$, then we say that f behaves like *min on input* (\prec, \prec) .
- $(f(u_1, u_2), f(v_1, v_2)) \in N$ if and only if $(u_1, v_1) \in N$ and $(u_2, v_2) \in N$, then we say that f behaves like *max on input* (\prec, \prec) .
- $(f(u_1, u_2), f(v_1, v_2)) \in E$ if and only if $(u_1, v_1) \in E$, then we say that f behaves like *p_1 on input* (\prec, \prec) .
- $(f(u_1, u_2), f(v_1, v_2)) \in E$ if and only if $(u_2, v_2) \in E$, then we say that f behaves like *p_2 on input* (\prec, \prec) .

Analogously, we define behavior on input (\prec, \succ) using pairs $(u_1, u_2), (v_1, v_2) \in V^2$ with $u_1 \prec v_1$ and $u_2 \succ v_2$.

Of course, we could also have defined “behavior on input (\succ, \succ) ” and “behavior on input (\succ, \prec) ”; however, behavior on input (\succ, \succ) equals behavior on input (\prec, \prec) , and behavior on input (\succ, \prec) equals behavior on input (\prec, \succ) . Thus, there are only two kinds of inputs to be considered, namely the “straight input” (\prec, \prec) and the “twisted input” (\prec, \succ) .

Proposition 50. *Let $f: V^2 \rightarrow V$ be injective and canonical as a function on the ordered random graph, and suppose it preserves E and N . Then it behaves like *min*, *max*, *p_1* or *p_2* on input (\prec, \prec) (and similarly on input (\prec, \succ)).*

Proof. By definition of the term canonical; one only needs to enumerate all possible types $\text{tp}_{\text{qf}}(x, y)$ of pairs $x, y \in V^2$ with respect to the ordered random graph. \square

We remark that the four possibilities correspond to the four binary operations g on the two-element domain $\{E, N\}$ that are *idempotent*, i.e., that satisfy that $g(E, E) = E$ and $g(N, N) = N$.

Definition 51. If $f: V^2 \rightarrow V$ behaves like X on input (\prec, \prec) and like Y on input (\prec, \succ) , where $X, Y \in \{\text{max}, \text{min}, p_1, p_2\}$, then we say that f is of *type* X/Y .

Observe that in Proposition 50, we did not care about the fact that a canonical injection $f: V^2 \rightarrow V$ also behaves regularly with respect to the order: The latter implies, for example, that f is either strictly increasing or decreasing with respect to the pointwise order, i.e., $x_1 \prec y_1$ and $x_2 \prec y_2$ either always implies $f(x_1, x_2) \prec f(y_1, y_2)$, or it always implies $f(x_1, x_2) \succ f(y_1, y_2)$. Fix from now on any automorphism α of the graph G that reverses the order on V . By applying α to f if necessary, we may assume that f is strictly increasing, which will be a tacit assumption from now on. Having that, one easily checks that f satisfies one of the implications

$$x_1 \prec y_1 \wedge x_2 \neq y_2 \rightarrow f(x_1, x_2) \prec f(y_1, y_2)$$

and

$$x_1 \neq y_1 \wedge x_2 \prec y_2 \rightarrow f(x_1, x_2) \prec f(y_1, y_2).$$

In the first case, we say that f *obeys p_1 for the order*, in the second case f *obeys p_2 for the order*. By switching the variables of f , we may always assume that f obeys p_1 for the order, and we shall do so from now on.

We will now prove that minimal binary canonical injections are never of mixed type, i.e., they have to behave the same way on straight and twisted inputs.

Proposition 52. *Let $f: V^2 \rightarrow V$ be injective and canonical as a function on the ordered random graph. Suppose moreover that f is of type X/Y , where $X, Y \in \{\max, \min, p_1, p_2\}$ and $X \neq Y$. Then f is not minimal.*

Proof. Suppose first that f is of type \max/p_i or of type p_i/\max , where $i \in \{1, 2\}$. We claim that f generates a binary injective canonical function g which is of type \max/\max . Clearly, all binary injective canonical functions generated by g then are also of type \max/\max , so g cannot generate f , which shows that f is not minimal. Assume without loss of generality that f is of type \max/p_i , and note that we assume that f obeys p_1 for the order. Set $h(u, v) := f(u, \alpha(v))$. Then h behaves like p_i on input (\prec, \prec) and like \max on input (\prec, \succ) ; moreover, $f(x_1, x_2) \prec f(y_1, y_2)$ iff $h(x_1, x_2) \prec h(y_1, y_2)$, for all $x_1 \neq y_1$ and $x_2 \neq y_2$. We then have that $g(u, v) := f(f(u, v), h(u, v))$ is of type \max/\max , finishing the proof of our claim.

If f is of type \min/p_i or of type p_i/\min , where $i \in \{1, 2\}$, then the dual proof works.

Consider the case where f is of type \max/\min or of type \min/\max . Assume without loss of generality that f is of type \max/\min , and remember that we assume that f obeys p_1 for the order. Consider $h(u, v) := f(f(u, v), \alpha(v))$. Then h is of type p_2/p_2 , so it cannot reproduce f . Hence f is not minimal.

To finish the proof, suppose that f is of type p_1/p_2 or of type p_2/p_1 . If f is of type p_1/p_2 , then $h(u, v) := f(f(u, v), \alpha(v))$ is of type p_2/p_2 and cannot reproduce f . If f is of type p_2/p_1 , then $g(u, v) := f(u, \alpha(v))$ is of type p_1/p_2 and still obeys p_1 for the order; hence, we are back in the first case. \square

This motivates the following definition.

Definition 53. Let $f: V^2 \rightarrow V$. We say that f behaves like $\min(\max, p_1, p_2)$ on input (\neq, \neq) iff it behaves like $\min(\max, p_1, p_2)$ both on input (\prec, \prec) and on input (\prec, \succ) . We also say that f is of type $\min(\max, p_1, p_2)$. If f is of type p_1 or p_2 then we also say that f is of type *projection*.

Our observations so far can be summarized as follows.

Proposition 54. *Let $f: V^2 \rightarrow V$ be essential and minimal. Then it is injective, canonical as a function on the (non-ordered) random graph and behaves like \min, \max, p_1 or p_2 on input (\neq, \neq) .*

Proof. We know that f generates a binary injection which is canonical as a function on the random ordered graph, hence it is itself such a function. Since by Proposition 52, f cannot have a “mixed” behavior, it behaves like \min, \max, p_1 or p_2 , and hence is also canonical as a function on the non-ordered random graph. \square

In the following, we will thus forget about the order that we imposed on the random graph, and use the terms “canonical” and “behavior” relative to G . We now consider further types of tuples $x, y \in V^2$: So far, we did not look at the case where $x_1 = y_1$ or $x_2 = y_2$.

Definition 55. Let $f: V^2 \rightarrow V$. We say that f behaves like $e_E(e_N, \text{id}, -)$ on input $(\neq, =)$ iff for every fixed $c \in V$, the function $g(x) := f(x, c)$ behaves like $e_E(e_N, \text{id}, -)$. Similarly we define behavior on input $(=, \neq)$.

If f is canonical and injective, then it behaves like one of the mentioned functions on input $(\neq, =)$ and $(=, \neq)$, respectively. This motivates the following

Definition 56. We say that $f: V^2 \rightarrow V$ is of type E/N iff f behaves like e_E on input $(\neq, =)$ and like e_N on input $(=, \neq)$. Similarly we define the types E/E , N/E , E/id , $E/-$, etc. Moreover, we say that f is *balanced* iff it is of type id/id , we say it is *E -dominated* iff it is of type E/E , and we say it is *N -dominated* iff it is of type N/N .

In the following proposition we finally characterize those canonical behaviors that yield minimal functions.

Proposition 57. *The essential minimal operations on G are precisely the binary injective canonical operations of the following types:*

- (1) *Projection and balanced.*
- (2) *max and balanced.*
- (3) *min and balanced.*
- (4) *max and E -dominated.*
- (5) *min and N -dominated.*
- (6) *Projection and E -dominated.*
- (7) *Projection and N -dominated.*
- (8) *p_2 and E/id , or p_1 and id/E .*
- (9) *p_2 and N/id , or p_1 and id/N .*

Moreover, these 9 different kinds of minimal functions do not generate one another. Furthermore, any two functions in the same group do generate one another.

Proof. By Proposition 54 we know that all minimal essential functions are necessarily canonical binary injections of type min, max, or projection. Therefore, we must show that out of those functions, the minimal ones are precisely those listed above. Let henceforth f be a canonical binary injection of type min, max, or projection.

Let us first prove that if f is listed above, then it is indeed minimal. To this end, observe first that by the homogeneity of G and local closure, f then generates all other functions in its class of the theorem. Next note the following facts which can easily be proven by a standard induction over terms.

- Any binary essential function generated by a binary canonical injection of type min, max, or projection, respectively, is of the same type.
- Any binary essential function generated by a binary canonical injection that is balanced and preserves E and N is balanced.

It follows immediately that the if f belongs to the first three classes of the proposition, then it is minimal.

It is easy to verify that any binary essential function generated by an E -dominated binary canonical injection of type max is E -dominated. Dually, any binary essential function generated by an N -dominated binary canonical injection of type min is N -dominated. These two facts imply minimality in case f belongs to items (4) or (5).

Observe next that any binary essential function generated by an E -dominated binary canonical injection of type projection is E -dominated. Dually, any binary essential function generated by an N -dominated binary canonical injection of type projection is N -dominated. This implies minimality for the case where f belongs to items (6) or (7).

To prove minimality for the case where f falls into items (8) or (9) we claim the following: Any binary essential function generated by a binary canonical injection of type E/id and p_2 is either of the same type or of type id/E and p_1 . Dually, any binary essential function generated by a binary canonical injection of type N/id and p_2 generates is either of the same

type or of type id/N and p_1 . To see this, let $f(u, v)$ be of type E/id and p_2 . $f(v, u)$ is of type id/E and p_1 . Both $f(u, f(u, v))$ and $f(v, f(u, v))$ are of type E/id and p_2 . So is $f(f(u, v), v)$. The function $f(f(u, v), u)$ is of type id/E and p_1 . Finally, $f(f(u, v), f(v, u))$ also is of type id/E and p_1 , so f cannot generate any new typesets.

Next we show that if f does not belong to any of the listed classes, then it is not minimal.

Suppose first it is of type max . We claim that if f is not balanced or E -dominated, then f is not minimal. We go through all possibilities: If f is of type E/id , then $g(x, y) := f(f(x, y), x)$ is E -dominated. By our observation above, g cannot reproduce f . If f is of type E/N , then g is E -dominated as well. So it is if f is of type $E/-$. If f is of type N/id , then $g(x, y) := f(x, f(x, y))$ is balanced, so f is not minimal by the above. If f is of type $N/-$, then g is balanced as well. If f is of type $\text{id}/-$ or of type $-/-$, then $g(x, y) := f(x, f(x, y))$ is of type E/id , which we have already shown not to be minimal. By symmetry, if we switch the arguments in a type of f , e.g., if f is of type id/E , then f is not minimal either. We have thus covered all possible types.

The dual argument works if f is a binary canonical injection of type min : If f is not balanced or N -dominated, then f is not minimal.

Suppose now that f is of type p_1 . We claim that if f is not balanced, E -dominated, N -dominated, of type id/E , or of type id/N , then f is not minimal. To see this, we distinguish all possible cases: If f is of type E/id , $E/-$, $-/\text{id}$, or $-/-$, then $g(x, y) := f(x, f(x, y))$ is balanced and cannot reproduce f . If it is of type E/N or $\text{id}/-$, then g is of type E/id , and we are back in the preceding case. Dually, if f is of type N/id or $N/-$, then g is balanced. If it is of type N/E , then g is of type N/id , bringing us back to the preceding case. If it is of type $-/E$, then g is of type id/E and p_1 , and hence cannot reproduce f by the above. The dual argument works if f is of type $-/N$.

Finally, let f be of type p_2 . Then the same argument as above shows that if f is not balanced, E -dominated, N -dominated, of type E/id , or of type N/id , then f is not minimal. This finishes the proof. \square

To summarize, we now restate and prove Theorems 3 and 5 in the following

Theorem 58 (Summary of Theorems 3 and 5). *Any minimal function e on the random graph is equivalent to exactly one of the following operations: a constant operation; e_N ; e_E ; $-$; sw ; or*

- (6) *a binary injection of type p_1 that is balanced in both arguments;*
- (7) *a binary injection of type max that is balanced in both arguments;*
- (8) *a binary injection of type max that is E -dominated in both arguments;*
- (9) *a binary injection of type p_1 that is E -dominated in both arguments;*
- (10) *a binary injection of type p_1 that is balanced in the first and E -dominated in the second argument;*

or to one of the duals of the last four operations (the operation in (6) is self-dual).

Proof. If e is not essential, then it generates, and hence is equivalent to, a constant operation, e_N , e_E , $-$ or sw ; this is the content of Theorem 3, which we already proved in Section 6. We also have argued that these functions do not generate each other.

If e is essential, then it must preserve E and N , by Lemma 39, and it must be binary and injective by Theorem 38. By Proposition 57 it must be canonical and of one of the 9 types listed there; moreover, Proposition 57 shows that functions of different types do not generate each other. Observe that in Proposition 57, class (1) is (6) here, (2) is (7) here, (3) is the

dual of (7) here, (4) is (8) here, (5) is the dual of (8) here, (6) is (9) here, (7) is the dual of (9) here, (8) is (10) here, and (9) is the dual of (10) here. \square

We conclude this paper by remarking that the relations that are preserved by one of the essential operations in Theorems 3 and 5 (i.e., the relations of the structures in Corollary 11) also have syntactic descriptions. For instance, it is not hard to show (see [3]) that a relation R with a first-order definition in G is preserved by a binary operation of type min that is N -dominated in both arguments if and only if R can be defined by a quantifier-free Horn formula over $(V; E, =)$ (i.e., by a quantifier-free formula in conjunctive normal form where each clause contains at most one literal of the form $E(x, y)$ or $x = y$).

REFERENCES

- [1] Fred G. Abramson and Leo Harrington. Models without indiscernibles. *Journal of Symbolic Logic*, 43(3):572–600, 1978.
- [2] Manuel Bodirsky and Hubert Chen. Oligomorphic clones. *Algebra Universalis*, 57(1):109–125, 2007.
- [3] Manuel Bodirsky, Hubie Chen, Jan Kára, and Timo von Oertzen. Maximal infinite-valued constraint languages. *Theoretical Computer Science (TCS)*, 410:1684–1693, 2009. A preliminary version appeared at ICALP’07.
- [4] Manuel Bodirsky, Hubie Chen, and Michael Pinsker. The reducts of equality up to primitive positive interdefinability. *Journal of Symbolic Logic*, 75(4):1249–1292, 2010.
- [5] Manuel Bodirsky, Peter Jonsson, and Timo von Oertzen. Horn versus full first-order: a complexity dichotomy for algebraic constraint satisfaction problems. *Journal of Logic and Computation*, 2010. To appear.
- [6] Manuel Bodirsky and Jan Kára. The complexity of equality constraint languages. *Theory of Computing Systems*, 3(2):136–158, 2008. A conference version appeared in the proceedings of CSR’06.
- [7] Manuel Bodirsky and Jan Kára. The complexity of temporal constraint satisfaction problems. *Journal of the ACM*, 57(2):41 pp, 2009. An extended abstract appeared in the proceedings of STOC’08.
- [8] Manuel Bodirsky and Jaroslav Nešetřil. Constraint satisfaction with countable homogeneous templates. *Journal of Logic and Computation*, 16(3):359–373, 2006.
- [9] Manuel Bodirsky and Michael Pinsker. Reducts of Ramsey structures. *AMS Contemporary Mathematics*, vol. 558 (*Model Theoretic Methods in Finite Combinatorics*), pages 489–519, 2011.
- [10] Manuel Bodirsky and Michael Pinsker. Schaefer’s theorem for graphs. In *Proceedings of STOC*, pages 655–664, 2011. Preprint of the long version available at arxiv.org/abs/1011.2894.
- [11] Manuel Bodirsky and Michael Pinsker. Topological Birkhoff. Preprint, arxiv.org/abs/1203.1876, 2012.
- [12] Manuel Bodirsky, Michael Pinsker, and Todor Tsankov. Decidability of definability. In *Proceedings of LICS*, pages 321–328, 2011. Preprint of full journal version available from arxiv.org/abs/1012.2381.
- [13] Anthony Bonato and Dejan Delić. The monoid of the random graph. *Semigroup Forum*, 61:138–148, 2000.
- [14] Anthony Bonato, Dejan Delić, and Igor Dolinka. All countable monoids embed into the monoid of the infinite random graph. *Accepted to Discrete Mathematics*, 2009.
- [15] Andrei A. Bulatov. A dichotomy theorem for constraint satisfaction problems on a 3-element set. *Journal of the ACM*, 53(1):66–120, 2006.
- [16] Peter J. Cameron. Transitivity of permutation groups on unordered sets. *Mathematische Zeitschrift*, 148:127–139, 1976.
- [17] Peter J. Cameron. *Oligomorphic Permutation Groups*. Cambridge University Press, Cambridge, 1990.
- [18] Peter J. Cameron. The random graph. *Algorithms and Combinatorics*, 14:333–351, 1997.
- [19] Peter J. Cameron. The random graph revisited. In *Proceedings of the European Congress of Mathematics*, volume 201, pages 267–274. Birkhäuser, 2001.
- [20] Dejan Delić and Igor Dolinka. The endomorphism monoid of the random graph has uncountably many ideals. *Semigroup Forum*, 69:75–79, 2004.
- [21] Martin Goldstern and Michael Pinsker. A survey of clones on infinite sets. *Algebra Universalis*, 59:365–403, 2008.
- [22] Wilfrid Hodges. *A shorter model theory*. Cambridge University Press, Cambridge, 1997.

- [23] Markus Junker and Martin Ziegler. The 116 reducts of $(\mathbb{Q}, <, a)$. *Journal of Symbolic Logic*, 74(3):861–884, 2008.
- [24] Jaroslav Nešetřil and Vojtěch Rödl. Partitions of finite relational and set systems. *J. Combinatorial Theory Series A*, 22(3):289–312, 1977.
- [25] Jaroslav Nešetřil. Ramsey theory. *Handbook of Combinatorics*, pages 1331–1403, 1995.
- [26] Jaroslav Nešetřil and Vojtěch Rödl. Ramsey classes of set systems. *Journal of Combinatorial Theory, Series A*, 34(2):183–201, 1983.
- [27] Jaroslav Nešetřil and Vojtěch Rödl. The partite construction and Ramsey set systems. *Discrete Mathematics*, 75(1-3):327–334, 1989.
- [28] Michael Pinsker. More sublattices of the lattice of local clones. *Order*, 27(3):353–364, 2010.
- [29] Reinhard Pöschel and Lev A. Kalužnin. *Funktionen- und Relationenalgebren*. Deutscher Verlag der Wissenschaften, 1979.
- [30] Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of STOC*, pages 216–226, 1978.
- [31] Ágnes Szendrei. *Clones in universal Algebra*. Séminaire de Mathématiques Supérieures. Les Presses de l'Université de Montréal, 1986.
- [32] Simon Thomas. Reducts of the random graph. *Journal of Symbolic Logic*, 56(1):176–181, 1991.
- [33] Simon Thomas. Reducts of random hypergraphs. *Annals of Pure and Applied Logic*, 80(2):165–193, 1996.

LABORATOIRE D'INFORMATIQUE (LIX), CNRS UMR 7161, ÉCOLE POLYTECHNIQUE, 91128 PALAISEAU, FRANCE

E-mail address: bodirsky@lix.polytechnique.fr

URL: <http://www.lix.polytechnique.fr/~bodirsky/>

ÉQUIPE DE LOGIQUE MATHÉMATIQUE, UNIVERSITÉ DIDEROT-PARIS 7, UFR DE MATHÉMATIQUES, 75205 PARIS CEDEX 13, FRANCE

E-mail address: marula@gmx.at

URL: <http://dmg.tuwien.ac.at/pinsker/>