Divisibility of Dedekind Finite Sets

David Blair*
Mathematics Department
University of Michigan
Ann Arbor, MI 48109–1109, U.S.A.
dblair@umich.edu

Andreas Blass†
Mathematics Department
University of Michigan
Ann Arbor, MI 48109–1109, U.S.A.
ablass@umich.edu

Paul Howard
Mathematics Department
Eastern Michigan University
Ypsilanti, MI 48197, U.S.A.
phoward@emunix.emich.edu

April 20, 2002

Abstract

A Dedekind-finite set is said to be divisible by a natural number n if it can be partitioned into pieces of size n. We study several aspects of this notion, as well as the stronger notion of being partitionable

*Partially supported by the NSF's Research Experiences for Undergraduates program.
†Partially supported by NSF grants DMS–9505118 and DMS–0070723. Thanks also to the Mittag-Leffler Institute for its support of a visit in October, 2000.

into n pieces of equal size. Among our results are that the divisors of a Dedekind-finite set can consistently be any set of natural numbers (containing 1 but not 0), that a Dedekind-finite power of 2 cannot be divisible by 3, and that a Dedekind-finite set can be congruent modulo 3, to all of 0, 1, and 2 simultaneously. (In these results, 2 and 3 serve as typical examples; the full results are more general.)

1 Introduction

We begin with definitions of the terms in the title.

Definition 1.1 A set X (or its cardinality) is *Dedekind-finite* if X has no countably infinite subset.

In the presence of the axiom of choice (AC), Dedekind finiteness is equivalent to (ordinary) finiteness, i.e., to being in one-to-one correspondence with $\{1, 2, \ldots, n\}$ for some natural number n. But we shall work in set theory without AC, so it is possible to have infinite Dedekind-finite sets. We shall primarily consider divisibility in the following sense.

Definition 1.2 A set X (or its cardinality) is *divisible* by a natural number n if X can be partitioned into pieces each having cardinality n.

We shall also consider the following natural alternative notion, along with some other variants in Section 6.

Definition 1.3 A set X (or its cardinality) is *strongly divisible* by a natural number n if X can be partitioned into n pieces, all of the same cardinality.

Typical examples of the sorts of results that we shall prove are the following. A Dedekind-finite cardinal of the form 2^m cannot be divisible by 3, but it is possible for 3 to divide 2^m with m Dedekind-finite. Every set of natural numbers that contains 1 but not 0 occurs as the set of divisors of an infinite, Dedekind-finite set in some model of set theory. It is possible for a Dedekind-finite cardinal to be divisible by 3 and remain divisible by 3 when any finite number is added or subtracted.

To help orient the reader, we list in the next proposition some well-known equivalents of Dedekind-finiteness, but we omit the proof since these equivalents will be used only in incidental remarks, not in any of our proofs.

Proposition 1.4 *For any set X, the following are equivalent.*

- *X is Dedekind-finite.*
- *All well-orderable subsets of X are finite.*
- *$X \not\cong X + 1$.*
- *For all Y, Z, if $X + Y \cong X + Z$ then $Y \cong Z$.*

Notation

As is customary in set theory, we identify each natural number n with the set $\{0, 1, \ldots, n-1\}$ of its predecessors. We assume that cardinal numbers $|X|$ of sets X have been defined so that $|X| = |Y|$ if and only if $X \cong Y$, which means that there is a bijection between X and Y. We also assume that the cardinal numbers of finite sets are the natural numbers, $|n| = n$. The details of the definition of cardinal numbers of infinite sets won't matter. (A common convention is "Scott's trick": $|X|$ is the collection of all sets that are $\cong X$ and of smallest possible rank. Since we don't assume the axiom of choice, we cannot use the more common convention that $|X|$ is the first ordinal $\cong X$, for there might be no such ordinal.)

We write, as usual, $|X| \leq |Y|$ if there is a one-to-one map from X into Y. Thus, X is Dedekind-finite if and only if $\aleph_0 \not\leq |X|$.

We shall use the same symbols for the arithmetic operations on cardinal numbers and for the corresponding operations on sets. Thus, $X + Y$ is the disjoint union of X and Y. (If X and Y are not disjoint, replace Y by a bijective copy. It won't matter which copy is used, since we shall be interested only in properties invariant under bijections.) $X \times Y$ is the set of all pairs (x, y) with $x \in X$ and $y \in Y$. And X^Y is the set of all functions from Y to X.

We shall use the notation $X - n$, when X is an infinite set and n is a natural number, to mean the result of removing n elements from X. The result has the same cardinality regardless of which particular n points are removed. This follows from Proposition 1.4 since n is Dedekind-finite, but it can also be proved very easily by induction on n. Of course, the complementary notation, $X + n$, is covered by the preceding paragraph; it is the result of adding n new elements to X.

2 Permutation Models

We shall work in set theory without the axiom of choice, and our standard reference is Jech's book [3].

There are two versions of set theory without choice, called ZF and ZFA in [3]. They differ in that the latter allows atoms, objects that are not sets but can be members of sets, while the former requires all objects to be sets. All our results will work for either version. Our proofs of "outright" theorems (in contrast to consistency results) will be formalizable in the weaker theory ZFA; it will never matter in these proofs whether atoms are present or not. Our consistency results will be given by the method of permutation models [3, Chapter 4], so they directly give models of only the weaker theory ZFA. Nevertheless, they indirectly give, via the Jech-Sochor theorem [4, 3], models of ZF. The Jech-Sochor theorem provides embeddings of arbitrarily long initial segments of ZFA models into ZF models. All the statements whose consistency we prove depend only on a (very small) initial segment of the ZFA model, so they are preserved by the embedding and we thus obtain their consistency with ZF.

As mentioned above, we use the method of permutation models when giving consistency proofs. Actually, we shall need only a rather special case of this method. In all our constructions, we shall begin with a countably infinite set A of atoms, equipped with some structure (for example some relations on A). We write $\mathfrak{A}$ to mean the set A together with the specified structure. The permutation model will then be determined (as in [3, Chapter 4]) by the group of all automorphisms of $\mathfrak{A}$ and the normal filter corresponding to finite supports. More explicitly, this means the following.

First, build a set-theoretic universe $V(A)$ by starting at level 0 with the set A of atoms, forming at level 1 all sets of atoms, and in general forming at any level α all sets whose members are at earlier levels. Here α ranges over all the ordinal numbers (but in fact only the first few levels will be relevant to our results). Because *all* sets are added at each stage, $V(A)$ satisfies the axiom of choice, as well as the axioms of ZFA.

Every automorphism π of $\mathfrak{A}$ (indeed, every permutation of A) extends to an automorphism of $V(A)$ by

$$\pi(x) := \{\pi(y) : y \in x\}.$$

A set or atom x is *symmetric* if there is a finite set $E \subseteq A$ such that, whenever π is an automorphism of $\mathfrak{A}$ fixing E pointwise, then $\pi(x) = x$. Then E is

called a *support* of x. (In particular, every atom is symmetric, supported by any finite set of atoms containing it.) Very roughly, "E supports x" means that the set x depends only on the atoms in E, because other atoms can be moved by π without affecting x. The permutation model M, determined by the set A and its given structure, consists of the atoms in A and all those sets x that are *hereditarily symmetric*, i.e., x is symmetric, and so are all its elements, all their elements, etc.

It is well known (see [3]) that M is a model of ZFA. In addition, each automorphism of $\mathfrak{A}$, when extended to $V(A)$ as above, maps M into itself and in fact is an automorphism of M. This means that it preserves the membership relation and therefore anything set-theoretically definable.

When we apply this method of model construction, we shall usually just describe A and its structure and then say something like "let M be the resulting permutation model."

In our consistency proofs, we shall have to verify that certain sets are Dedekind-finite in permutation models. The following proposition, a special case of a well-known characterization of well-orderability in permutation models, will be a useful tool in these verifications.

Proposition 2.1 *Let M be a permutation model as above and X a set in M. Then X is Dedekind-finite in M if and only if, for every finite $E \subset A$, only finitely many members of X are supported by E.*

Proof Suppose first that X is not Dedekind-finite. Let $f : \mathbb{N} \to X$ be a one-to-one map in M and let E be a support of f. Any automorphism π of $\mathfrak{A}$ that fixes E pointwise therefore fixes f. It also fixes every natural number n, because natural numbers are pure sets, involving no atoms. So π fixes every value $f(n)$ of f. (To see this, remember that π is an automorphism of M, so it preserves all set-theoretically definable relations, in particular the ternary relation between function, argument, and value.) But then all the infinitely many values of f are supported by E. This proves the "if" half of the proposition.

For the "only if" half (which we shall not need but include for completeness), suppose some finite $E \subseteq A$ supports infinitely many members of X. In $V(A)$, where the axiom of choice holds, these infinitely many members must include a countable subfamily, which can be enumerated by some one-to-one $f : \mathbb{N} \to X$. Since E supports all the arguments and values of this f, it is

easy to check that it also supports f and that thus $f \in M$. So X is not Dedekind-finite in M. $\square$

We close this section with a brief review of relevant properties of some well-known permutation models.

Example 2.2 [Basic Fraenkel Model] Let $\mathfrak{A}$ be just the countably infinite set A, with no structure. So the automorphisms are all of the permutations of A. The resulting permutation model M is called the basic Fraenkel model [2, 3]. It follows immediately from Proposition 2.1 that A is Dedekind-finite in M, for the only atoms supported by a finite set $E \subseteq A$ are the members of E.

If a partition $\mathcal{P}$ of A is supported by a finite set $E \subseteq A$ then its restriction to $A - E$ is trivial, i.e., either all or no pairs of distinct elements of $A - E$ lie in blocks of $\mathcal{P}$. This follows immediately from the fact that permutations of A fixing E pointwise can send any pair of distinct elements of $A - E$ to any other such pair.

Therefore, the only natural number that divides A in the basic Fraenkel model is 1.

Example 2.3 [Second Fraenkel Model] Let $\mathfrak{A}$ be a countably infinite set A together with a partition into pairs P_n and with an enumeration of those pairs, $n \mapsto P_n$. So the automorphisms must map each pair P_n to itself; an automorphism is completely specified by telling in which P_n's it leaves the two members fixed and in which it interchanges them. The resulting permutation model is called the second Fraenkel model [2, 3].

A finite set $E \subseteq A$ supports only finitely many atoms, namely its members and the atoms paired with them in P_n's. So A is Dedekind-finite in M. (Notice that its quotient obtained by collapsing each pair P_n to a single point is not Dedekind-finite.)

Clearly, 2 divides A in this model, since the partition consisting of the P_n's is supported by the empty set. In fact, every even natural number $2d \neq 0$ divides A, for there are symmetric partitions in which each block consists of d of the P_n's.

No odd natural number $c > 1$ divides A. To see this, suppose we had a partition $\mathcal{Q}$ of A into blocks of size c, and suppose it had a finite support $E \subseteq A$. Enlarging E (to at most twice its original size), we assume that E is a union of some of the P_n's. As it is finite, E meets only finitely many blocks of $\mathcal{Q}$, so consider a block $Q \in \mathcal{Q}$ disjoint from E. As c is odd, find

some P_n such that Q contains exactly one of the two members of P_n, and call that member a. Notice that, since $c > 1$, Q has another member b in some other P_m. Let π be the automorphism that interchanges the two members of P_n while fixing all the other atoms. It fixes E pointwise, so it fixes the partition $\mathcal{Q}$. It also fixes b, so it fixes the block of $\mathcal{Q}$ that contains b, namely Q (because Q is set-theoretically definable from $\mathcal{Q}$ and b). But it moves the element a of Q to an atom outside Q, contrary to the fact that it preserves membership.

Example 2.4 Again let A be partitioned into countably many pairs P_n, but let the structure $\mathfrak{A}$ include only the partition, not the enumeration of the P_n's as in the previous example. Thus, an automorphism of $\mathfrak{A}$ can send one P_n to a different P_m. In this permutation model M, the set of atoms is again Dedekind-finite, and it is obviously divisible by 2 since the partition into the blocks P_n is symmetric. But it is not divisible by any natural number $d > 2$.

To see this, suppose we had a partition $\mathcal{Q}$ of A into pieces of size d, and suppose it had a finite support $E \subseteq A$. As before, we can find some $Q \in \mathcal{Q}$ disjoint from E and from the P_n's meeting E. Since $d > 2$, this Q meets at least two of the P_n's, say P_r and P_s. Also, as Q is finite, we can choose some t such that Q and E don't meet P_t. Now consider an automorphism π of $\mathfrak{A}$ that interchanges the pairs P_s and P_t while fixing all the other atoms. It fixes E pointwise and therefore fixes $\mathcal{Q}$. It fixes the atoms in P_r, one of which is in Q. So it must fix Q. But it moves an atom in Q in P_s to one outside Q in P_t, a contradiction.

Example 2.5 Let d be any natural number ≥ 2. The preceding example can be modified by beginning with a partition of A into sets P_n of cardinality d. The same proof as above shows that, in the resulting permutation model M, the set A of atoms is divisible by d but not by any larger natural number. In fact, it isn't divisible by any smaller natural number either, except of course by 1.

To see this, suppose A had a partition $\mathcal{Q}$ into pieces of size c with $2 \leq c < d$, and suppose $\mathcal{Q}$ were supported by a finite set $E \subseteq A$. As before, we find $Q \in \mathcal{Q}$ disjoint from E and from every P_n meeting E. If Q meets two or more of the P_n's, then we reach a contradiction as in the preceding example. It remains to consider the possibility that Q is entirely included in a single P_m. Since $c \geq 2$, let x and y be distinct elements of Q. Since $c < d$, let z be an element of $P_m - Q$. Let π be the permutation that interchanges y and

z while fixing all other atoms. It is an automorphism of $\mathfrak{A}$ (because y and z are in the same P_m) and it fixes E pointwise. So it fixes the partition $\mathcal{Q}$ and, since it also fixes x, it must fix the block of $\mathcal{Q}$ containing x, namely Q. But it moves an element y of Q to a non-element z of Q, so we again have a contradiction.

3 Congruence Classes Can Overlap

Stephen Hechler asked whether it is possible for a Dedekind-finite set X and $X+1$ to both be divisible by 3. The main results of this section provide an affirmative answer (with or without $X+2$ also being divisible by 3). We shall also present some related results, to put this answer and Hechler's question into perspective. Notice that Dedekind-finiteness is an essential ingredient of the question because, if X were not Dedekind-finite, then there would be a bijection between X and $X+1$ (Proposition 1.4), so if X were divisible by 3 then $X+1$ would be also.

Theorem 3.1 *It is consistent that there exists a Dedekind-finite set X such that $X \pm n$ is divisible by 3 for all natural numbers n.*

Proof It suffices to consider X, $X-1$, and $X-2$, because every integer is congruent modulo 3 to one of $0, -1, -2$, and a partition of a set Y into pieces of size 3 can be trivially modified to give such partitions for $Y \pm 3$.

We use the permutation model M determined by the structure of a complete binary tree on the set A of atoms. In more detail, this means the following. Let the set A of atoms be in one-to-one correspondence (in $V(A)$) with the standard binary tree consisting of all finite sequences of 0's and 1's; write $a(s)$ for the atom corresponding to the sequence s. The structure we use on A is the "immediate successor" relation R; this is the binary relation that holds between $a(s)$ and $a(t)$ if and only if t is $s^\frown\langle 0\rangle$ or $s^\frown\langle 1\rangle$, i.e., if and only if s is obtained from t by deleting its last term.

We need some easily verified observations about the automorphisms of the structure $\langle A, R\rangle$. First, if an automorphism π sends $a(s)$ to $a(t)$, then the sequences s and t have the same length, and π sends $a(s^\frown\langle 0\rangle)$ and $a(s^\frown\langle 1\rangle)$ to $a(t^\frown\langle 0\rangle)$ and $a(t^\frown\langle 1\rangle)$, not necessarily respectively. Second, if the length of s is greater than n, then there is an automorphism that fixes $a(t)$ for all t of length n or smaller but moves $a(s)$. (For example, π could simply

interchange 0 and 1 in the $n+1^{\text{st}}$ position of all the sequences that label the atoms.)

Since the number of binary sequences of length at most n is finite, the second of these observations implies, via Proposition 2.1, that A is Dedekind-finite in M.

We show next that A is divisible by 3 in M, by explicitly exhibiting a symmetric partition into 3-element sets. Each of the sets in the partition is trivially symmetric, being supported by itself. So the partition is hereditarily symmetric and therefore available in M.

We use the partition consisting of the sets $\{a(s), a(s^\frown\langle 0\rangle), a(s^\frown\langle 1\rangle)\}$ for all sequences s of even length (including the empty sequence). In the terminology traditionally used with trees, we can describe the blocks of this partition as each consisting of a node at an even level together with its two (odd level) children. Because of our observation that automorphisms preserve levels, it is easy to check that this partition is symmetric, supported by the empty set. Thus, A is divisible by 3 in M.

If we interchange the roles of even and odd in the preceding paragraph, we obtain a partition of all of A except the root $a(\varnothing)$. This partition is also in M, so we have that $A-1$ is divisible by 3.

Finally, consider the partition of A into the sets $\{a(s), a(s^\frown\langle 0\rangle), a(s^\frown\langle 1\rangle)\}$ where s is either a sequence of even length starting with 0 or a sequence of odd length starting with 1. This partitions all of A except the root and $a(\langle 0\rangle)$. It is symmetric, being supported by $\{a(\langle 0\rangle)\}$. Therefore $A-2$ is also divisible by 3 in M. □

Convention 3.2 The reader will have noticed that the bijection a in the preceding proof, sending sequences s to the corresponding atoms $a(s)$, tends to interfere with easy reading. It would be convenient to identify atoms with the corresponding sequences and thus omit all mention of a. Technically, this would be inaccurate, since sequences are sets and atoms are not. Nevertheless, we shall, from now on, make this identification and analogous identifications for other sets of atoms. If a situation arises where the identification could lead to real confusion, we shall revert to a notation with an explicit bijection, like the a in the preceding proof. In all other situations, we avoid this notational baggage.

Corollary 3.3 *For any $d \geq 3$, it is consistent to have a Dedekind-finite set X such that $X \pm n$ is divisible by d for all natural numbers n.*

Proof Just like the proof of the theorem, with a $(d-1)$-branching tree instead of a binary one. □

Notice that $d \geq 3$ is important in the corollary; the proof doesn't work when $d = 2$ because a 1-branching tree (a single chain) has no nontrivial automorphisms. Not only the proof but the corollary itself fail for $d = 2$.

Theorem 3.4 *If X and $X - 1$ are both divisible by 2, then X is Dedekind-infinite.*

Proof Suppose $\mathcal{P}$ is a partition of X into 2-element sets and $\mathcal{Q}$ is a partition of $X - \{a\}$ into 2-element sets, where a is some element of X. We define the *$\mathcal{P}$-partner* of an element of X to be the other element in the same piece of the partition $\mathcal{P}$; the *$\mathcal{Q}$-partner* of an element of $X - \{a\}$ is defined analogously. Define a function $f : \mathbb{N} \to X$ by

$$f(0) = a \quad \text{and} \quad f(n+1) = \begin{cases} \text{the } \mathcal{P}\text{-partner of } f(n) & \text{if } n \text{ is even} \\ \text{the } \mathcal{Q}\text{-partner of } f(n) & \text{if } n \text{ is odd.} \end{cases}$$

A picture should make it clear that this f is well-defined (i.e., the "n odd" clause never encounters $f(n) = a$ so that the $\mathcal{Q}$-partner would be undefined) and one-to-one. But for the sake of completeness we give a proof.

We show by induction on n that $f(0), \ldots, f(n)$ are well-defined and distinct. This is obvious for $n = 0$. Assume, toward a contradiction, that it holds for a certain n but fails for $n+1$. What fails cannot be well-definedness, i.e., we cannot have n odd and $f(n) = a$, for the latter can happen only if $n = 0$ (by the assumed distinctness up to n) and 0 isn't odd. So what fails is the distinctness; $f(n+1) = f(k)$ for some $k \leq n$. In fact $k < n$, because $f(n)$ and $f(n+1)$ are partners (under $\mathcal{P}$ or under $\mathcal{Q}$) and therefore distinct.

Suppose temporarily that $k = 0$. Thus, $f(n+1) = f(0) = a$, which is not a $\mathcal{Q}$-partner of anything. If n were odd, then $f(n+1)$ would be the $\mathcal{Q}$-partner of $f(n)$. So we know that n must be even and so $a = f(n+1)$ is the $\mathcal{P}$-partner of $f(n)$. Therefore $f(n)$ is the $\mathcal{P}$-partner of a, namely $f(1)$. Since $f(0), \ldots, f(n)$ are all distinct yet $f(n) = f(1)$, we must have $n \leq 1$. But n is even and $k < n$, so this is impossible.

Our temporary supposition that $k = 0$ has led to a contradiction. So we know that $k = l + 1$ for some l. Let us write $\mathcal{R}$ to mean $\mathcal{P}$ if n is even and $\mathcal{Q}$ if n is odd; so $f(n+1)$ and $f(n)$ are $\mathcal{R}$-partners. If l had the same parity as n, then $f(l+1) = f(k)$ and $f(l)$ would also be $\mathcal{R}$-partners. Since

$f(n+1) = f(k)$, this would mean that $f(n) = f(l)$. But $l < n$, so this contradicts the induction hypothesis that $f(0), \ldots, f(n)$ are distinct. So l must have the opposite parity from n. Then $k = l+1$ has the same parity as n, and so $f(k)$ and $f(k+1)$ are (like $f(n)$ and $f(n+1)$) $\mathcal{R}$-partners. From $f(k) = f(n+1)$ it now follows that $f(n) = f(k+1)$. Applying again the fact that $f(0), \ldots, f(n)$ are distinct, we conclude that $k+1 \geq n$. But this is impossible, since $k < n$ and they have the same parity. This contradiction completes the proof that f is well-defined and one-to-one, and therefore X is Dedekind-infinite. □

This result explains, of course, why Hechler chose 3 as the divisor in his question. The following result explains why he asked about divisibility rather than strong divisibility.

Theorem 3.5 *If X and $X-1$ are both strongly divisible by 3, then X is Dedekind-infinite.*

We defer the proof to Section 6 (see Example 6.16), where we shall obtain this result as a corollary of a more general one. The following result, whose proof we also defer (to Example 6.14) for the same reason, fills in a gap between Theorems 3.1 and 3.5

Theorem 3.6 *It is consistent to have a Dedekind-finite set X such that X is strongly divisible by 3 and $X-1$ and $X-2$ are divisible by 3.*

Having seen in Corollary 3.3 that a Dedekind-finite set can belong simultaneously to all congruence classes modulo a prescribed $d \geq 3$, it is natural to ask about Dedekind-finite sets belonging to just some (or none) of the congruence classes. Although we do not have a complete answer, the following two examples take care of some easy cases, and the subsequent theorem and its corollaries handle some more difficult situations.

Example 3.7 The set A of atoms in the basic Fraenkel model, described in Example 2.2, is not congruent to any integer modulo any integer $d \geq 2$. Indeed, if any finitely many elements are added to or removed from A, any partition of the resulting set must contain either an infinite piece or infinitely many one-element pieces, by the argument given in Example 2.2.

Example 3.8 We saw in Example 2.5 that the set A of atoms in the model defined there is a Dedekind-finite set divisible by d but no other divisor ≥ 2. The proof given there can be adapted easily to show that A is not congruent modulo d to any of $1, 2, \ldots, d-1$ and is not congruent to any finite number modulo any $d' \geq 2$ other than d. Indeed, with the notation of Example 2.5, we see that if $\mathcal{Q} \in M$ is a partition of A plus or minus a finite set, into finite pieces that are not singletons, then all but finitely many of the pieces in $\mathcal{Q}$ must coincide with pieces P_n of the partition used in defining the model. This immediately implies that there are no congruences modulo divisors other than d. And by considering the finitely many pieces of $\mathcal{Q}$ and the finitely many P_n's that don't coincide, we easily find that the number of points added to or removed from A must be a multiple of d.

It is trivial to modify this example to get an infinite, Dedekind-finite set that belongs to any one prescribed congruence class modulo d. Just add an appropriate finite set to the set A considered above. (If we omitted "infinite" from the last sentence, then of course there would also be finite examples.)

Theorem 3.9 *For any integer $d \geq 3$, it is consistent to have a Dedekind-finite set X such that X and $X+1$ are divisible by d but $X+k$ is not divisible by d for any i in the range $2 \leq i < d$.*

Proof Let $d \geq 3$ be given, and consider the permutation model M built from the following structure $\mathfrak{A}$ of atoms. Begin with a countably infinite set A of atoms, partitioned into pieces P_n ($n \in \mathbb{N}$) of size $d-1$. The desired structure on A is this partition together with the enumeration $n \mapsto P_n$ of its pieces. Thus, this model is related to the second Fraenkel model (Example 2.3) and $d-1$ as Example 2.5 is to Example 2.4 and d. We construct a set X in the model such that X and $X+1$ are divisible by d and $X+i$ is not divisible by d for $2 \leq i < d$.

For $n \in \mathbb{N}$, let $X_n = \{\{x\} : x \in P_n\} \cup \{\{x, y\} : x \in P_n \text{ and } y \in P_{n+1}\}$ and let $X = \bigcup_{n \in \mathbb{N}} X_n$. It is easy to see, using Proposition 2.1, that X is Dedekind-finite. Indeed, if E is a finite subset of A, then $E \subseteq \bigcup_{n<k} P_n$ for some $k \in \mathbb{N}$, and every element of X supported by E is a subset of this union. So E supports only finitely many elements of X.

X is divisible by d since the sets $\{\{x\}\} \cup \{\{x, y\} : y \in P_{n+1}\}$ for $n \in \mathbb{N}$ and $x \in P_n$ form a partition of X into d-element sets. $X-(d-1)$ is divisible by d (and therefore so is $X+1$) because the sets $\{\{x\}\} \cup \{\{x, y\} : y \in P_{n-1}\}$

for $n \in \mathbb{N} - \{0\}$ and $x \in P_n$ form a partition of $X - \{\{x\} : x \in P_0\}$ into d-element sets.

We now prove by contradiction that (in the model) the sets $X - r$ for $1 \leq r < d - 1$ (and hence also $X + i$ for $2 \leq i < d$) are not divisible by d. Assume that for such an r, $X - r$ is divisible by d and that $\mathcal{Q}$ is a partition of $X - R$ into d-element sets where R is an r-element subset of $\{\{x\} : x \in P_0\}$. Let $\bigcup_{n<k} P_n$ be a support of $\mathcal{Q}$. For the remainder of the proof we shall denote by G the group of automorphisms of $\mathfrak{A}$ which fix this support pointwise. Choose a natural number m so large that for $j \geq m$ if $z \in X_j$ then the $\mathcal{Q}$-block containing z does not meet $\bigcup_{n<k} X_n$ and hence

$$\text{for every } z' \text{ in the } \mathcal{Q} \text{ block containing } z, z' \cap (\bigcup_{n<k} P_n) = \varnothing. \tag{1}$$

We will say that two elements z and z' of X are adjacent if for some natural number n, z and z' both meet P_n. Note that if z and z' are adjacent with $z \in X_n$ and $z' \in X_{n'}$ then $|n' - n| \leq 1$. We also note that if z and z' are two non-adjacent elements of X neither of which meets $\bigcup_{n<k} P_n$, then there are permutations ϕ_i and ψ_i, $1 \leq i \leq d-2$, in G such that $\phi_i(z) = z$, $\psi_i(z') = z'$ and the elements $z, z', \phi_1(z'), \phi_2(z'), \ldots, \phi_{d-2}(z'), \psi_1(z), \psi_2(z), \ldots, \psi_{d-2}(z)$ are distinct. (Assuming that $z' \cap P_n = \{a_0\}$ where $P_n = \{a_0, a_1, \ldots, a_{d-2}\}$ we can take take ϕ_i to be the permutation that interchanges a_0 and a_i. A similar construction can be used to obtain the ψ_i's.) It follows that two such non-adjacent elements cannot be in the same $\mathcal{Q}$-block. For if it were the case that z and z' were in the same $\mathcal{Q}$-block Q, the condition $\phi_i(z) = z$ would imply that $\phi_i(Q) = Q$. Similarly, $\psi_i(Q) = Q$. Therefore Q would have to contain the $2d-2$ elements $z, z', \phi_1(z'), \phi_2(z'), \ldots, \phi_{d-2}(z'), \psi_1(z), \psi_2(z), \ldots, \psi_{d-2}(z)$. This is a contradiction since the $\mathcal{Q}$-blocks have cardinality d and $d > 2$. We summarize the result of the discussion above with:

Lemma 3.10 *For all z and z' in X, if z and z' are disjoint from $\bigcup_{n<k} P_n$ and z and z' are in the same $\mathcal{Q}$ block, then z and z' are adjacent.*

Next we argue that

Lemma 3.11 *Either*

1. *For all $x \in P_m$, the $\mathcal{Q}$-block containing $\{x\}$ is $\{\{x\}\} \cup \{\{t, y\} : y \in P_{m+1}\}$ for some $t \in P_m$ or*

2. *For all $x \in P_m$, the $\mathcal{Q}$-block containing $\{x\}$ is $\{\{x\}\} \cup \{\{t,y\} : y \in P_{m-1}\}$ for some $t \in P_m$.*

Proof Assume $x \in P_m$; by (1) every z in the $\mathcal{Q}$-block containing $\{x\}$ is disjoint from $\bigcup_{n<k} P_n$. Therefore by 3.10 $\{x\}$ and z are adjacent. This implies that z contains some element of P_m, that is z is in $\{\{t\} : t \in P_m\} \cup \{\{t,y\} : t \in P_m \text{ and } y \in P_{m+1}\} \cup \{\{t,y\} : t \in P_m \text{ and } y \in P_{m-1}\}$. Since $\{\{t\} : t \in P_m\}$ has only $d-1$ elements and each $\mathcal{Q}$-block has d elements, the $\mathcal{Q}$-block containing $\{x\}$ must contain either an element of $\{\{t,y\} : t \in P_m \text{ and } y \in P_{m+1}\}$ or an element of $\{\{t,y\} : t \in P_m \text{ and } y \in P_{m-1}\}$. We will argue that in the first case alternative 1 of the lemma holds. The argument that alternative 2 holds in the second case is similar.

Assume that $\{t,y\}$ is in the $\mathcal{Q}$-block Q containing $\{x\}$ where $t \in P_m$ and $y \in P_{m+1}$. For each $y' \in P_{m+1}$, the permutation ϕ interchanging y and y' is in G and fixes $\{x\}$. It therefore fixes Q and it follows (since ϕ fixes t) that $\phi(\{t,y\}) = \{t,y'\} \in Q$. The $\mathcal{Q}$ block Q therefore contains $\{\{x\}\} \cup \{\{t,y'\} : y' \in P_{m+1}\}$. Since Q has d elements it follows that $Q = \{\{x\}\} \cup \{\{t,y'\} : y' \in P_{m+1}\}$. Although it is not needed for the proof of the theorem, we note that for $d > 3$, it must be the case that $t = x$.

Now for any $x' \in P_m$, the permutation ϕ interchanging x and x' interchanges the $\mathcal{Q}$-block containing x and the $\mathcal{Q}$-block containing x'. Therefore the $\mathcal{Q}$-block containing x' is $\{\{\phi(x)\}\} \cup \{\{\phi(t),y'\} : y' \in P_{m+1}\} = \{\{x'\}\} \cup \{\{\phi(t),y'\} : y' \in P_{m+1}\}$. Alternative 1 follows since $\phi(t) \in P_m$. $\square$

To complete the proof of the theorem we shall show that either of the two possibilities given in lemma 3.11 leads to a contradiction. Suppose first that alternative 1 of 3.11 holds. From this assumption it follows that X_m is partitioned into d-element sets by $\mathcal{Q}$-blocks and therefore the $\mathcal{Q}$ block containing z for any $z \in \bigcup_{j>m} X_j$ is disjoint from X_m. It is also the case that the $\mathcal{Q}$-block containing such a z is disjoint from X_i for $i < m$ for if z' is in this $\mathcal{Q}$-block then by (1) z' is disjoint from $\bigcup_{n<k} P_n$. By lemma 3.10 z and z' are adjacent. Since z' is not in X_m it must be in some X_i where $i > m$.

We have thus shown no $\mathcal{Q}$-block can meet both $\bigcup_{n\leq m} X_n - R$ and $\bigcup_{n>m} X_n$, and so we conclude that $\bigcup_{n\leq m} X_n - R$ is partitioned into d-element sets by $\mathcal{Q}$-blocks and is therefore divisible by d. On the other hand $\bigcup_{n\leq m} X_n$ is a finite set divisible by d since the sets $\{\{x\}\} \cup \{\{x,y\} : y \in P_{n+1}\}$ for $n \leq m$ and $x \in P_n$ form a partition of it into sets of cardinality d. But this means that $\bigcup_{n\leq m} X_n - R$ is not divisible by d since $1 \leq |R| < d-1$.

We arrive at a similar contradiction in the case that alternative 2 of lemma 3.11 holds by considering the set $Z = \left(\bigcup_{n<m} X_n \cup \{\{x\} : x \in P_m\right) - R$. An argument similar to the one in the previous paragraph shows that Z is partitioned into d-element sets by $\mathcal{Q}$-blocks. On the other hand $W = \left(\bigcup_{n<m} X_n \cup \{\{x\} : x \in P_m\}\right) - \{\{x\} : x \in P_0\}$ is partitioned into sets of cardinality d by the sets $\{\{x\}\} \cup \{\{x, y\} : y \in P_{n-1}\}$ for $0 < n \leq m$ and $x \in P_n$. This is a contradiction since the finite sets Z and W differ by $\{\{x\} : x \in P_0\} - R$, a set whose cardinality is strictly between 1 and d. $\square$

Corollary 3.12 *Let $d \geq 3$ and $1 \leq r < d-1$ be given. It is consistent to have a Dedekind-finite set X such that X, $X+1$, ... , $X+r$ are divisible by d but $X+r+1$, ... , $X+d-1$ are not.*

Proof We modify the proof of the preceding theorem by using r "layers" of atoms, each of which looks like the A of the theorem. That is, we begin with a countable set A of atoms, partitioned first into r countably infinite layers A_ξ ($\xi < r$), with each layer A_ξ partitioned into $(d-1)$-element pieces $P_{\xi,n}$. Let $\mathfrak{A}$ be the structure consisting of the set A with these partitions and the labeling $(\xi, n) \mapsto P_{\xi,n}$. Let M be the associated permutation model. In this model, let

$$\begin{aligned} X_{\xi,n} &= \{\{x\} : x \in P_{\xi,n}\} \cup \{\{x, y\} : x \in P_{\xi,n} \text{ and } y \in P_{\xi,n+1}\} \\ X_\xi &= \bigcup_{n \in \mathbb{N}} X_{\xi,n} \\ X &= \bigcup_{\xi<r} X_\xi. \end{aligned}$$

As in the proof of the theorem, the model contains, for each ξ, a partition of X_ξ into pieces of size d (namely $\{\{x\}\} \cup \{\{x, y\} : y \in P_{\xi,n+1}\}$ for $n \in \mathbb{N}$ and $x \in P_{\xi,n}$) and a partition of $X_\xi - \{\{x\} : x \in P_{\xi,0}\}$ into pieces of size d (namely $\{\{x\}\} \cup \{\{x, y\} : y \in P_{\xi,n-1}\}$ for $n \in \mathbb{N} - \{0\}$ and $x \in P_{\xi,n}$). By using one of these partitions for some levels ξ and the other for the rest, we can assemble a partition, into pieces of size d, for any set obtained from X by removing all or some or none of the $(d-1)$-element sets $\{\{x\} : x \in P_{\xi,0}\}$. These partitions show that $X - q(d-1)$ and therefore $X + q$ are divisible by d for $0 \leq q \leq r$. It remains to verify the non-divisibility assertion of the corollary.

Suppose this non-divisibility assertion failed. So we would have a partition $\mathcal{Q}$, into blocks of size d, for a set of the form $X - R$, where R is a finite set of cardinality not congruent modulo d to any $q(d-1)$ (equivalently to $-q$) for $0 \leq q \leq r$. Reducing R if necessary by a multiple of d elements, we can arrange that $R \subseteq \{\{x\} : x \in P_{0,0}\}$. As in the proof of the theorem, let E support $\mathcal{Q}$, let k be large enough so that $E \subseteq \bigcup_\xi \bigcup_{n<k} P_{\xi,n}$, and let m be so large that no block of $\mathcal{Q}$ has an element meeting a set of the form $P_{\xi,n}$ with $n \leq k$ and an element meeting a set of the form $P_{\xi,n}$ with $n \geq m$.

If all the members (in A) of members (in X) of a block of $\mathcal{Q}$ are in pieces $P_{\xi,n}$ with $n < p$, then we say that this block of $\mathcal{Q}$ *lives below* p. To *live above* p is defined analogously. Thus, for example, our choice of m ensures that every block of $\mathcal{Q}$ lives either below m or above k.

A block $Q \in \mathcal{Q}$ that has elements meeting different levels A_ξ cannot live above k. Otherwise, by considering permutations that fix E pointwise, fix one of the relevant levels pointwise, but move elements of Q from another level, we would find at least $2(d-1)$ members in Q, contrary to $|Q| = d \geq 3$. In particular, the blocks in $\mathcal{Q}$ that don't live below m must lie entirely in one level.

Consider, for each level ξ, the analogs of the $\bigcup_{n \leq m} X_n - R$ and Z from the proof of the theorem, namely

$$Y_\xi = \{\{x\} : x \in P_{\xi,n}, n < m\} - R \cup \{\{x,y\} : x \in P_{\xi,n}, y \in P_{\xi,n+1} \text{ and } n < m\}$$

and

$$\begin{aligned} Z_\xi &= Y_\xi \cup \{\{x\} : x \in P_{\xi,m}\} \\ &= \{\{x\} : x \in P_{\xi,n}, n \leq m\} - R \cup \\ &\qquad \{\{x,y\} : x \in P_{\xi,n}, y \in P_{\xi,n+1} \text{ and } n < m\} \end{aligned}$$

Notice that each Y_ξ would be divisible by d if we hadn't removed R; therefore the actual cardinality of Y_ξ is congruent modulo d to $-|R \cap X_\xi|$. Similarly, each Z_ξ would be congruent to $d-1$ modulo d if we hadn't removed R, so its actual cardinality is congruent modulo d to $-1-|R \cap X_\xi|$. (It is true but unimportant that our normalization of R makes $R \cap X_\xi = \varnothing$ for all $\xi \neq 0$; the important aspect of the normalization is that it ensures $R \subseteq \{\{x\} : x \in P_{\xi,n}, n < m\}$, since $m > 0$.)

Define a level ξ to be Y-nice if every block of $\mathcal{Q}$ that meets Y_ξ is included in Y_ξ; define Z-nice analogously. The proof of the theorem shows that every

level must be Y-nice or Z-nice. Let q be the number of Z-nice levels; of course we have $0 \leq q \leq r$. Consider the union of the Z_ξ's from the Z-nice levels and the Y_ξ's from the remaining levels, which are Y-nice. The cardinality of this union is, by the calculations in the preceding paragraph, congruent modulo d to $-q-|R|$. On the other hand, by definition of "nice," this union is a union of blocks of $\mathcal{Q}$, so its cardinality is divisible by d. Therefore, $|R|$ is congruent modulo d to $-q$ and thus to $q(d-1)$, contrary to our choice of R. □

Corollary 3.13 *Let C be any set of consecutive congruence classes modulo $d \geq 3$. It is consistent to have a Dedekind-finite set X such that $X + k$ is divisible by d if and only if $k \in C$.*

Proof Being a set of consecutive congruence classes, C can be obtained from a set of the form $\{0, 1, \ldots, r\}$ (or $\varnothing$) by adding an appropriate element c modulo d. We have an example of a Dedekind-finite X such that $X + k$ is divisible by d if and only if $k \in \{0, 1, \ldots, r\}$; this comes from the preceding corollary if $r \geq 1$, from Example 3.8 if $r = 0$, and from Example 3.7 if $C = \varnothing$. Removing c elements from X, we get the set required for the given C. □

We do not know whether the last corollary remains true without the hypothesis "consecutive."

4 Powers

The preceding section showed that a Dedekind-finite cardinal can behave quite differently from a finite one, for example by being congruent modulo 3 to all of 0, 1, and 2. The main result of the present section shows that some other aspects of Dedekind-finite divisibility work exactly as for finite cardinals.

Theorem 4.1 *If 2^X is Dedekind-finite, then it is not divisible by 3.*

This theorem is a special case of the following.

Theorem 4.2 *If c and d are finite numbers, if c^X is Dedekind finite, and if c^X is divisible by d, then some finite power of c is divisible by d.*

The conclusion of the theorem can be equivalently formulated as: Every prime divisor of d also divides c. In particular, if d is prime (or a product of distinct primes) then it must divide c.

Proof The theorem is trivial if c is 0 or 1, so we assume that $c \geq 2$.

Assume that d divides c^X, and fix a partition $\mathcal{P}$ of c^X into pieces of size d. Assume further that d divides no finite power of c. We shall produce a countable sequence of distinct subsets of X, thereby showing that 2^X is not Dedekind finite. Since 2^X is a subset of c^X, it will follow immediately that c^X is not Dedekind-finite.

We shall be working with certain partitions Π of X (not to be confused with the fixed partition $\mathcal{P}$ of c^X). We write $C(\Pi)$ for the set of those functions from X to c (i.e., members of c^X) that are constant on all the pieces of the partition Π. Each Π that we consider will have only finitely many pieces, so $C(\Pi)$ will be finite. More precisely, if Π has n pieces then the cardinality of $C(\Pi)$ is exactly c^n, which is, by hypothesis, not divisible by d. Therefore, there will be members f of c^X that are not in $C(\Pi)$ but are in the same block of $\mathcal{P}$ as some member of $C(\Pi)$. Since there are only finitely many such f's (at most $(d-1)c^n$ of them), we can refine Π to a partition Π', still with finitely many pieces, such that each of these f's is in $C(\Pi')$. For definiteness, we take the coarsest possible Π'; two elements x and y of X are in the same piece of Π' if and only if $f(x) = f(y)$ for every function f that is in the same $\mathcal{P}$-block with a member of $C(\Pi)$ (including functions that are themselves in $C(\Pi)$). The crucial fact about this construction is that Π' properly refines Π.

For the rest of the proof, we use only the operation $\Pi \mapsto \Pi'$ defined in the preceding paragraph and the fact that it sends every partition of Π with finitely many pieces to another such partition properly refining the first. Starting with the partition Π_0 consisting of the single piece X, we iterate this operation to produce a sequence of partitions $\Pi_{n+1} = \Pi_n{}'$.

Temporarily fix some element $x \in X$, and let $B_n(x)$ be the piece of partition Π_n that contains x. Thus

$$X = B_0(x) \supseteq B_1(x) \supseteq B_2(x) \supseteq \ldots.$$

If infinitely many of the inclusions in this sequence are proper, then the distinct $B_n(x)$'s form a countable sequence of subsets of X. So in this case the proof is complete.

Now un-fix x. If there is even a single $x \in X$ for which the $B_n(x)$ sequence has infinitely many proper inclusions, then the proof is complete.

So we assume henceforth that, for every x, there is an n such that the $B_k(x)$ sequence has no proper inclusions beyond the n^{th} term. Define $s(x)$ to be the smallest such n.

If the function $s : X \to \mathbb{N}$ takes infinitely many distinct values, then the nonempty sets of the form $s^{-1}(\{n\})$ are countably many (enumerate them in order of increasing n) distinct subsets of X. So again the proof is complete in this case.

There remains only the case that the function s takes only finitely many values. In this case, let m be an upper bound for all these values, and notice that $B_m(x) = B_{m+1}(x)$ for all x. But this means that $\Pi_{m+1} = \Pi_m$, which contradicts the fact that Π' always properly refines Π. This contradiction shows that the present case cannot arise, so the proof is complete. □

In the theorem just proved, the assumption that c^X is Dedekind-finite cannot be replaced by the weaker assumption that X is Dedekind-finite.

Theorem 4.3 *It is consistent to have a Dedekind-finite X such that 2^X is divisible by all positive integers.*

Proof The model we use is the second Fraenkel model from Example 2.3 above; we use the same notation here as in that example. We saw there that the set A of atoms is Dedekind-finite in this model M. So it remains to exhibit a partition of the power set of A into sets of cardinality d.

First, we need a convenient description of subsets S of A in M. In the first place, any such S determines three subsets $S_0, S_1, S_2 \subseteq \mathbb{N}$ by

$$S_i := \{n \in \mathbb{N} : |S \cap P_n| = i\}.$$

Then S is the union of all the P_n's for $n \in S_2$ plus one element from each P_n for $n \in S_1$. We write S^* for the set of these single elements for $n \in S_1$, i.e.,

$$S^* = S \cap \bigcup_{n \in S_1} P_n = S - \bigcup_{n \in S_2} P_n.$$

Notice that (S_0, S_1, S_2) and S^* completely determine S.

If E supports S, then it must meet P_n for every $n \in S_1$, for otherwise the permutation that interchanges the members of P_n (and fixes all other atoms) would fix E pointwise, would therefore fix S, but would move the atom in $S \cap P_n$ to one outside S, a contradiction. Since every $S \in M$ has a finite support, we conclude that S_1 is always finite.

Now let an arbitrary positive integer d be given; we seek a partition (in M) of $\mathcal{P}(A)$ into pieces of size d. Consider first the collection of all triples (S_0, S_1, S_2) as above, i.e., the collection of all ordered, 3-piece partitions of $\mathbb{N}$ in which the middle piece is finite. Let Π be a partition of this collection of triples into sets of size d with the additional property that the d triples in any block of Π all have the same middle component. It is trivial to find such a partition in $V(A)$, where the axiom of choice holds; just partition the set of triples first into blocks according to their middle components, note that each of these blocks is infinite (of the cardinality of the continuum), and so partition each of these blocks into subblocks of size d. But, although it was found in $V(A)$, our partition Π is in fact in M. Indeed, it is hereditarily supported by $\varnothing$ because it is a set of sets of ordered triples of sets of natural numbers, so no atoms are involved in it.

Notice also that the functions sending any $S \subseteq A$ to (S_0, S_1, S_2) and to S^* are in M because they are invariant under all automorphisms of $\mathfrak{A}$ and thus supported by $\varnothing$. Therefore, the following equivalence relation on $\mathcal{P}(A)$ is also in M.

$$S \sim T \iff (S_0, S_1, S_2) \text{ and } (T_0, T_1, T_2) \text{ lie in the same } \Pi\text{-block and } S^* = T^*.$$

Since each equivalence class with respect to $\sim$ has size d, the proof is complete. □

5 Prescribed Divisors

What implications are there between statements of the form "n divides X" for different n and the same X? When X is finite, there are many obvious implications of this sort. If n divides X then so do all divisors of n. If two relatively prime numbers divide X then so does their product. We have already seen that implications of the first sort need not be correct when X is merely Dedekind-finite. Example 2.5 gives Dedekind-finite sets divisible by any n and none of its non-trivial divisors. In this section, we shall show that for Dedekind-finite X there are no necessary connections between different divisors; anything can happen.

Theorem 5.1 *Let $D \subseteq \mathbb{N} - \{0, 1\}$. It is consistent that there is a Dedekind-finite set divisible by all members of D and by no other natural numbers ≥ 2.*

Proof If D is empty or consists of a single number, then Example 2.2 or 2.5 is as required. So we assume from now on that D has at least two elements.

Given D, let A_* be the free product of cyclic groups of all orders in D,

$$A_* := \mathbb{Z}/d_1 * \mathbb{Z}/d_2 * \cdots$$

where $D = \{d_1, d_2, \dots\}$. Recall that this group can be constructed as follows. First, form all *words*, i.e., finite sequences whose terms, also called *letters* in this context, are non-identity elements of the groups $\mathbb{Z}/d$. (We assume these groups are disjoint; for the traditional definition of $\mathbb{Z}/d$ as a quotient of $\mathbb{Z}$ this will automatically be true.) Call a word *reduced* if no two adjacent letters in it come from the same $\mathbb{Z}/d$. Then A_* is the set of reduced words, with multiplication defined by concatenation and simplification. That is, to multiply two reduced words, we write one after the other and reduce the result by performing the following operation as often as needed to produce a reduced word. ("As often as needed" might be not at all, if the concatenation is already reduced.) If two adjacent letters come from the same group $\mathbb{Z}/d$ then replace them with their product in that group provided this product is not the identity, and simply delete them if the product is the identity. This multiplication operation on A_* is associative, the empty word serves as an identity element, and the inverse of a reduced word is obtained by replacing each letter by its inverse (in the group $\mathbb{Z}/d$ that it came from) and reversing the order of the letters. So A_* is a group.

For each $d \in D$, the words consisting of one letter from $\mathbb{Z}/d$ form, together with the identity, an isomorphic copy of $\mathbb{Z}/d$ in A_*. We shall simplify notation by identifying $\mathbb{Z}/d$ with this copy. Let $\mathcal{P}_{d*}$ be the partition of A_* into the (left) cosets $x(\mathbb{Z}/d)$ of this subgroup $\mathbb{Z}/d$. Thus, two reduced words x and y are in the same block of $\mathcal{P}_{d*}$, and we call them *d-adjacent* or *d-neighbors*, if and only if $x^{-1}y \in \mathbb{Z}/d$, if and only if either

- one is obtained from the other by appending a single letter from $\mathbb{Z}/d$ or
- they both end with letters from $\mathbb{Z}/d$ and are identical except for this last letter.

We say that x and y are *neighbors* if they are d-neighbors for some $d \in D$. Thus, the neighbors of a reduced word x are formed by either deleting the last letter, appending one additional letter, or changing the last letter to another

one from the same subgroup $\mathbb{Z}/d$. In each case, if the deleted, added, or altered letter is from $\mathbb{Z}/d$, then we have a d-neighbor. We shall refer to this description of the partitions and the adjacency relations as the "right end" description, since it refers only to the right ends of the reduced words.

Notice that A_* is infinite (because $|D| \geq 2$) and that each block of $\mathcal{P}_{d*}$ has cardinality d.

Notice also that A_*, considered as a graph with the adjacency relation, is connected; we can go from any word to any other by first deleting all the letters of the first word one at a time and then appending the letters of the second word. Of course, if the two words have a (nonempty) common initial segment, then we can be more efficient, stopping the deletions and starting the additions when we reach that common segment. If it happens that the next letter after the longest common initial segment (empty or not) comes from the same $\mathbb{Z}/d$ in both words, then we can be slightly more efficient yet, for instead of deleting this letter from the first word and then adding the corresponding letter from the second word, we can change the one letter to the other in a single adjacency step. It is easy to check that no further efficiency is possible. This discussion can be succinctly described as follows.

Lemma 5.2 *Consider two reduced words $x = wx'$ and $y = wy'$ whose longest common initial segment is w. If x' and y' begin with letters from the same $\mathbb{Z}/d$, then the distance in the adjacency graph between x and y is length(x') + length$(y') - 1$. Otherwise, the distance is length(x') + length(y'). In either case, this distance is the length of the reduced form of $x^{-1}y$.*

Proof The preceding discussion proves all but the last sentence. So consider what happens when we form $x^{-1}y = x'^{-1}w^{-1}wy'$. First, the w parts cancel. Then the first letter q of y' is next to the inverse p^{-1} of the first letter p of x'. If p and q come from different $\mathbb{Z}/d$'s, then we already have a reduced word and its length is length(x') + length(y'). If p and q come from the same $\mathbb{Z}/d$, then $p^{-1}q$ is replaced by a single letter from that $\mathbb{Z}/d$ — not deleted, because $p \neq q$. (If p and q were equal, they would have been part of w rather than of x' and y'.) This decreases the length by 1 and leaves us a reduced word whose length is length(x') + length$(y') - 1$. □

Corollary 5.3 *Suppose $x_0, x_1, \ldots, x_l$ are reduced words, with each consecutive pair x_{i-1}, x_i d_i-adjacent. Suppose further that $d_i \neq d_{i+1}$ for all i. Then the distance from x_0 to x_l in the adjacency graph is l.*

Proof For each i we have $x_i = x_{i-1}p_i$ for some letter $p_i \in \mathbb{Z}/d_i$. Therefore $x_l = x_0 p_1 p_2 \ldots p_l$, and $x_0^{-1} x_l = p_1 p_2 \ldots p_l$. This product of p_i's is a reduced word because consecutive d_i's are distinct. By the lemma, the distance between x_0 and x_l is the length of this word, namely l. □

Define $A = \mathbb{N} \times A_*$, the disjoint union of countably many copies of A_*, and for each $d \in D$ define $\mathcal{P}_d$ to be the partition of A obtained by copying $\mathcal{P}_{d*}$ in each component. That is, a block of $\mathcal{P}_d$ has the form $\{i\} \times B$ where $i \in \mathbb{N}$ and B is a block of $\mathcal{P}_{d*}$. Let $\mathfrak{A}$ be the structure given by A with all these partitions. Let M be the resulting permutation model. (See Convention 3.2 for how the elements of A became atoms.) The partitions $\mathcal{P}_d$ are symmetric and witness that every $d \in D$ divides A in M. We shall show that, in M, A is Dedekind-finite and not divisible by any natural number ≥ 2 that is not in D; this will complete the proof of the theorem. We use the terminology "d-adjacent" and "adjacent" in A to refer to these concepts within each copy of A_* separately. These copies are thus the components of the graph given by the adjacency relation.

We shall need some information about automorphisms of $\mathfrak{A}$, particularly that there are enough automorphisms. To establish this information, we first consider automorphisms of $\mathfrak{A}_*$, the structure given by A_* and the partitions $\mathcal{P}_{d*}$. Notice that the group structure of A_*, though used in defining these partitions, is not part of the structure $\mathfrak{A}_*$; it need not be preserved by the automorphisms.

Lemma 5.4 *Any member of A_* can be mapped to any other member of A_* by some automorphism of $\mathfrak{A}_*$.*

Proof If a is any element of A_* then the operation of left multiplication (in the group A_*), $x \mapsto ax$ maps left cosets of any subgroup to left cosets of the same subgroup. So it is an automorphism of $\mathfrak{A}_*$. To send any specified x to any specified y by such an automorphism, just use $a = yx^{-1}$. □

Lemma 5.5 *Let $d \in D$ and let σ be any permutation of the nonidentity elements of $\mathbb{Z}/d$. There is an automorphism of $\mathfrak{A}_*$ that fixes the identity element and agrees with σ on the remaining elements of the subgroup $\mathbb{Z}/d$.*

Proof The desired automorphism is defined on reduced words w as follows. If the first letter in w is not from $\mathbb{Z}/d$, or if $w = 1$ so there is no first letter, then map w to itself. If the first letter is in $\mathbb{Z}/d$, then apply σ to this

(occurrence of this) letter and leave the remaining letters unchanged. If we recall the right end descriptions of the partitions $\mathcal{P}_{e*}$, then it becomes easy to check that the map we defined preserves the partitions and is therefore an automorphism of $\mathfrak{A}_*$. □

Lemma 5.6 *Let $d \in D$, let $a \in A_*$, and let σ be any permutation of the elements other than a in the $\mathcal{P}_{d*}$-block that contains a. Then there is an automorphism of $\mathfrak{A}_*$ that fixes a and agrees with σ on the rest of its $\mathcal{P}_{d*}$-block.*

Proof Use Lemma 5.4 to reduce the problem to the case $a = 1$, which is covered by Lemma 5.5. In more detail, let π be an automorphism of $\mathfrak{A}_*$ sending a to 1. It sends the rest of the $\mathcal{P}_{d*}$-block containing a to the rest of the block containing 1, namely the set N of non-identity elements of $\mathbb{Z}/d$. Therefore, $\pi\sigma\pi^{-1}$ is a permutation of N. By Lemma 5.5, let τ be an automorphism of $\mathfrak{A}_*$ that fixes 1 and agrees with $\pi\sigma\pi^{-1}$ on N. Then $\pi^{-1}\tau\pi$ fixes a and agrees with σ on the rest of its $\mathcal{P}_{d*}$-block. □

Lemma 5.7 *Let n be a natural number and w a reduced word of length at least $n+2$. Then there is an automorphism of $\mathfrak{A}_*$ that moves w but fixes all reduced words of length $\leq n$.*

Proof Let the length of w be $l+2$, where $l \geq n$, and let $w = w'pq$ where w' is a word of length l and p and q are letters. Suppose first that p comes from $\mathbb{Z}/d$ with $d \neq 2$. Let p' be a different non-identity element of the same $\mathbb{Z}/d$. Then the desired automorphism can be defined as follows. Given an arbitrary word, if it has p or p' as its $l+1^{\text{st}}$ letter, then change it to p' or p, respectively. Otherwise, leave w unchanged. The right end description of the partitions makes it clear that this is an automorphism, and it clearly moves w (to $w'p'q$) while fixing all reduced words of length at most l.

If $p \in \mathbb{Z}/2$, then $q \in \mathbb{Z}/d$ for some $d \neq 2$, because consecutive letters in a reduced word cannot come from the same $\mathbb{Z}/d$. So we can find a different non-identity element $q' \in \mathbb{Z}/d$ and define the desired automorphism as follows. Given an arbitrary word, if it has q or q' as its $l+2^{\text{nd}}$ letter, then change it to q' or q, respectively. Otherwise, leave w unchanged. □

Lemma 5.8 *Let $d \in D$ with $d \neq 2$, and let w be a word in which a letter from $\mathbb{Z}/d$ occurs. Then there is an automorphism of $\mathfrak{A}_*$ that moves w but fixes all words containing no letter from $\mathbb{Z}/d$.*

Proof Suppose a letter $p \in \mathbb{Z}/d$ occurs in w at position l. Since $d > 2$, let p' be a different non-identity element of $\mathbb{Z}/d$. The desired automorphism can be defined as follows. Given an arbitrary word, if it has p or p' as its l^{th} letter, then change it to p' or p, respectively. Otherwise, leave the word unchanged. □

The preceding lemmas provide certain automorphisms of $\mathfrak{A}_*$. They yield automorphisms of $\mathfrak{A}$, because they can be used in any component $\{i\} \times A_*$, each component being isomorphic to $\mathfrak{A}_*$. In addition, automorphisms of $\mathfrak{A}$ can be obtained by permuting the components; that is, for any permutation π of $\mathbb{N}$, there is an automorphism of $\mathfrak{A}$ defined by $(i, w) \mapsto (\pi(i), w)$.

We can now prove easily that A is Dedekind-finite in M, by applying the criterion in Proposition 2.1. If E is a finite subset of A, then we shall show that every atom (i, w) supported by E must satisfy all the following conditions.

- There is an element of E in the same component $\{i\} \times A_*$.
- $\text{length}(w) \leq \text{length}(w') + 1$ for some $(i, w') \in E$.
- Every letter in w comes from $\mathbb{Z}/2$ or from the same $\mathbb{Z}/d$ as at least one letter in w' for at least one $(i, w') \in E$.

Indeed, if the first condition were violated, then we could move (i, w) while fixing E pointwise, by permuting components. If the second condition were violated, then Lemma 5.7 provides an automorphism of the i^{th} component moving (i, w) while fixing the elements of E in this component; extend it by the identity on the remaining components, so as to get an automorphism fixing all of E. Finally, if the third condition were violated, proceed analogously, using Lemma 5.8.

There are only finitely many elements (i, w) satisfying these three conditions. Indeed, by the first condition, they all come from just finitely many components. By the third condition, they use letters from only finitely many of the groups $\mathbb{Z}/d$; as these groups are finite, they use only finitely many letters. And by the second condition, their lengths are bounded.

Thus, E supports only finitely many atoms, and by Proposition 2.1 we know that A is Dedekind-finite in M. It remains to prove that no $q \geq 2$ outside D divides A in M.

Suppose, therefore, that $q \geq 2$ and that $\mathcal{Q}$ is a partition of A into q-element sets, supported by a finite set $E \subseteq A$. We must show that $q \in D$.

Call a component $\{i\} \times A_*$ of A *free* if it contains no element of E. Notice that all but finitely many components are free.

If a block of the partition $\mathcal{Q}$ meets a free component, then it is entirely included in that component. Indeed, suppose $Q \in \mathcal{Q}$ meets a free component C and also another component C'. Since Q is finite, there is a component C'' that meets neither Q nor E. Then a permutation that interchanges C and C'' while fixing all other atoms will fix E pointwise, will therefore fix $\mathcal{Q}$, will fix an element of Q in C', will therefore fix Q, but will move an element of Q in C to an element outside Q in C'', a contradiction.

From now on, we concentrate on a single free component. For simplicity of notation, we identify this component $\{i\} \times A_*$ with A_*, suppressing all mention of i and writing w instead of (i, w). By the preceding paragraph, $\mathcal{Q}$ restricts to a partition, still called $\mathcal{Q}$, of this component A_* into pieces of size q. Furthermore, this partition of A_* is invariant under all automorphisms of $\mathfrak{A}_*$, because any automorphism of this component could be extended to all of A by fixing all other atoms; it then fixes E pointwise and therefore fixes $\mathcal{Q}$. Our task is thus reduced to analyzing a partition $\mathcal{Q}$ of A_* into q-element sets that is invariant under all automorphisms of $\mathfrak{A}_*$. We shall frequently use the fact that, since $\mathcal{Q}$ is invariant, any automorphism π applied to a block $Q \in \mathcal{Q}$ produces again a block of $\mathcal{Q}$. In particular, if $x \in Q$ and $\pi(x) \in Q$ then $\pi(Q)$, being a block and having an element $\pi(x)$ in common with Q, must coincide with Q.

Let l be the maximum distance, in the adjacency graph, between any two elements of A_* that lie in the same block of $\mathcal{Q}$. If this maximum is 0, then the blocks of $\mathcal{Q}$ are singletons, contrary to our assumption that $q \geq 2$.

Suppose next that $l = 1$ and consider an arbitrary block $Q \in \mathcal{Q}$ and (as $q \geq 2$) two arbitrary, distinct elements $x, y \in Q$. As $l = 1$, x is d-adjacent to y for some $d \in D$. Any other element $z \in Q$ is, for the same reason, d'-adjacent to x for some $d' \in D$. If $d \neq d'$ then y and z would not be adjacent (apply Corollary 5.3 to the sequence y, x, z). So $d = d'$ and all of Q is therefore included in the $\mathcal{P}_{d*}$-block B of x. Suppose, toward a contradiction, that Q were properly included in B. Then Lemma 5.6 would provide an automorphism fixing x and sending y to an element of $B - Q$. Fixing x, this automorphism would have to fix Q, yet it moves an element $y \in Q$ out of Q, a contradiction. Therefore $Q = B$ and in particular $q = d$. (Notice that, when $l = 1$, we have shown that $\mathcal{Q}$ must coincide on free components with $\mathcal{P}_{d*}$ for some $d \in D$. Not only is the divisor q one of the divisors we wanted, but the partition $\mathcal{Q}$ is, on free components, one of the partitions we explicitly

put into our model.)

Suppose next that $l = 2$. Let x and y be two elements at distance 2 in some $\mathcal{Q}$-block Q. So there is some z adjacent to both x and y; say it is d-adjacent to x and d'-adjacent to y. We know $d \neq d'$, for otherwise the distance between x and y would be only 1. In particular, at least one of d, d' is not 2; without loss of generality assume $d \neq 2$. So the $\mathcal{P}_d$-block containing x and z also contains another element z'. Lemma 5.6 provides an automorphism π fixing x and sending z to z'. Notice that

- y is d'-adjacent to z,
- z is d-adjacent to z', and
- z' is d'-adjacent to $\pi(y)$ (as π preserves d'-adjacency).

Since $d \neq d'$, we can apply Corollary 5.3 to conclude that the distance from y to $\pi(y)$ is 3. On the other hand, since π fixes x, it fixes the $\mathcal{Q}$-block Q containing it. Since $y \in Q$, it follows that $\pi(y) \in Q$. So we have y and $\pi(y)$ both in Q yet at distance 3. This contradicts the assumption that $l = 2$, which is therefore impossible.

Suppose next that $l = 3$. Let x and y be two elements at distance 3 in some $\mathcal{Q}$-block Q. Say x is d-adjacent to z, which is d'-adjacent to w, which is d''-adjacent to y. To abbreviate this situation, we write

$$x - (d) - z - (d') - w - (d'') - y.$$

Of course d' is distinct from both d and d'', as otherwise we'd have a shorter path from x to y, skipping z or w.

If $d \neq 2$ then we can argue much as in the case $l = 2$. Let z' be another element in the d-block of x and z and let π be an automorphism fixing x and therefore Q but sending z to z'. Then

$$y - (d'') - w - (d') - z - (d) - z' - (d') - \pi(w) - (d'') - \pi(y).$$

By Corollary 5.3, the distance from y to $\pi(y)$ is 5, contradicting, since $\pi(y) \in \pi(Q) = Q$, the fact that $l = 3$. This contradiction shows that we must have $d = 2$. (In particular, the case $l = 3$ cannot arise if $2 \notin D$.)

A symmetrical argument shows that $d'' = 2$. Of course then $d' \neq 2$. Let B be the block of $\mathcal{P}_{d'*}$ that contains z and w, and let $\tilde{B}$ be the set of 2-neighbors of the elements of B. Notice that every element has a unique 2-neighbor, so

$\tilde{B}$ has the same cardinality d' as B. As z, w are in B, their 2-neighbors x, y are in $\tilde{B}$ as well as being in Q.

Consider any other element of $\tilde{B}$, say v, the 2-neighbor of some $u \in B$ distinct from z, w. By Lemma 5.6, there is an automorphism π that fixes z and sends w to u. Since it preserves $\mathcal{P}_{2*}$, this π must also fix x and send y to v. Fixing x it must fix its $\mathcal{Q}$-block Q. Since $y \in Q$ it follows that $v = \pi(y) \in Q$.

This proves that $\tilde{B} \subseteq Q$. Our next goal is to establish that $\tilde{B} = Q$. To prepare for this, we first observe that any element of B can be mapped to any other by an automorphism of $\mathfrak{A}$ that maps B onto itself; this follows immediately from Lemma 5.6 since the cardinality d' of B is greater than 2. Furthermore, any such automorphism maps $\tilde{B}$ into itself, since it preserves $\mathcal{P}_{2*}$. It therefore maps some members of Q (namely those in $\tilde{B}$) into Q and therefore must map Q into itself. Summarizing, we have that automorphisms of $\mathfrak{A}$ that fix all of B, $\tilde{B}$, and Q act transitively on B. In the following we shall refer to this observation as "transitivity on B."

Suppose now, toward a contradiction, that we had some element $t \in Q - \tilde{B}$. Consider the shortest path from x to t. Its length is at least one (since x is in $\tilde{B}$ and t isn't), so we can exhibit it as

$$x - (e) - x' - \cdots - t,$$

where $- \cdots -$ denotes a path of some length 0, 1, or 2 (since the distance from x to t is at most $l = 3$). If e were different from 2, then the path

$$y - (2) - w - (d') - z - (2) - x - (e) - x' - \cdots - t$$

would show, by Corollary 5.3, that the distance from y to t is at least 4, contradicting $l = 3$. So $e = 2$ and therefore $x' = z$. Thus, t is at a distance at most 2 from z and, a fortiori, at a distance at most 2 from B. We consider the possibilities for this last distance.

Case 0: $t \in B$. By transitivity on B, we may suppose without loss of generality that $t = z$. There is, by Lemma 5.4, an automorphism π that sends x to z. Since it preserves 2-adjacency, it must also send z to x. And since it preserves $\mathcal{Q}$ and sends one element x of Q to another, z, it must fix Q. Therefore, from $y \in Q$ we infer $\pi(y) \in Q$. But the path

$$y - (2) - w - (d') - z - (2) - x = \pi(z) - (d') - \pi(w) - (2) - \pi(y)$$

shows, by Corollary 5.3, that the distance from y to $\pi(y)$ is 5, contradicting $l = 3$. So Case 0 is impossible.

Case 1: t is at distance 1 from B. By transitivity on B, we may suppose without loss of generality, that t is adjacent to z. So we have

$$t - (e) - z - (d') - w - (2) - y$$

for some e. If $e = d'$ then $t \in B$ and we are back in Case 0. So $e \neq d'$, and the distance between t and y is 3 (by Corollary 5.3 again). But then we can argue for t, y exactly as we did for x, y above, to conclude that $e = 2$ and therefore t is the 2-neighbor of z, namely x. But $x \in \tilde{B}$ and $t \notin \tilde{B}$, so Case 1 is also impossible.

Case 2: t is at distance 2 from B. By transitivity on B, we may suppose without loss of generality, that

$$t - (e) - s - (e') - z$$

for some e, e' and some s. Of course $e \neq e'$ for otherwise the distance from t to B would be only 1, via e-adjacency to z. Consider the path

$$t - (e) - s - (e') - z - (d') - w - (2) - y.$$

If Corollary 5.3 applies to it, then the distance between t and y is 4, contradicting $l = 3$. So this corollary must not apply; the only way for this to happen is that $e' = d'$. But then $s \in B$, t is at distance 1 from B, and we are back in the previous case.

Having exhausted all the possibilities, we conclude that no such t can exist. Therefore $Q = \tilde{B}$. Therefore

$$q = |Q| = |\tilde{B}| = |B| = d' \in D.$$

This completes the proof in the case that $l = 3$.

Finally, suppose $l \geq 4$. Let x and y be two elements at distance l in a $\mathcal{Q}$-block Q. Consider the path of length l joining them, say

$$x - (d) - z - (d') - w - \cdots - y,$$

where d and d' are distinct; in particular they are not both 2. If $d \neq 2$ then, by Lemma 5.6, we can find an automorphism π fixing x and therefore Q

while sending z to a different element in the same $\mathcal{P}_{d*}$-block with x and z. Then Corollary 5.3 applied to the path

$$\pi(y) - \cdots - \pi(w) - (d') - \pi(z) - (d) - z - (d') - w - \cdots - y$$

(where the second $-\cdots-$ is as before and the first is obtained from it by applying π and reversing the order) shows that the distance between y and $\pi(y)$ is $2l-1$. Yet y and $\pi(y)$ are both in Q, where the maximum distance is l. So $2l-1 \leq l$, which contradicts $l \geq 4$.

There remains the case that $d = 2$. Then $d' \neq 2$, so Lemma 5.6 provides an automorphism π that fixes z while sending w to a different element of the same $\mathcal{P}_{d'*}$-block. Fixing z, this π must also fix its unique 2-neighbor x, and so it must fix the $\mathcal{Q}$-block Q of x. Thus, from $y \in Q$ we infer $\pi(y) \in Q$. Applying Corollary 5.3 to

$$\pi(y) - \cdots - \pi(w) - (d') - w - \cdots - y,$$

we find that the distance from y to $\pi(y)$ is $2l-3$. Since the maximum distance in any $\mathcal{Q}$-block is l, we must have $2l-3 \leq l$. But this contradicts $l \geq 4$. So this case cannot arise.

Summarizing, we have that the only possible values for l are 1 and 3, and in these cases q has to be a member of D. □

The preceding proof showed that, in the model M constructed there, not only is every divisor (≥ 2) of A in the specified set D, but any partition into q-element pieces must agree in all free components either with one of the "built in" partitions $\mathcal{P}_d$ of A or with one of the closely related partitions $\{\tilde{B} : B \in \mathcal{P}_d\}$ that we found in the analysis of the case $l = 3$.

The use of cyclic groups in this proof was inessential; any groups of the same cardinalities would serve as well.

6 Intermediate Divisibilities

Recall that strong divisibility of a Dedekind-finite set X by a natural number n was defined to mean that X is the union of n sets all of the same cardinality. The following proposition gives some easy equivalent formulations of this notion, one of which will motivate some additional notions of divisibility.

Proposition 6.1 *For any Dedekind-finite set X and any natural number n, the following are equivalent.*

1. *X is strongly divisible by n.*

2. *$X \cong n \times Y$ for some Y.*

3. *There exist a partition $\mathcal{P}$ of X into pieces of size n and a function assigning to each piece $P \in \mathcal{P}$ a linear ordering of P.*

Proof If $n = 0$ then all three assertions are trivially equivalent to $X = \varnothing$. So we assume from now on that $n \neq 0$.

To prove that (1) implies (2), assume that X is partitioned into n pieces $Y_0, \ldots, Y_{n-1}$ all of the same size. For each i in the range $0 \leq i < n$, let f_i be a bijection between Y_0 and Y_i. (Since i takes only finitely many values, choosing the f_i's doesn't require AC.) Then the map

$$n \times Y_0 \to X : (i, y) \mapsto f_i(y)$$

is a bijection as required in (2).

To prove that (2) implies (3), notice that $n \times Y$ has the structure required in (3), namely a partition into the pieces $n \times \{y\}$ and a linear ordering of each piece by putting $(i, y) < (j, y)$ if $i < j$ in $\mathbb{N}$. Then transport this structure to X via the bijection given by (2).

Finally, to prove that (3) implies (1), assume that $\mathcal{P}$ and f are given as in (3). Define Y_i for $i = 0, 1, \ldots, n-1$ to be the set of points $x \in X$ that have, in the piece $P \in \mathcal{P}$ containing them, exactly i strict predecessors with respect to the ordering $f(P)$. Then X is partitioned into these sets Y_i, and they all have the same cardinality because there is a bijection from Y_i to Y_j sending each element of Y_i to the (unique) element of Y_j in the same block of $\mathcal{P}$. □

Corollary 6.2 *Strong divisibility implies divisibility.*

Proof Immediate from part (3) of the proposition. □

Part (3) of the proposition shows exactly how strong divisibility goes beyond divisibility. In addition to the partition $\mathcal{P}$ required for divisibility, strong divisibility requires a certain structure, namely a linear ordering, to be specified in each piece of $\mathcal{P}$. It suggests other forms of divisibility defined by requiring other sorts of structure on the pieces. For example, we could ask that X be partitionable into pieces of size 3 with a choice of a cyclic ordering

of each piece. It is easy to see that this notion, which we call *divisibility by cyclic 3*, is implied by strong divisibility by 3 (for from a linear ordering of a piece we can produce a cyclic ordering in a canonical way) and in turn implies divisibility by 3.

In order to discuss this notion of "divisibility by n with some additional structure" in any generality, we need a general concept of "structure." Fortunately, such a notion, well suited to our purposes, is provided by Joyal's theory [5] of species of structures. It is quite abstract, requiring only that one is given, for each n-element set, a set of things called structures on that set, together with a reasonable way of transporting structures from one set to another via bijections between the sets. We shall impose an additional requirement, called atomicity, on the species we consider, to avoid such unnatural (for our purposes) notions of structure as "either a cyclic ordering or a partition into two parts."

Definition 6.3 An *atomic species* $\mathcal{S}$ of n-element structures, or for brevity an *n-species*, consists of an assignment of

- to each n-element set F a nonempty set $\mathcal{S}(F)$ called the set of $\mathcal{S}$-structures on F and
- to each bijection $f : F \to E$ between n-element sets a map, $\mathcal{S}(f)$, but sometimes simply called f, from $\mathcal{S}(F)$ to $\mathcal{S}(E)$

subject to the following requirements.

- If f is an identity map then so is $\mathcal{S}(f)$.
- $\mathcal{S}(f \circ g) = \mathcal{S}(f) \circ \mathcal{S}(g)$ whenever $f \circ g$ is defined.
- If $s_0, s_1 \in \mathcal{S}(F)$ then there is a permutation π of F such that $\mathcal{S}(\pi)(s_0) = s_1$.

The requirements other than the last one say simply that $\mathcal{S}$ is a functor from the category of n-element sets and bijections into the category of sets. They imply that $\mathcal{S}(f)$ is always a bijection and that $\mathcal{S}(f)^{-1} = \mathcal{S}(f^{-1})$. Given two sets, each equipped with an $\mathcal{S}$-structure, we call a bijection between them an isomorphism (of the structured sets) if it sends the given structure on one set to the given structure on the other. The last clause in the definition is the atomicity requirement, saying that all structures on a single set are isomorphic; there is only one sort of structure in $\mathcal{S}$.

Example 6.4 The simplest example of an atomic species of structures is no structure at all. It assigns to each n-element set some 1-element set; then there's no choice about what to assign to bijections.

A second example is the species of linear orders of n-element sets, assigning to each such set the set of all its linear orderings. If $f : F \to E$ is a bijection and $\leq$ is a linear ordering of F, then $\mathcal{S}(f)(\leq)$ is the linear ordering of E defined by $f^{-1}(x) \leq f^{-1}(y)$.

A closely related example is the set of cyclic orderings, with the obvious notion of transport along bijections.

A simpler but useful example is the species of pointed (or rooted) sets, sending each n-element set F and each bijection to itself, i.e. the identity functor. A structure of this species on a set F is simply a distinguished element in F, and an isomorphism is a bijection $F \to E$ sending the distinguished element of F to that of E.

There is an atomic species of structures of 3-element sets, where a structure is a partition of the set into two pieces. Since such a partition must have pieces of sizes 1 and 2 and since it's completely determined by the 1-element piece, this species is isomorphic, in a sense to be formally defined in a moment, to the species of pointed 3-element sets.

The species of 2-piece partitions of a 4-element set is not atomic, for partitions into two pieces of size 2 are not isomorphic to partitions into pieces of sizes 1 and 3. If one specifies the size of the pieces, then one gets atomic species. The species of partitions into 1+3 is isomorphic to the species of distinguished points. The species of partitions into 2+2 is not isomorphic to any of the preceding examples.

Definition 6.5 A *morphism* from one atomic species of n-element structures $\mathcal{S}$ to another $\mathcal{T}$ consists of an assignment i of a function $i_F : \mathcal{S}(F) \to \mathcal{T}(F)$ for every n-element set F such that, for any bijection $f : F \to E$ between n-element sets,

$$\mathcal{T}(f) \circ i_F = i_E \circ \mathcal{S}(f).$$

If all the functions i_F are bijections, then i is an *isomorphism* of species.

The definition of isomorphism simply requires the $\mathcal{S}$-structures on any set to be matched with the $\mathcal{T}$-structures (on the same set) in a way that is preserved by transport of structures along bijections. In the language of category theory, it says that i is a natural isomorphism from the functor $\mathcal{S}$ to $\mathcal{T}$. Similarly, morphisms of species are natural transformations.

Atomic species of structures on n-element sets can be classified up to isomorphism; they correspond to conjugacy classes of subgroups of the symmetric group. The following definition shows how to obtain species from subgroups. We use the symbol $\mathrm{Sym}(n)$ for the symmetric group of all permutations of the set $n = \{0, 1, \ldots, n-1\}$.

Definition 6.6 Let G be a subgroup of $\mathrm{Sym}(n)$. The species $\mathcal{B}_G$ of *bijections modulo* G is defined as follows. For any n-element set F, let $\mathrm{Bij}(n, F)$ be the set of bijections from n to F, and let $\sim_G$ be the equivalence relation on $\mathrm{Bij}(n, F)$ given by

$$b \sim_G b' \iff (\exists g \in G)\, b' = b \circ g.$$

Then $\mathcal{B}_G(F)$ is the set of equivalence classes $[b]_G$. For a bijection $f : F \to E$, we define $\mathcal{B}_G(f)([b]_G) = [f \circ b]_G$. When only a single G is under consideration, we shall often omit the subscript G.

It is easy to check that $\mathcal{B}_G$ is an n-species. The next proposition says that, up to isomorphism, there are no other n-species. It and the subsequent proposition are due to Joyal [5].

Proposition 6.7 *Let $\mathcal{S}$ be an atomic species of n-element structures, let $a \in \mathcal{S}(n)$, and let $G = \{\pi \in Sym(n) : \mathcal{S}(\pi)(a) = a\}$. Then G is a subgroup of $Sym(n)$ and the species $\mathcal{S}$ and $\mathcal{B}_G$ are isomorphic.*

Proof That G is a subgroup follows immediately from the requirement, in the definition of species, that composition and identities be respected. An isomorphism i from $\mathcal{B}_G$ to $\mathcal{S}$ is defined by

$$i_F([b]_G) = \mathcal{S}(b)(a).$$

This is well-defined because, if $[b] = [b']$ then $b' = bg$ for some $g \in G$ and therefore $\mathcal{S}(b')(a) = \mathcal{S}(b)(\mathcal{S}(g)(a)) = \mathcal{S}(b)(a)$, where the last equation comes from the definition of G. To see that i respects transport along bijections $f : F \to E$ we compute that, for any $[b] \in \mathcal{B}_G(F)$,

$$\mathcal{S}(f)(i_F([b])) = \mathcal{S}(f)(\mathcal{S}(b)(a)) = \mathcal{S}(fb)(a) = i_E([fb]) = i_E(\mathcal{B}_G(f)([b])).$$

To show that i_F is one-to-one, suppose $i_F([b]) = i_F([b'])$, i.e., $\mathcal{S}(b)(a) = \mathcal{S}(b')(a)$. Then a is fixed by $\mathcal{S}(b)^{-1}\mathcal{S}(b') = \mathcal{S}(b^{-1}b')$. So $b^{-1}b' \in G$ and therefore $b \sim_G b(b^{-1}b') = b'$ and $[b] = [b']$.

Finally, to show that i_F is surjective, consider any $s \in \mathcal{S}(F)$ and any bijection $b : n \to F$. By the atomicity condition on $\mathcal{S}$, there is a permutation π of F sending $i_F([b])$ to s. Then, in view of the computation above showing that i respects transport along bijections, we have

$$s = \mathcal{S}(\pi)(i_F([b])) = i_F([\pi b]) \in \text{Range}(i_F).$$

□

For a complete classification of n-species we need, in addition to the preceding proposition, a description of which G's lead to isomorphic species. The following proposition provides this and a bit more information that will be useful later.

Proposition 6.8 *Let G and G' be subgroups of Sym(n). There is a morphism from $\mathcal{B}_G$ to $\mathcal{B}_{G'}$ if and only if G is conjugate in Sym(n) to a subgroup of G'. The species $\mathcal{B}_G$ and $\mathcal{B}_{G'}$ are isomorphic if and only if G is conjugate to G'.*

Proof Suppose first that G is conjugate to a subgroup of G', say $\pi^{-1}G\pi \subseteq G'$, where π is some permutation of n. Then a morphism m from $\mathcal{B}_G$ to $\mathcal{B}_{G'}$ is defined by $m_F([b]_G) = [b\pi]_{G'}$. This is well-defined because, if $b' = bg$ with $g \in G$, then $b'\pi = bg\pi = (b\pi)(\pi^{-1}g\pi)$ with $\pi^{-1}g\pi \in G'$. The verification that m_F respects transport amounts to $(fb)\pi = f(b\pi)$.

Notice that m_F is always surjective, for every bijection $n \to F$ is of the form $b\pi$ for some b. If G is conjugate to G' via π then, applying the preceding paragraph with the roles of G and G' interchanged and with π replaced by π^{-1}, we obtain a morphism from $\mathcal{B}_{G'}$ to $\mathcal{B}_G$ that is obviously a two-sided inverse for m. So when the subgroups are conjugate the corresponding species are isomorphic. (Alternative argument: If $G' = \pi^{-1}G\pi$ then G' is the stabilizer of the element $[\pi^{-1}]_G \in \mathcal{B}_G(n)$. So, by the proof of Proposition 6.7, $\mathcal{B}_G \cong \mathcal{B}_{G'}$.)

Conversely, suppose there is a morphism $m : \mathcal{B}_G \to \mathcal{B}_{G'}$. Let π be an element of $\text{Bij}(n, n) = \text{Sym}(n)$ such that $[\pi]_{G'} = m_n([1]_G)$, where 1 means the identity permutation of n. Then, for every $g \in G$, we have, since morphisms respect transport by g and since, by definition of $\sim_G$, $[1]_G = [g]_G$,

$$\begin{aligned}[\pi]_{G'} = m_n([1]_G) = m_n([g]_G) &= m_n(\mathcal{B}_G(g)([1]_G)) \\ &= \mathcal{B}_{G'}(g)(m_n([1]_G) = \mathcal{B}_{G'}(g)([\pi]_{G'}) = [g\pi]_{G'}.\end{aligned}$$

This means that $g\pi = \pi g'$ for some $g' \in G'$. But then $g' = \pi^{-1} g \pi$. So we have shown that $\pi^{-1} G \pi \subseteq G'$, as required.

Finally, if we have an isomorphism between $\mathcal{B}_G$ and $\mathcal{B}'_G$, then each of G and G' is conjugate to a subgroup of the other. Since they are finite, they must be conjugate. (It is possible to give a slightly longer argument that doesn't depend on finiteness.) □

By the preceding propositions, isomorphism classes of n-species correspond exactly to conjugacy classes of subgroups of $\mathrm{Sym}(n)$. In particular, the subgroup $\mathrm{Sym}(n)$ corresponds to the structure giving no information, and the subgroup $\{1\}$ corresponds to linear ordering (for the stabilizer of a linear order on a finite set is trivial). For $n = 2$ there are no other examples. For $n = 3$ there are two. The cyclic subgroup of order 3 in $\mathrm{Sym}(3)$ corresponds to the structure of cyclic ordering, and the three (conjugate) subgroups of order 2 correspond to the structure of a distinguished point. We do not describe all the 4-species, since there are 11 of them, but we mention that the species of partitions into 2+2 corresponds to the 2-Sylow subgroup of $\mathrm{Sym}(4)$.

It follows immediately from the definition of $\mathcal{B}_G$ that the number of its structures on any n-element set is the index of G in $\mathrm{Sym}(n)$.

We now return to considerations of divisibility.

Definition 6.9 Let $\mathcal{S}$ be an n-species. A Dedekind-finite set X is *divisible by n with structure $\mathcal{S}$* if there exist a partition $\mathcal{P}$ of X into pieces of size n and a function assigning to each piece $P \in \mathcal{P}$ an $\mathcal{S}$-structure on it, i.e., an element of $\mathcal{S}(P)$. We sometimes use the alternative terminology "divisible by $\mathcal{S}$-structured n."

Thus, ordinary divisibility and strong divisibility are the special cases where $\mathcal{S}$ is no structure and linear ordering, respectively.

Proposition 6.10 *If there is a morphism $i : \mathcal{S} \to \mathcal{T}$ and if X is divisible by $\mathcal{S}$-structured n, then it is also divisible by $\mathcal{T}$-structured n.*

Proof Given a partition $\mathcal{P}$ and an $\mathcal{S}$-structure $s(P)$ on each piece P, assign to each P the $\mathcal{T}$-structure $i_P(s(P))$. □

We are now in a position to extend Corollary 3.3 to structured divisibility.

Theorem 6.11 *Let G be a subgroup of $Sym(d)$ such that,*

$$(\exists a \in d)(\forall c \in d - \{a\})(\exists g \in G)\,(g(a) = a \text{ and } g(c) \neq c).$$

Then it is consistent to have a Dedekind-finite X such that $X \pm n$ is divisible by d with $\mathcal{B}_G$-structure for all natural numbers n.

Proof Fix $a \in d$ as in the hypothesis, and let H be its stabilizer in G,

$$H := \{g \in G : g(a) = a\}.$$

It will suffice to arrange for $\mathcal{B}_H$-structured divisibility, because there is a morphism from $\mathcal{B}_H$ to $\mathcal{B}_G$, by Proposition 6.8. From now on, we ignore G and work only with H. It will also be convenient to assume, without loss of generality, that $a = 0$. Thus, H leaves 0 fixed, but does not fix any other element of d.

As in the proof of Corollary 3.3, let A be a complete $d-1$-branching tree. For each node s let $b_s : d \to A$ be a one-to-one map sending 0 to s and sending $1, \ldots, d-1$ onto the children of s in the tree A. Thus b_s is a bijection from d to the family F_s consisting of s and its children; so $[b_s]_H$ is a $\mathcal{B}_H$-structure on F_s. Let $\mathfrak{A}$ be the set A together with its tree structure and with all these structures $[b_s]_H$ on the families F_s. Thus, an automorphism π of $\mathfrak{A}$ is an automorphism of the tree which not only maps F_s onto $F_{\pi(s)}$ (as any automorphism of the tree must) but also sends $[b_s]_H$ to $[b_{\pi(s)}]_H$. This means that $\pi \circ b_s = b_{\pi(s)} \circ h$ for some $h \in H$. Notice that part of this condition is automatically satisfied, by any automorphism of the tree. Namely,

$$\pi(b_s(0)) = \pi(s) = b_{\pi(s)}(0) = b_{\pi(s)}(h(0)).$$

So the additional requirement imposed by structure preservation, beyond being a tree automorphism, is that the children of any node s be mapped "correctly" to the children of the image node.

Given any natural number l and any node t at depth greater than l in the tree, we have an automorphism π of the tree that moves t while fixing all nodes of A of depth at most l. To see this, let s be the parent of t, and let $t = b_s(c)$, where $c \in d - \{0\}$. The desired automorphism π will fix all nodes that are not descendants of s. That includes all nodes of depth at most l, as required. Our π will fix s but permute its children as follows. By hypothesis, we can choose an $h \in H$ that moves c. Then move members of F_s according to the rule $b_s(x) \mapsto b_s(h(x))$. This fixes s as claimed, but it moves $t = b_s(c)$ because h moves c (and b_s is one-to-one). It preserves the structure $[b_s]_H$ on F_s because it sends it to $[\pi \circ b_s]_H = [b_s \circ h]_H$ which equals $[b_s]_H$ because $h \in H$. Finally, extend the automorphism to the distant descendants of s by

induction on depth in the tree; once $\pi(r)$ is defined, for r a proper descendant of s, let $\pi(b_r(x)) = b_{\pi(r)}(x)$ for all $x \in d - \{0\}$. This clearly preseves the structure.

Now let M be the permutation model determined by $\mathfrak{A}$. The result of the preceding paragraph implies that any finite set $E \subseteq A$ supports only finitely many atoms, namely at most those whose depth in the tree is no bigger than that of the deepest member of E. Therefore, by Proposition 2.1, A is Dedekind-finite in M. Exactly as in the proof of Corollary 3.3 (i.e., as in the proof of Theorem 3.1), for each $n = 0, 1, \ldots, d-1$ we can partition $A - n$ into pieces of the form F_s, which have size d. In addition, these pieces have $\mathcal{B}_H$-structures $[b_s]_H$. The function assigning to each piece this structure is in M; it is symmetric simply because our $\mathfrak{A}$ includes these structures and so all automorphisms preserve them. Therefore, every $A \pm n$ is divisible in M by $\mathcal{B}_H$-structured d and therefore also by $\mathcal{B}_G$-structured d. □

Example 6.12 For $d = 2$, no G can satisfy the hypothesis of the theorem, for a permutation fixing one of the two elements of d must also fix the other. So our theorem does not apply to $d = 2$, as it must not in view of Theorem 3.4.

For $d = 3$, the hypothesis is satisfied by the full symmetric group Sym(3) and by the two-element subgroups (but not by the 3-element subgroup or the trivial group). For the full symmetric group, the theorem simply repeats Theorem 3.1. For the two-element subgroups, the theorem tells us that we can have Dedekind finite X and all $X \pm n$ divisible by pointed 3. This is not new information, since the partitions used in the proof of Theorem 3.1 actually had pointed pieces. Each piece consisted of a node and its two children; let the parent be the distinguished point. (This assignment of structure to each piece is in the model because the tree structure is in the model.)

For $d = 4$, four conjugacy classes of subgroups satisfy the hypothesis of the theorem. Two of them, namely the full Sym(4) and the subgroup Sym(3) fixing one point, are analogous to those in the $d = 3$ case. In addition, there is a subgroup of order 3 that fixes one point and permutes the other three cyclically, and there is the alternating group of order 12. The former gives us the consistency of a Dedekind-finite X where each $X \pm n$ is divisible by 4 with, in each piece, a distinguished point and a cyclic ordering of the other 3 points. The latter gives the weaker result, where each piece has simply an "orientation," i.e., an enumeration defined up to even permutations.

The last theorem can be improved somewhat, with only a minor modification of the proof.

Theorem 6.13 *Let d and G be as in the preceding theorem. It is consistent to have a Dedekind-finite X strongly divisible by d and all $X \pm n$ divisible by d with $\mathcal{B}_G$-structure.*

Proof We indicate only the modifications needed in the preceding proof. A is the same tree as before. For nodes s of odd depth, F_s has a $\mathcal{B}_G$-structure as before. For nodes s of even depth (including the root of depth 0) F_s has a $\mathcal{B}_{\{1\}}$-structure, which amounts to a linear ordering (by Proposition 6.8). $\mathfrak{A}$ is A with the tree structure and all these structures on the F_s's. M is the resulting permutation model. The partition of X into d-element sets, consisting of the F_s for s of even length, has a specified linear order on each piece (the assignment of linear orders to these pieces being symmetric because all automorphisms of $\mathfrak{A}$ preserve it). For $X - n$ with $n = 1, 2, \ldots, d-1$, we have a partition into pieces F_s with, in general, some odd and some even depth s's. Some of these pieces are equipped with a $\mathcal{B}_G$-structure and others with a $\mathcal{B}_{\{1\}}$-structure. But any structure of the latter sort produces one of the former, via the morphism from Proposition 6.8. So all pieces have been assigned a $\mathcal{B}_G$-structure. For other $X \pm n$ we reduce to the cases already considered by adding or subtracting multiples of d.

The proof of Dedekind-finiteness of A in Theorem 6.11 depended on the assertion that, given any natural number l and any node t of depth greater than l, we have an automorphism π moving t while fixing all nodes of depth at most l. The proof of this assertion depended on permuting F_s in a non-trivial way, so this argument is available for our present model when the length of s is odd, but not when it is even. Nevertheless, we get a permutation of the required sort providing the length of t exceeds l by at least 2. Then, if t has even length we can argue as before, while if t has odd length we can apply the previous argument with the parent node of t in place of t itself. (Note that an automorphism moving the parent of t necessarily moves t also.) It follows that a finite $E \subseteq A$ can support only atoms whose depth is at most one more than that of the deepest node in E. There are only finitely many such nodes, so A is Dedkind finite by Proposition 2.1. □

Example 6.14 It follows from this theorem that we can have a Dedekind-finite X strongly divisible by 3 while $X - 1$ and $X - 2$ are divisible by cyclic 3. This proves somewhat more than was claimed in Theorem 3.6.

The preceding two theorems concerned the situation where some point-stabilizer in G is fairly large. The next result handles the other extreme, where all point-stabilizers are trivial.

Theorem 6.15 *Let $d \geq 2$. Suppose G is a subgroup of $Sym(d)$ such that no element of G except the identity fixes any element of d. If both X and $X-1$ are divisible by d with $\mathcal{B}_G$-structure, then X is Dedekind-infinite.*

Proof Let G be as in the hypothesis, and consider first an arbitrary d-element set F with a $\mathcal{B}_G$-structure. Recall that this structure is an equivalence class $[b] = [b]_G$ of bijections from n to F. We claim that, for any $i \in d$ and any $x \in F$, there is at most one element of $[b]$ that maps i to x. Indeed, if $b_1 = bg_1$ and $b_2 = bg_2$ were two such elements, where $g_1, g_2 \in G$ by definition of $\sim_G$, then

$$g_1^{-1}g_2(i) = g_1^{-1}b^{-1}bg_2(i) = b_1^{-1}b_2(i) = b_1^{-1}x = i.$$

By the hypothesis on G, it follows that $g_1^{-1}g_2$ is the identity, so $g_1 = g_2$ and $b_1 = b_2$, as claimed.

Still considering a set F with a $\mathcal{B}_G$-structure $[b]$, suppose we are given an element $x \in F$. We claim that, from these data, we can canonically define (without arbitrary choices) a linear ordering $\preceq = \preceq_{[b],x}$ of F. To do this, let i be the smallest number in $d = \{0, 1, \ldots, d-1\}$ such that some member of $[b]$ maps i to x. Of course such an i exists, because the members of $[b]$ are bijections. Having determined i, let b' be the member of $[b]$ that maps i to x; it is unique by the preceding paragraph. Use this b' to transport the standard ordering of the natural numbers from d to F; that is, $b'(j) \preceq b'(k)$ if and only if $j \leq k$.

Now consider a set X such that X and $X-1$ are both divisible by d with $\mathcal{B}_G$-structure. Fix a partition $\mathcal{P}$ of X and a partition $\mathcal{Q}$ of $X - \{a\}$, for a certain $a \in X$, into d-element pieces, such that there are functions assigning to each piece $P \in \mathcal{P}$ and to each piece $Q \in \mathcal{Q}$ a $\mathcal{B}_G$-structure on that piece. Fix all these structures for the rest of the argument.

We shall prove that X is Dedekind-infinite by constructing a countably infinite sequence of distinct members of X. The construction will proceed in stages. At each stage, we shall have a finite sequence s of distinct members of X, and we shall extend it by appending finitely many new elements (at least one, possibly more) to it. The construction begins with s consisting of just the one term a, the element omitted for the partition $\mathcal{Q}$.

At any stage, let s be the sequence constructed so far. If the length of s is not divisible by d, then proceed as follows. Since all blocks of the partition $\mathcal{P}$ have size d, there must be a term in the sequence s whose $\mathcal{P}$-block does not consist entirely of terms of s. Let x be the first such term in s, and let P be its $\mathcal{P}$-block. From x and the given $\mathcal{B}_G$-structure on P, we obtain as above a linear ordering $\preceq$ of P. Append to s all the members of P that are not already terms of s, listing them in the order given by $\preceq$.

On the other hand, if the length of s is divisible by d, then consider s without its first term a; this is a sequence of elements of the set $X - \{a\}$ partitioned by $\mathcal{Q}$, and its length is not divisible by d (since $d \geq 2$). So there must be a term in the sequence s whose $\mathcal{Q}$-block does not consist entirely of terms of s. Now proceed as in the preceding paragraph, using the first such element, its $\mathcal{Q}$-block, and the linear ordering $\preceq$ of that block induced by this element and the given $\mathcal{B}_G$-structure.

Since at least one new element is appended to s at every stage, and no element is ever repeated, the construction defines a countably infinite sequence of distinct elements of X. □

Example 6.16 The trivial group $\{1\}$ satisfies the hypothesis of this theorem, so a Dedekind-finite set X and $X - 1$ cannot both be strongly divisible by d. This establishes Theorem 3.5.

The cyclic subgroup of $\mathrm{Sym}(d)$ generated by a cycle of length d also satisfies the hypothesis of the theorem. Therefore, a Dedekind-finite X and $X - 1$ cannot both be divisible by d with cyclic ordering.

These theorems cover the cases of groups $G \subseteq \mathrm{Sym}(n)$ with either very large or very small point stabilizers. There are many intermediate cases, and we do not have general results concerning these. The next theorem describes what happens for one special class of such groups.

Theorem 6.17 *Let $d = p_1 + \cdots + p_k$ with all $p_i \geq 2$, and let $\mathcal{S}$ be the d-species where a structure on F is a partition of F into pieces $F_1, \ldots, F_k$ of sizes $p_1, \ldots, p_k$ together with a cyclic ordering of each piece. If the greatest common divisor of $p_1, \ldots, p_k$ is 1, then it is consistent to have a Dedekind-finite X such that all $X \pm n$ are divisible by d with structure $\mathcal{S}$. If the greatest common divisor of $p_1, \ldots, p_k$ is greater than 1, then a Dedekind-finite set X and $X - 1$ cannot both be divisible by d with structure $\mathcal{S}$.*

To avoid possible confusion, we emphasize that in an $\mathcal{S}$-structure the pieces F_i are supposed to be numbered. So even if two of the p_i's are the same, interchanging the corresponding F_i's would produce a different $\mathcal{S}$-structure.

Proof Suppose first that the greatest common divisor is 1, so we have an equation of the form $1 = p_1 r_1 + \cdots + p_k r_k$ where the r_i's are integers. Transposing the terms with negative r_i's and renumbering for notational convenience, we have

$$1 + p_1(-r_1) + \cdots + p_l(-r_l) = p_{l+1} r_{l+1} + \cdots + p_k r_k$$

where now all the coefficients $-r_i$ for $i \leq l$ and r_i for $i > l$ are positive.

On a countably infinite set A of atoms, put the following structure:

- a partition of A into pieces P_n ($n \in \mathbb{N}$) of size d,
- for each piece P_n, a partition into subpieces $P_{n,i}$ of sizes p_i ($1 \leq i \leq k$),
- the function $(n, i) \mapsto P_{n,i}$ labeling the subpieces, and
- a cyclic ordering of each subpiece.

Let $\mathfrak{A}$ be the set A with all this structure. Thus an automorphism of $\mathfrak{A}$ must map each subpiece onto itself, preserving the given cyclic order. Let M be the resulting permutation model. Since the partition into pieces P_n and the assignment, to each P_n, of its partition into subpieces and the cyclic orderings of the subpieces are symmetric, A is divisible by d with structure $\mathcal{S}$ in M.

We show next that $A+1$ has the same divisibility property. Let x be the element added to A to form $A+1$. Let Z be the set consisting of x and, for each $i \leq l$, the elements of $P_{n,i}$ for the first $(-r_i)$ values of n. The cardinality of this set is, by our choice of the r_i's, $p_{l+1} r_{l+1} + \cdots + p_k r_k$, so we can partition it into r_j blocks of size p_j for $j = l+1, \ldots, k$, and we can fix (in M) a cyclic ordering of each of these blocks (finitely many blocks, so AC is not used). Compared to the original structure $\mathfrak{A}$ on A, we have, by adding the point x and rearranging finitely much of the structure, gained r_i cyclically ordered blocks of size p_i, for all i, where a negative gain means a loss. But now it is easy to rearrange these new blocks along with the surviving old subpieces into a partition $\mathcal{Q}$ of $A+1$ into pieces of size d equipped with $\mathcal{S}$-structures. The n^{th} block of $\mathcal{Q}$ consists of the blocks $P_{n-r_i,i}$, where for positive r_i the subpieces with negative first subscript mean the blocks in Z newly created

by the rearrangement. (For negative r_i, the first $-r_i$ numbers do not occur as first subscript, which is correct, since those subpieces were lost into Z in the rearrangement.)

To show the same divisibility property for $A + n$ with $n > 1$, use the same technique, but rearrange n times as many subpieces. For $A - n$, add a multiple of d to return to the case where points are added to A. Thus, all $A \pm n$ are divisible by d with $\mathcal{S}$-structure.

It remains to show that A is Dedekind-finite in M, but this is clear by Proposition 2.1. If E is a finite subset of A then it supports only atoms in the same subpieces $P_{n,i}$ as members of E. Any other atom can be moved by a cyclic permutation of its subpiece, fixing all other atoms. This completes the proof of the theorem's first assertion.

To prove the second assertion, suppose the greatest common divisor of the p_i's is $d' \geq 2$, and suppose both X and $X - 1$ are divisible by d with $\mathcal{S}$-structure. Thus, both of these sets are partitioned into finite subpieces of sizes divisible by d' with a cyclic ordering on each subpiece. Thus, as soon as a point is specified in any such subpiece, a linear ordering of that whole subpiece is available: start at the given point and proceed in the given cyclic order.

In this situation, we can produce a countably infinite sequence of distinct members of X exactly as in the proof of Theorem 6.15, only with d' in the role previously played by d. □

Remark 6.18 The proof of the theorem can be used to get the following slightly more precise information about the structure $\mathcal{S}$ considered there. Let d' be the greatest common divisor of the p_i's. Then it is consistent to have a Dedekind-finite set X such that $X \pm n$ is divisible by d with $\mathcal{S}$-structure for all natural numbers n that are divisible by d'. On the other hand, if n is not divisible by d' and if both X and $X - n$ are divisible by d with structure $\mathcal{S}$, then X is Dedekind-infinite.

7 Quotients

It seems reasonable that a discussion of divisibility should include a discussion of quotients, the result of division. If X is divisible by d, say $\mathcal{P}$ is a partition of X into pieces of size d, then the cardinality of $\mathcal{P}$, i.e., the number of pieces, seems to be a reasonable notion of quotient when X is divided by d. The

question arises whether this quotient is well-defined. That is, if $\mathcal{Q}$ is another partition of X into pieces of size d, must we have $\mathcal{P} \cong \mathcal{Q}$?

In the case of strong divisibility, where the pieces of both partitions are equipped with specified linear orderings, an affirmative answer was given by Lindenbaum and Tarski [6]; see also [7] and [1].

Theorem 7.1 (Lindenbaum and Tarski) *If d is a positive integer and $d \times X \cong d \times Y$ then $X \cong Y$.*

This theorem holds for arbitrary X and Y, whether or not they are Dedekind-finite.

The assumption of strong divisibility in the theorem is necessary, even for Dedekind-finite X and Y. No weaker form of divisibility with structure will do.

Theorem 7.2 *Let $d \geq 2$, and let $\mathcal{S}$ be any d-species not isomorphic to the species of linear orderings. Then it is consistent to have a Dedekind-finite set X with two partitions, $\mathcal{P} \ncong \mathcal{Q}$ into d-element pieces, together with functions assigning to each piece $P \in \mathcal{P}$ an $\mathcal{S}$-structure on P and to each piece $Q \in \mathcal{Q}$ a linear ordering of Q.*

Recall that, by Proposition 6.8, a linear ordering of Q provides a structure of any other d-species on Q. Thus, the theorem says that quotients are not unique, no matter how much structure is imposed on the pieces of the two partitions, unless one imposes linear orders on the pieces of both partitions. In other words, the Lindenbaum-Tarski result is the best possible one along these lines.

Proof We shall need to know that every $\mathcal{S}$-structured set F has a non-trivial automorphism. To show this, we may assume, by Proposition 6.7, that $\mathcal{S} = \mathcal{B}_G$ for some subgroup G of $\mathrm{Sym}(d)$. Since $\mathcal{S}$ is not isomorphic to the species of linear orderings, G is not the trivial subgroup $\{1\}$; let $g \in G - \{1\}$. Now consider any $\mathcal{B}_G$-structure $[b]_G$ on F; here b is a bijection $d \to F$. Let $\pi = bgb^{-1}$. This is a permutation of F, and it is not the identity because $g \neq 1$. It is an automorphism of the structure $[b]_G$ because it sends this structure to

$$[\pi b]_G = [(bgb^{-1})b]_G = [bg]_G = [b]_G$$

where the last equality is because $g \in G$.

Now we turn to the construction of the required model. We use a variant of the model in Example 2.5. As in that example, let the countably infinite set A of atoms be partitioned into pieces P_n ($n \in \mathbb{N}$) of size d. Choose (in $V(A)$ where AC is available) an $\mathcal{S}$-structure on each piece P_n. Let $\mathfrak{A}$ be A together with this partition and this choice of $\mathcal{S}$-structures on the pieces, and let M be the resulting permutation model. Let $X = d \times A$.

This X is Dedekind-finite by Proposition 2.1, for a finite set E of atoms supports only those elements of X that have the form (i, a) with $i \in d$ and a in a P_n that meets E (so E supports at most $d^2|E|$ members of X).

Let $\mathcal{P}$ be the partition of X induced by the original partition $\mathcal{P}_0 = \{P_n : n \in \mathbb{N}\}$ of A. That is, the pieces of $\mathcal{P}$ are the sets $\{i\} \times P_n$ for $i \in d$ and $n \in \mathbb{N}$. Since $\mathcal{P}_0$ and the chosen assignment of $\mathcal{S}$-structures to its pieces P_n are clearly symmetric, it follows that $\mathcal{P} \in M$ and that its pieces can be assigned $\mathcal{S}$-structures by transferring the given $\mathcal{S}$-structures of the P_n's. Thus, we have a partition and assignment witnessing that X is divisible by d with $\mathcal{S}$-structure. The set of pieces of this partition is in one-to-one correspondence (in M) with $d \times \mathcal{P}_0$, piece $\{i\} \times P_n$ corresponding to (i, P_n). So we can take the quotient to be $d \times \mathcal{P}_0$. (To avoid possible confusion we point out that the bijection $n \mapsto P_n$ was not included in the structure $\mathfrak{A}$, and it is easy to see that $\mathcal{P}_0$ is not countably infinite in M; in fact it is Dedekind-finite by an easy application of Proposition 2.1.)

Our second partition $\mathcal{Q}$ of X consists of the pieces $d \times \{a\}$ for all $a \in A$. This witnesses strong divisibility of X by d (see Proposition 6.1), and the pieces are in obvious one-to-one correspondence with A.

Thus, to complete the proof, it suffices to show that there is no bijection $f : d \times \mathcal{P}_0 \to A$ in M. Suppose, toward a contradiction, that f were such a bijection, supported by a finite set $E \subseteq A$.

Consider any P_n not meeting E and consider $f(i, P_n)$ for any $i \in d$. We claim it must be an element of P_n. To see this, suppose it were in some other P_m, and choose some P_k that is distinct from both P_n and P_m and does not meet E. Then $\mathfrak{A}$ has an automorphism π that interchanges P_n and P_k while fixing all other atoms. This π fixes E pointwise and therefore fixes f. It also fixes $f(i, P_n)$ since it is in a block P_m not moved by π. And of course it fixes the natural number i. But it sends P_n to P_m. Thus,

$$f(i, P_n) = \pi(f(i, P_n)) = f(i, P_m),$$

which contradicts the fact that f is one-to-one. This establishes our claim.

So we now know that f maps each $d \times \{P_n\}$ into (and therefore onto) P_n, as long as P_n is disjoint from E. Consider any such P_n and let π be an automorphism of $\mathfrak{A}$ that fixes all atoms outside this P_n but moves at least one atom, say a, in P_n. The existence of such an automorphism is guaranteed by the first paragraph of this proof. Now $a = f(i, P_n)$ for some $i \in d$. Since π fixes E pointwise, it fixes f. It also fixes i and P_n. (Of course it doesn't fix P_n pointwise, but it fixes it as a set.) Therefore,

$$a = f(i, P_n) = \pi(f)(\pi(i), \pi(P_n)) = \pi(f(i, P_n)) = \pi(a),$$

which contradicts the fact that π moves a. This contradiction shows that no such f can exist, so $\mathcal{P} \not\cong \mathcal{Q}$. □

References

[1] J. Conway and P. Doyle, Division by three, `http://math.dartmouth.edu/~doyle/docs/three/`

[2] A. Fraenkel, Über den Begriff "definit" und die Unabhängigkeit des Auswahlaxioms, *Sitzungsberichte der Preußischen Akademie der Wissenschaften (Berlin), Physik.-math. Klasse* (1922) 253–257.

[3] T. Jech, *The Axiom of Choice*, North-Holland, 1973.

[4] T. Jech and A. Sochor, Applications of the Θ-model, *Bull. Acad. Polon. Sci. Sér. Sci. Math. Astronom. Phys.* 14 (1966) 351–355.

[5] A. Joyal, Une théorie combinatoire des séries formelles, *Adv. Math.* 42 (1981) 1–82.

[6] A. Lindenbaum and A. Tarski, Communication sur les recherches de la théorie des ensembles, *Comptes Rendus des Séances de la Société des Sciences et Lettres de Varsovie, Classe III* 19 (1926) 299–330.

[7] A. Tarski, Cancellation laws in the arithmetic of cardinals, *Fund. Math.* 36 (1949) 77–92.