1
Introduction

1.1 Overview

The purpose of this work is to lift the notion of restricted orbit equivalence to the category of free and ergodic actions of discrete amenable groups. We mean *lift* in two senses. First of course we will generalize the results in [43], where Rudolph developed a theory of restricted orbit equivalence for $\mathbb{Z}$-actions, and in [25] where both authors later established a similar theory for actions of $\mathbb{Z}^d, d \geq 1$ to actions of these more general groups. However, we will also *lift* in the sense that we will develop the axiomatics and argument structures in what we feel is a far more natural and robust fashion. Both [43] and [25] were based on axiomatizations of a notion called a "size" measuring the degree of distortion of a box in $\mathbb{Z}^d$ caused by a permutation. It is not evident that on their common ground, $\mathbb{Z}$-actions, these two theories agree. Hence we refer to the first as a 1-size and the second as a p-size, p for "permutation".

Here we will establish the axiomatics of what we will simply call a size. We ask that the reader accept this new definition. In the Appendix we show that any equivalence relation that arose from a p-size will arise from a size as we define it here. The same is not done for 1-sizes, but for a slight strengthening of this axiomatics that includes all the examples in [43].

We will work on the level of countable and discrete amenable groups, as the work of Ornstein and Weiss [37] has shown this to be a natural level on which all the basic dynamics and ergodic theory we need holds sway. We will not push beyond this to locally-compact amenable groups as the formalism of both orbit equivalence and entropy theory require basic work before our approach appears feasible.

Outlining our approach, in Section 2.1 we will set up the basic vocabulary we will use for our work, the vocabulary of arrangements and rearrangements of orbits, and describe the natural topologies on these spaces. In Section 2.2 we will establish the axiomatics of a size m and the nature of the associated equivalence relation, m-equivalence. One sees immediately here the change in perspective from [43] and [25] in that a size m is now a family of pseudometrics on the full-group of a free and ergodic action, one for each arrangement of the orbit as an action of the group G. The m-equivalence class of an arrangement will appear here as a certain G_δ subset of the completion of the full-group relative to this pseudometric. We end Chapter 2 with a list of seven equivalence relations, some well known others not so well known, which can be described as m-equivalences for an appropriate m. We also present one "non-example" that uses the methods developed here but does not fall directly under our development and indicates one of several directions in which to further broaden this approach.

In Chapter 3 we present the fundamental results that we will need from the Ornstein and Weiss work on the ergodic theory of actions of amenable groups. In Chapter 4 we present a variety of copying lemmas that will be essential to our progress both in developing an entropy theory for restricted orbit equivalences and for our proof of the equivalence theorem. Chapter 5 contains our development of an entropy theory for restricted orbit equivalences. We define an entropy, called m-entropy, associated with each size as the infimum of the classical entropy on the m-equivalence class. The principle result we obtain, (as was done in [43] and [25] for the cases they considered) is that a restricted orbit equivalence is either entropy-preserving, in that m-entropy is simply the classical entropy, or entropy-free in that on a residual subset of the equivalence class the entropy is zero and hence the m-entropy of all actions is zero.

From this point our goal is to prove the natural generalization of Ornstein's isomorphism theorem for Bernoulli shifts for our restricted orbit equivalences. That is to say, we wish to show that there are certain distinguished free and ergodic actions, intrinsically recognizable, for which m-entropy is a complete invariant of m-equivalence. As we indicated earlier our goal in *lifting* results is both to demonstrate that they hold in the more general context and also to raise the general level of argument to a more robust form. In particular the approach to the equivalence theorem we take is to bring to bear the categorical approach that Burton and Rothstein [42] brought to the isomorphism theorem.

To accomplish this it is necessary to build up a certain topological perspective. We will be considering Polish spaces and Polish actions. Recall that these are topological spaces that can be imbedded as a G_δ subset of a compact metric space, and homeomorphisms of them. Chapter 6 presents this development from its foundations through to the proof that the space of m-joinings of free and ergodic G-actions form a Polish space of measures. Although we cannot expect the reader to have any idea at this point what precisely an m-joining is, the point of such a result should be clear to the reader familiar with the Burton–Rothstein approach. Any m-equivalence between two actions will sit as a subset of this space of m-joinings. Just as Burton and Rothstein show that for Bernoulli actions of equal entropy, the conjugacies are a residual subset of the space of joinings, our aim is to show that for any two m-Bernoulli actions of equal m-entropy, the m-equivalences sit as a residual subset of the space of m-joinings.

To obtain this in Section 7.2 we introduce the notion of an m-finitely determined action and develop some of its basic properties, in particular that it is an m-equivalence invariant and is inherited by factor actions. What remains to complete the equivalence theorem is to define a list of open sets in the space of m-joinings whose intersection is precisely the m-equivalences, and to show that in the case of m-finitely determined actions they are dense. The first part of this is easy. It is the second part that takes some work. Central to this proof of density is of course the copying lemmas which we developed in Chapter 4. With them we show how to *perturb* an arbitrary m-joining to lie in one of our open sets. Section 7.1 presents the basic structure theory for the notion of *perturbation* we will use. Section 7.3 finally completes the equivalence theorem.

To connect this work with earlier work, in the Appendix we will demonstrate that p-sizes give sizes in our sense with the same equivalence classes, and that 1-sizes essentially do in that all known examples do, and in general the m-finitely determined classes are the same. This means the large classes of examples discussed there are restricted orbit equivalences according to the definition used here.

1.2 A roadmap to the text

We offer the readers an indication of how they might best benefit from this text depending on the level and nature of their interest. This text is not intended as an introduction to the isomorphism theory of Ornstein. A reasonable preparation is necessary before the material presented here will be comfortably accessible. To the reader interested in a broad overview of dynamics we recommend the texts of Walters [62], Petersen [39], Cornfeld, Fomin and Sinai [7], Hasselblatt and Katok [17], and Brin and Stuck [4]. We also recommend the reader consult Ornstein [35], Shields [50], Rudolph [44], and Ornstein, Rudolph and Weiss [38]. Of these [44] is perhaps the most important reference as it is intended as a basic technical introduction to fundamentals at work here. For the reader seeking a deep understanding of our presentation these older works will provide valuable perspective and background.

We envision two audiences for this text, those seeking a basic technical overview of the tools and methods of this theory for use in their own work but without an intention to work in this area and students and researchers seeking a deep understanding of this area with the intention of working in it. We give here an abbreviated path through the text for the former audience and a recommended first reading for the latter.

All readers should spend time on Chapter 2. For a reader only interested in what the text is about, Chapter 2 offers a sufficient treatment. The material of Chapter 3, through Theorem 3.0.6, is the now classical treatment of the ergodic theory of discrete amenable groups due to Ornstein and Weiss [37]. Understanding this is a must for anyone interested in modern ergodic theory. The next few pages introduce the vocabulary of names for the entropy theory of these actions and it is important to understand them. Beyond this, the important conclusions of Chapter 3 are Theorem 3.0.9 and Corollary 3.0.10 and the reader should be familiar with their meaning.

Chapter 4 is quite technical. What is essential are the two Theorems 4.0.5 and 4.0.13 and one should understand their meaning. The technicalities of proof can be absorbed later if needed. What is essential from Chapter 5 concerning m-entropy is contained in its first paragraph.

Chapters 6 and 7 are really only appropriate for the reader wanting a detailed understanding of the equivalence theory. Chapter 6 develops the space of m-joinings of two actions as a Polish space of measures on a symbolic representation. This provides a framework on which the equivalence theorem of Chapter 7 can be reached without much ado. Chapter 6 though is quite heavy going. The reader should read Section 6.1 as an introduction to Polish spaces ending with a brief and vague description of the succession of spaces constructed to reach the notion of an m-joining. From here one can proceed to Chapter 7, skipping over Section 7.1 and instead focusing on the definitions and results of Section 7.2. Continue with Section 7.3 through the statement of Theorem 7.3.3.

The path just described gives a complete overview of the text. To continue the reader should now return for a careful reading of Chapter 6 and then Section 7.1.

1.3 History and references

The notion of restricted orbit equivalence can be traced back some decades, and the techniques used here can be traced back somewhat further. We take this opportunity to outline our understanding of the history of these ideas and to acknowledge the sources of our work. There are many significant parts of this broad area that we will not mention here. We focus on those particular ideas central to the evolution of this particular part of the theory, the construction of general orbit equivalence theorems, and the tools of those constructions.

Certainly the first pieces, historically, of this story do not relate directly to notions of orbit equivalence, but rather are the basic technical point-set tools of ergodic theory, in particular the Rokhlin lemma [31], [41], the mean ergodic theorem [61], and the entropy theory of Kolmogorow and Sinai [28], [29], [51] (we cite the earliest references we know to the basic results and methods). Of these three the Rokhlin lemma may seem the most mundane but in the long run it is in fact the deepest of the three. Looking forward to the seminal work of Ornstein and Weiss [37] on actions of amenable groups, it is their ability to realize a Rokhlin lemma that in fact carries forward their entire program.

Dye's proof, in 1959 [10], that any two ergodic measure preserving actions on non-atomic standard probability spaces are orbit-equivalent is the first instance of an orbit equivalence theorem, and of course a very startling one. The core technical pieces of this proof are of course the Rokhlin lemma and the ergodic theorem. Dye's original techniques are still visible in all the copying lemmas that follow. The only issue that he did not have to address was entropy. Dye's result was simply too profound and final in the case of measure-preserving actions and so has had much more impact in non-singular dynamics and Von Neuman algebras than in ergodic theory per se. Krieger [30], [54] was able to characterize the orbit equivalence classes of all non-singular ergodic actions by careful study of the obstacles to Dye's argument. The obstacle to moving our work in the direction of non-singular dynamics is the issue of entropy. Certainly, though, one might consider developing the notion of entropy-free restricted orbit equivalences in this setting. On the other hand, perhaps the attempt to lift these methods to the non-singular category will give some insight into how to approach entropy in this category as entropy in the measure-preserving category can be defined by a restricted orbit equivalence (see Section 2.3).

The next step, and truly the pivotal one, in this development was Ornstein's proof that any two Bernoulli shifts of equal entropy were isomorphic [32], [33], [34]. One must remember Sinai [52] had already shown they were weakly isomorphic, that is to say each sat as a factor action of the other. Ornstein's real contribution, from our perspective, is his general characterization of the Bernoulli shifts via the notion of finitely determined actions and more generally the wide range of powerful constructive tools he laid out. At this time the notion of finitely determined was given in a finitistic form, not in terms of joinings as is more common now; but the lift to the joinings perspective is a modest contribution in comparison to the significance of the original concept. This theorem is a phenomenal piece of technical work. Others have had profound insights since then but this result showed what the path to an equivalence theorem would look like. All one had to do was see how to take the steps.

One also should note that the Ornstein machinery provides tools for showing not only that Bernoulli shifts of equal entropy are conjugate, but also that various collections of actions are not conjugate. Parallel to the positive side of equivalence/isomorphism results, one could use these methods to develop non-isomorphism results. For example, Ornstein and Shields [36] constructed uncountable families of non-conjugate ergodic K-systems, all of the same entropy. Relating back to Sinai's theorem, Polit [40] produced a pair of weakly-isomorphic but non-isomorphic actions (in this case zero-entropy mixing actions). This side of ergodic theory is extensive.

During the same period Vershik [60] began an investigation building on Dye's work, but taking a different focus. In orbit equivalence terms he was considering actions of groups other than $\mathbb{Z}$, in particular actions of infinite sums of finite cyclic groups. Such a group is the union of a sequence of finite groups $\mathscr{H}_n$, where each $\mathscr{H}_n/\mathscr{H}_{n+1}$ is a cyclic group of order r_n. The simplest non-trivial case is where the r_n's are all two, the dyadic case. What he considered was in fact a notion of restricted orbit equivalence, asking that the orbit equivalence between two such actions should be an orbit equivalence of each of the $\mathscr{H}_n$-subgroups. Heicklen has completely translated this work into the vocabulary of restricted orbit equivalence as we develop it here (see [19], [18] and Example 6 in Section 2.3). In truth this is more naturally described in terms of the $\mathscr{H}_n$-invariant σ-algebras $\mathscr{F}_n$. Notice that these form a decreasing sequence (or reverse filtration) of algebras $\mathscr{F}_{n+1} \subseteq \mathscr{F}_n$. Ergodicity of the action is equivalent to the algebras intersecting to the trivial algebra. Vershik's notion of equivalence is simply that there should be a map between the measure spaces respecting the two reverse filtrations. In this sense Vershik's study has a similar feel to that of Dye, the natural generalizations move away from measure-preserving actions. One is led to consider general reverse filtrations, without regard to their arising from a measure-preserving action.

In terms of orbit equivalence theory though, this particular development has a particularly prescient nature. Vershik showed that there was an entropy associated with this relation and, that if the r_n's did not grow too quickly, it was the standard entropy of the action, and if they grew quickly enough it would be zero. Using this observation both he and Stepin [53] were able to construct non-equivalent reverse filtrations as they had distinct entropies. Central to Vershik's study is the notion of standardness. This also points to a central aspect of later work, and certainly ours, that there will be certain distinguished classes, the "Bernoulli" class of the given equivalence relation. (Vershik's "standard" class is the zero-entropy m-finitely determined class for the associated size m.) This work gives the first clear indication of how entropy might enter in a general picture.

The next major contribution in the direction of our work also had two sides, one in the west and one in the east. The notion of Kakutani equivalence had arisen some decades earlier in the study of measurable cross-sections of measure-preserving actions of $\mathbb{R}$ [2], [3], [22]. It was known that although entropy was not an invariant of Kakutani equivalence, entropy class was (zero, finite, or infinite entropy) and that entropy changed in a simple way when moving among equivalent actions (Abramov's formula [1]). What Feldman did in the west and what Katok did in the former Soviet Union was to introduce the $\bar{f}$-metric (Feldman's notation) on names. Feldman used this to show that there were many distinct Kakutani equivalence classes of the same entropy class [11]. He also saw the possibility that $\bar{f}$ might plug into Ornstein's isomorphism machinery and lead to an equivalence theory parallel to Ornstein's conjugacy theory. Katok [26], in part jointly with Sataev [49], and Ornstein, Weiss and Rudolph fulfilled that expectation [38]. Katok and Sataev, of course, were working completely independently of Feldman, Ornstein and Weiss.

Building on Feldman's original examples, and Ornstein's observation that the Cartesian square of a rank-1 and mixing map would not be finitely fixed, many exotic examples were constructed. This constructive side was also pursued by Katok, leading to his construction, via this theory, of the first smooth K and not Bernoulli action [26].

Building on the existence of three theorems, Dye's, Ornstein's, and the Feldman, Katok, Ornstein, Weiss–Kakutani equivalence theorem, Feldman proposed in 1975 the potential for a general theory of equivalence relations based on the common structures in these three results. There were two essential gaps in the picture. Kakutani equivalence is not a restricted orbit equivalence in that it is not an orbit equivalence. This seemed a minor issue in that inducing on a subset is not so far from an orbit equivalence. More critical though was to understand what role entropy would play. The isomorphism theorem and Kakutani equivalence theorem both use entropy in much the same way, following Ornstein's basic plan. But entropy does not enter Dye's theorem at all. Of course Vershik's work had already indicated what the answer might be, but this was not well known in the west at the time. The one real gap in the picture was exactly how to phrase a general theorem, although the thought certainly was to create the needed material to apply Ornstein's method, that is to say, define some analogue of $\bar{d}$ or $\bar{f}$.

On a more technical level, in the late 1970s Burton and Rothstein gave an approach to Ornstein's proof that recast the focus to joinings [42], [43]. In these terms what one sees is not the detailed construction of a single conjugacy, but rather the description of all conjugacies as a residual subset of a space of joinings. Although in a pure sense there was nothing really new in what Rothstein did, it made the whole path to the result much clearer. The critical technical piece of the isomorphism theorem was the copying lemma which forced the denseness of certain open sets in a space of joinings. The intersection of these sets were precisely the factor maps, or projections, of some fixed but arbitrary system of sufficient entropy onto a finitely determined system. In this light one could say that Sinai had it right, that although weak isomorphism did not imply isomorphism, if the class of projections of a general system on a finitely determined system of equal entropy could be shown large enough (that is, shown residual in the compact space of all joinings) then there would be lots of isomorphisms to choose from between two finitely determined systems of equal entropy and residual set of them. It was not clear until our work here that this perspective could be lifted to the more general restricted orbit equivalence level. In particular [25], [37] and [43] all follow the original direct constructive approach of Ornstein.

When del Junco and Rudolph showed that Kakutani equivalence could be given a natural characterization in terms of orbit equivalence [9] the general picture became much clearer. As we pointed out earlier, from Abramov's formula we know that the entropy of an induced map varies inversely proportionally to the measure of the set on which one induces. If one defines a notion of "even equivalence" of two actions to mean that one induces conjugate actions in each on subsets of the same measure, then this equivalence relation preserves entropy. Moreover, they showed that the conjugacy could be extended from the subsets to the rest of the two ambient spaces as an orbit equivalence precisely because the two sets have the same measure. An orbit equivalence arising in this fashion could be characterized in a variety of ways; in particular in ways that extended to higher dimensional actions. A variety of authors have pursued this area (Katok [27] for example). The last piece in this particular development is Hasfura-Buenaga's proof of an equivalence theorem in $\mathbb{Z}^d$ [16].

One could now see all three, orbit equivalence, even Kakutani equivalence, and conjugacy, as restricted orbit equivalences and look for the common thread in the corresponding equivalence theorems. This is what was attempted in [43]. The basic structure laid out there was that one should axiomatize a notion that measures how badly one is distorting an orbit. In both [43] and [25] this is based on an axiomatization of how one would measure the wildness of a permutation of a large block of the acting group $\mathbb{Z}$ or $\mathbb{Z}^d$. These axiomatizations were rather elaborate and technical, especially in [43].

Very quickly it became evident that the formulation in [43] was quite flawed. The author admits there that the basic structures will not lift reasonably to larger group actions, and natural examples arose of equivalence relations, that were restrictions on orbit equivalence, that "ought" to be but could not be brought into the framework of [43]. For example, α-equivalences described here in Examples 4 and 5 cannot arise from a 1-size. At a more subtle level, Fieldsteel indicated quite rightly that the definition given for an m-joining was not sufficiently robust. This is because an m-joining was required to be a joining perturbed by a "bounded coboundary", not by an m-equivalence. In particular an m-equivalence itself was not necessarily an m-joining.

The reason for this concern over basics is that as we indicated earlier there is a good deal more to the isomorphism theory than just Ornstein's theorem itself. There are all the constructive examples built to show it attains the best one could hope for. There is the "relativized theory" of

Thouvenot [55] leading to his deep study of the weak-Pinsker property [56]. In particular one has Thouvenot's result that the property of being Kakutani equivalent to a map with the weak-Pinsker property is inherited by factor actions. Fieldsteel and one of the authors generalized this to the [43] theory of entropy-preserving sizes [14]. One also has the theory of isometric and affine extensions of Bernoulli actions [23], [24], [46]. One could hope to generalize all of this work from conjugacy, and Bernoulli actions, to restricted orbit equivalences and m-finitely determined actions. This has been done for some examples [43]. Fieldsteel [12] showed that if one took compact group extensions of two ergodic actions, by the same group, so that the extensions were ergodic, then one could construct an orbit equivalence between the two extensions that preserved the group extension structures.

In a result which really deserves deeper study Fieldsteel and Friedman [13] showed that Belinskaya's amazing theorem (if the generating functions of an orbit equivalence were integrable, then the equivalence was essentially trivial) was false in higher dimensions even if *integrable* was replaced with *bounded*. This indicates that there is perhaps a non-trivial L^1 and even L^∞ orbit-equivalence theory in higher dimensions which is vacuous in 1. That is to say, there may be families of sizes in $\mathbb{Z}^d$ for all d, which for some d, say $d = 1$, give the same equivalence relations, but for larger d do not.

At this point the authors put forward the development in [25] as a restricted orbit equivalence theory for actions of $\mathbb{Z}^d$. This work is quite parallel to [43]. In [43] one takes an injection from a block of integers (i, j) and "pushes together" the range set to obtain a permutation of (i, j). In [25] one constructs permutations of boxes in $\mathbb{Z}^d$ from injections by "filling in" that is by taking those points that the map throws out of the box and placing them on the points in the box which have no preimage. These are similar, but far from the same notion. Certainly this new picture was much more generalizable, but still was tied to basic structures in the group $\mathbb{Z}^d$, in particular boxes.

In the meantime Ornstein and Weiss [37] had finished their seminal study of the ergodic theory of amenable group actions. This profound work makes it evident that the natural level at which the three basic tools of ergodic theory (Rokhlin lemma, ergodic theorem and a Shannon–McMillan theorem) apply is that of amenable groups. To be more precise, it is clear that they all apply at the level of discrete amenable groups. It still remains unclear to what degree the entropy theory for general amenable groups is well-founded. As our work on orbit equivalence will