The Number of Path-Components of a Compact Subset of $\mathbb{R}^n$

Howard Becker *

Department of Mathematics
University of South Carolina
Columbia, SC 29208 USA
becker@math.sc.edu

§0. Introduction

This paper is concerned with the following question. Assume $\neg CH$; does there exist a compact set $K \subset \mathbb{R}^n$ such that K has exactly $\aleph_1$ path-components? For $\mathbb{R}^3$, the answer is yes. For $\mathbb{R}^2$, the answer is no, assuming a weak large cardinal axiom (which may or may not be necessary).

The proof of both results is descriptive set theoretic. Indeed, the motivation for asking the question is descriptive set theoretic. The same question for components, rather than path-components, would be a silly question; it is obvious (at least to descriptive set theorists) that the answer is no. It is also obvious that it is not possible that $2^{\aleph_0} \geq \aleph_3$ and that there is a compact $K \subset \mathbb{R}^n$ with $\aleph_2$ path-components. But the question as posed above does not seem to be a silly question. One of the purposes of this paper is to present the descriptive set theoretic point of view, and hopefully convince the reader that these "obvious" facts really are obvious. Two references for descriptive set theory are Kechris [13] and Moschovakis [17], and we follow their notation and terminology.

In both the $\mathbb{R}^3$ and $\mathbb{R}^2$ cases, we have results that are stronger than those stated above. In both cases, the size of the continuum is irrelevant and the theorem – properly stated – is nontrivial even if CH is true. These theorems will be given in §2. For $\mathbb{R}^3$, there is a more general theorem, a precise version of the following: Any $\boldsymbol{\Sigma}^1_1$ equivalence relation can be coded up as the equivalence relation of being in the same path-component of K, for some compact $K \subset \mathbb{R}^3$. From this it easily follows that there is a $K \subset \mathbb{R}^3$ with $\aleph_1$ path-components. That general theorem has other applications as well, one of which answers a question of Kunen-Starbird [14]. This paper is largely an explanation of the statement of these stronger theorems, and of the larger mathematical theory of which they are a part, that is, the descriptive set theory of equivalence relations. In the $\mathbb{R}^3$ case we say virtually nothing about the proof. In the $\mathbb{R}^2$ case we give an outline of the proof (§§6,7), containing several gaps, and using a stronger large cardinal axiom than required.

The author plans to some day write a long paper about path-connectedness, simple connectedness and descriptive set theory (Becker [3]). The results announced here will appear there with complete proofs. Most of Becker [3] will be concerned with calculating the complexity, with respect to the projective

---

* Partially supported by NSF Grant DMS–9505505

CAMBRIDGE

Cambridge University Press
978-1-107-16790-2 — Logic Colloquium '95
Edited by Johann A. Makowsky , Elena V. Ravve
Excerpt
More Information

hierarchy, of the following pointsets in the space $\mathcal{K}(\mathbb{R}^n)$ of compact subsets of $\mathbb{R}^n$:

$$PC_n = \{K \in \mathcal{K}(\mathbb{R}^n) : K \text{ is path-connected}\},$$
$$SC_n = \{K \in \mathcal{K}(\mathbb{R}^n) : K \text{ is simply connected}\}.$$

Several theorems of this sort were announced in Becker [2, Example 16 ff.], and proofs of some of them have appeared in Kechris [13, Theorems 33.17 and 37.11]. (*Remark.* There has been one new result since these publications appeared. Darji [7] and Just [10], independently, proved that PC_2 is not $\boldsymbol{\Sigma}^1_1$.) This topic is related to the results in this paper. The proof that there is a $K \in \mathcal{K}(\mathbb{R}^3)$ with $\aleph_1$ path-components has much in common with the proof that PC_2 *is not* $\boldsymbol{\Pi}^1_1$. The proof that, assuming large cardinals and $\neg CH$, there is no such K in $\mathcal{K}(\mathbb{R}^2)$, has much in common with the proof that SC_2 *is* $\boldsymbol{\Pi}^1_1$.

We work in ZFC. When anything more is used in a theorem it will be explicitly stated in the hypothesis.

§1. Path-components

Our basic reference for topological matters is Kuratowski [15]. Our terminology is standard, and mostly consistent with that reference.

Set theorists have a habit of calling practically anything "the reals". But here, topology actually matters, so the *reals* always means the reals. It is denoted by $\mathbb{R}$. The letter K will always denote a compact subset of $\mathbb{R}^n$ for some n. While our main interest is in such a space K, we give the definitions in more generality.

Definition 1. Definition Let X be a topological space and let $\mathbf{p}$, $\mathbf{q}$ be points in X. A *path from* $\mathbf{p}$ *to* $\mathbf{q}$ *in* X is a continuous function $\gamma : [0,1] \to X$ such that $\gamma(0) = \mathbf{p}, \gamma(\mathbf{1}) = \mathbf{q}$. An *arc* is a one-to-one path.

We sometimes abuse the language and refer to the pointset $Im(\gamma)$ as "the path γ" or "the arc γ".

For any topological space X, let $\approx_X$ denote the following equivalence relation on X:

$$\mathbf{p} \approx_{\mathbf{X}} \mathbf{q} \iff \text{there exists a path from } \mathbf{p} \text{ to } \mathbf{q} \text{ in } \mathbf{X}.$$

The $\approx_X$-equivalence classes are called the *path-components* of X. X is *path-connected* if it has only one path-component.

Path-connectedness and path-components should not be confused with a different notion: connectedness and components. (*Connected* means no nontrivial clopen sets, and a *component* is a maximal connected subset.) While path-connectedness implies connectedness, the converse is false, even for compact subsets of $\mathbb{R}^2$. The standard counterexample is $K^* = A_1 \cup A_2$, where

$$A_1 = \{(x, y) : -1 \leq x < 0 \text{ and } y = sin(1/x)\},$$
$$A_2 = \{(x, y) : x = 0 \text{ and } -1 \leq y \leq 1\}$$

(see Figure 1). K^* is connected. But K^* is not path-connected; it has exactly two path-components, A_1 and A_2.

Cambridge University Press
978-1-107-16790-2 — Logic Colloquium '95
Edited by Johann A. Makowsky , Elena V. Ravve
Excerpt
More Information

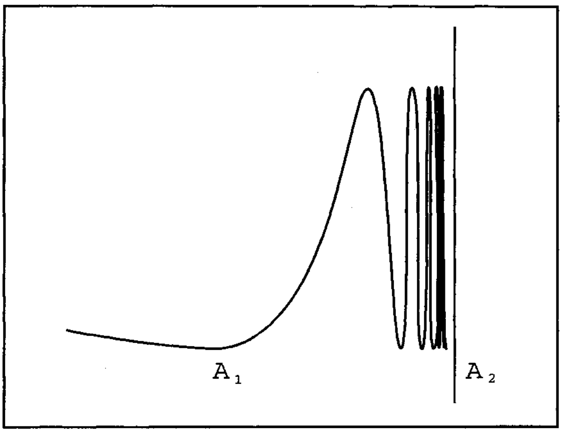

Figure 1

Theorem 1.1 Let $\mathbf{p}, \mathbf{q} \in \mathbb{R}^n$ and let $\gamma : [0,1] \longrightarrow \mathbb{R}^n$ be a path from $\mathbf{p}$ to $\mathbf{q}$. If $\mathbf{p} \neq \mathbf{q}$ then there is an arc γ' from $\mathbf{p}$ to $\mathbf{q}$ such that $Im(\gamma') \subset Im(\gamma)$.
Proof. See Kuratowski [15, §50, I, Theorem 2 and II, Theorem 1]. □

By 1.1, for any $X \subset \mathbb{R}^n$, path-components are the same thing as *arc-components* and path-connectedness the same as *arc-connectedness.* (In fact, for any Hausdorff space, the two concepts coincide.)

§2. Statement of theorems

We have two theorems, 2.1 and 2.2, below, which answer the question posed at the beginning of this paper.
Theorem 2.1 There is a compact set $K \subset \mathbb{R}^3$ with the following properties.

(a) K has exactly $\aleph_1$ path-components.
(b) There does not exist a nonempty perfect set $P \subset K$ such that any two distinct points of P are in different path-components of K.

The above theorem is proved in ZFC. The next theorem is not quite proved in ZFC, but rather in $ZFC + \epsilon$. (A precise description of ϵ is given below.)
Theorem 2.2 Assume ϵ. For any compact set $K \subset \mathbb{R}^2$, one of the following holds:

(i) K has only countably many path-components;
(ii) There is a nonempty perfect set $P \subset K$ such that any two distinct points of P are in different path-components of K. (Hence K has $2^{\aleph_0}$ path-components.)

The axiom ϵ is the following statement:

Every uncountable $\boldsymbol{\Sigma}^1_2$ set of reals contains a nonempty perfect subset.

By a theorem of Solovay (see Kanamori [11, Theorem 14.10]) ϵ is equivalent to:

$$\text{For all } a \subset \omega,\ \aleph_1^{L[a]} < \aleph_1.$$

The axiom ϵ is equiconsistent with the existence of an inaccessible cardinal (see Kanamori [11, Theorem 11.6]), and thus it is a "large cardinal axiom" by virtue of its consistency strength, although it does not, of course, imply the actual existence of large cardinals. Serious large cardinal axioms, e.g., the existence of a measurable cardinal, imply that ϵ is *true* (as opposed to merely consistent). Hence these large cardinal axioms imply that the conclusion of 2.2 is true. For more information on large cardinal axioms, see Kanamori [11].

This axiom has been around for a long time, and has been explicitly considered as a hypothesis of theorems, but does not seem to have ever been given a name. To rectify that oversight, I have decided to call it ϵ. Compared to the large cardinal axioms commonly used in set theory these days, this axiom is a *very* weak assumption – the name ϵ is entirely appropriate.

Theorem 2.2 leads to an interesting open question in reverse mathematics: Is 2.2 provable in weak subsystems of $ZFC + \epsilon$, such as ZFC? It is possible that it is provable in ZFC. But I would conjecture that it is not, and that, in fact, the following is provable in ZFC: There exists a compact $K \subset \mathbb{R}^2$ and a bijection between the path-components of K and $\aleph_1^L$. If this is the case, then in all models where $\aleph_1^L = \aleph_1 < 2^{\aleph_0}$, the answer to the question posed at the beginning of this paper would be yes, even for $\mathbb{R}^2$; hence a large cardinal axiom really would be necessary to get a no answer.

§3. Descriptive set theory and equivalence relations, I: Theorems of Silver and Burgess

If E is an equivalence relation on X and $Y \subset X$, Y is called *E-invariant* if for all $y, y' \in X$:

$$y \in Y \text{ and } yEy' \Longrightarrow y' \in Y.$$

Definition2. Definition Let X be a Polish space, let E be an equivalence relation on X, and let $Y \subset X$ be E-invariant. We say that Y has *perfectly many* E-equivalence classes if there is a nonempty perfect set $P \subset Y$ such that no two distinct points of P are E-equivalent.

Clearly perfectly many equivalence classes implies $2^{\aleph_0}$ equivalence classes. In fact, "perfectly many" is, in some sense, an effectivized version of "continuum many": Y has continuum many equivalence classes iff there is some (arbitrary) function f from the Cantor set $\mathcal{C}$ into Y, such that for $x, y \in \mathcal{C}$, if $x \neq y$ then $f(x) \not\!E f(y)$; Y has perfectly many equivalence classes iff there is a *continuous* f as above. "Perfectly many", unlike "continuum many", is absolute whenever E and Y are absolutely - Δ_2^1 (which is the only situation we consider in this paper). Therefore, the size of $2^{\aleph_0}$ is irrelevant to the question of whether there are perfectly many equivalence classes.

Cambridge University Press
978-1-107-16790-2 — Logic Colloquium '95
Edited by Johann A. Makowsky , Elena V. Ravve
Excerpt
More Information

In this terminology, Theorem 2.1(b) (respectively, Theorem 2.2(ii)) states that K does not have (respectively, does have) perfectly many path-components.

In the next two theorems, we consider this property in the case $Y = X$, when E is $\boldsymbol{\Pi}^1_1$ (coanalytic) and when E is $\boldsymbol{\Sigma}^1_1$ (analytic), where E is regarded as a pointset in the space $X \times X$.

Theorem 3.1 (Silver). Let X be a Polish space and let E be a $\boldsymbol{\Pi}^1_1$ equivalence relation on X. One of the following two cases holds:

(i) X has countably many E-equivalence classes;
(ii) X has perfectly many E-equivalence classes.

This dichotomy theorem is not true for $\boldsymbol{\Sigma}^1_1$ equivalence relations. The following equivalence relation E^* on $\mathcal{C}$ is a counterexample:

$$xE^*y \Longleftrightarrow [(x \notin WO \text{ and } y \notin WO) \text{ or } |x| = |y|],$$

where WO denotes the set of ordinal codes and $|x|$ denotes the ordinal encoded by x. Clearly there are exactly $\aleph_1$ E^*-equivalence classes, and the Boundedness Theorem implies that (even if CH is true) there are not perfectly many classes.

Theorem 3.2 (Burgess). Let X be a Polish space and let E be a $\boldsymbol{\Sigma}^1_1$ equivalence relation on X. One of the following three cases holds:

(i) X has countably many E-equivalence classes;
(ii) X has $\aleph_1$ and not perfectly many E-equivalence classes;
(iii) X has perfectly many E-equivalence classes.

As shown above, case (ii) of 3.2 can occur. Thus $\boldsymbol{\Sigma}^1_1$ equivalence relations come in three types. Assuming $\neg CH$, the three types are just three cardinalities for the set of equivalence classes: $\aleph_0, \aleph_1, 2^{\aleph_0}$. But if CH is true, we need a different way of distinguishing case (ii) from case (iii), and that is where the concept "perfectly many" comes in.

The original proof of Theorem 3.1 appeared in Silver [18]. A simpler proof, essentially due to Harrington, can be found in Martin-Kechris [16]. The original proof of Theorem 3.2 is in Burgess [5]. Shelah later discovered an extremely general theorem, of which both 3.1 and 3.2 are special cases – this can be found in Harrington-Shelah [9].

Although case (iii) is absolute, the distinction between cases (i) and (ii) of Theorem 3.2 is not, in general, absolute. For it is provable in ZFC that there is a $\boldsymbol{\Sigma}^1_1$ equivalence relation E^{**} on $\mathcal{C}$ and a bijection between the equivalence classes of E^{**} and $\aleph_1^L$. (*Proof.* Let $C_1 = \{x : x \in L_{\omega_1^x}\}$ be the largest thin Π^1_1 set – see Kechris [12] for details. Then define

$$xE^{**}y \Longleftrightarrow [(x \notin C_1 \text{ and } y \notin C_1) \text{ or } x = y].$$

Since $card(C_1) = card(\aleph_1^L)$, this works.) On the other hand, for some $\boldsymbol{\Sigma}^1_1$ equivalence relations, such as E^*, case (ii) holds in *every* model.

Now consider those Polish spaces K which are compact subsets of $\mathbb{R}^n$, and the equivalence relation $\approx_K$ on K of being in the same path-component. Clearly

$\approx_K$ is $\boldsymbol{\Sigma}^1_1$, since

$$\mathbf{p} \approx_K \mathbf{q} \Longleftrightarrow (\exists \gamma \in (C[0,1])^n) F(\mathbf{p},\mathbf{q},\gamma), \tag{3.3}$$

where F is the following closed subspace of the Polish space $K \times K \times (C[0,1])^n$:

$$F(\mathbf{p},\mathbf{q},\gamma) \Longleftrightarrow [Im(\gamma) \subset K \text{ and } \gamma(0) = \mathbf{p} \text{ and } \gamma(1) = \mathbf{q}].$$

Therefore Burgess's Theorem is applicable to $\approx_K$, and so, as pointed out in the introduction, it is not possible that $2^{\aleph_0} \geq \aleph_3$ and K has $\aleph_2$ path-components.

The equivalence relation of being in the same *component* of K is closed, hence $\boldsymbol{\Pi}^1_1$, and therefore Silver's Theorem is applicable. That is, for any K,

(3.4) K has either countably many or perfectly many components.

So, as was also pointed out in the introduction, assuming $\neg CH$, K cannot have $\aleph_1$ components. These facts about components can be proved directly, without going through Silver's Theorem.

But is the $\boldsymbol{\Sigma}^1_1$ equivalence relation $\approx_K$ also $\boldsymbol{\Pi}^1_1$? Note that by Suslin's Theorem, it is $\boldsymbol{\Pi}^1_1$ iff it is Borel.

It has been known since the work of Kunen-Starbird [14] in 1982 that there exists a compact $K \subset \mathbb{R}^3$ for which $\approx_K$ is not Borel (and that therefore Silver's Thorem is not, in general, applicable to the equivalence relation $\approx_K$). It is still an open question whether or not for every compact $K \subset \mathbb{R}^2$, $\approx_K$ is Borel. While it is possible that for all $K \in \mathcal{K}(\mathbb{R}^2)$, $\approx_K$ is Borel, it is not the case that $\approx_K$ is Borel *uniformly* in K. For if it was, PC_2 would be a $\boldsymbol{\Pi}^1_1$ set, which is not true (see Becker [2, Theorem 2.2]).

This is the background which motivated the question posed at the beginning of this paper. (That question was asked by the author in 1984 in several talks and in the circulated notes Becker [1], but never asked in print.) To summarize: We have a collection $\mathcal{E} = \{\approx_K : K \in \mathcal{K}(\mathbb{R}^n)\}$ of $\boldsymbol{\Sigma}^1_1$, generally non-Borel, equivalence relations; Theorem 3.2 classifies $\boldsymbol{\Sigma}^1_1$ equivalence relations into three types, all of which can occur; the question is whether type (ii) ($\aleph_1$, not perfectly many) can occur for equivalence relations in $\mathcal{E}$. There are many interesting questions (some solved, some open) of precisely this form: Given a proper subclass of the class of all $\boldsymbol{\Sigma}^1_1$ equivalence relations, can type (ii) occur in this subclass? For example, Vaught's Conjecture is such a question, since isomorphism for countable structures – restricted to the Borel set of models of a first-order theory – is a $\boldsymbol{\Sigma}^1_1$ equivalence relation.

Remark. For the equivalence relation of isomorphism, the distinction between cases (i) and (ii) of Theorem 3.2 is absolute. Thus if there is a counterexample to Vaught's Conjecture in L it remains a counterexample in V (even if $\aleph_1^L < \aleph_1$). See Becker-Kechris [4, §7.2]. In this respect, there is a descriptive set theoretic difference between Vaught's Conjecture and the analogous conjecture for path-components with which this paper is concerned.

§4. Path-components in compact subsets of $\mathbb{R}^3$

The question, as posed in §3, was whether case (ii) of Theorem 3.2 – which does occur among arbitrary $\boldsymbol{\Sigma}_1^1$ equivalence relations – can occur for a special sort of $\boldsymbol{\Sigma}_1^1$ equivalence relation, those of the form $\approx_K$. Of course, Theorem 2.1 says that it does. The way 2.1 is proved is to show that equivalence relations of the form $\approx_K$ are really not all that special; *any* $\boldsymbol{\Sigma}_1^1$ equivalence relation can be coded up as $\approx_K$ for some $K \in \mathcal{K}(\mathbb{R}^3)$. This is made precise in Theorem 4.1, below.

Let $\mathcal{C}$ denote the Cantor middle third set in $[0,1]$.

Theorem 4.1 Let E be a $\boldsymbol{\Sigma}_1^1$ equivalence relation on $\mathcal{C}$. There exists a compact set $K_E \subset \mathbb{R}^3$ satisfying the following three properties.

(a) For all $x \in \mathbb{R}$, $(x,0,0) \in K_E$ iff $x \in \mathcal{C}$.
(b) For all $\mathbf{p} \in \mathbf{K_E}$ there exists an $x \in \mathcal{C}$ such that $(x,0,0) \approx_{K_E} \mathbf{p}$.
(c) For all $x, y \in \mathcal{C}$, xEy iff $(x,0,0) \approx_{K_E} (y,0,0)$.

Both a proof of Theorem 4.1 and a magnificent 3-dimensional picture of K_E will appear in Becker [3].

Note that if the word "compact" was removed from 4.1, the proof would be quite easy. For each pair (x,y) such that xEy, we could pick a path $\gamma^{(x,y)}$ connecting x and y, and since we are in 3-dimensional space, there is enough room to pick these paths so that no two intersect except at the endpoints; then let K_E be the union of all these paths. However a K_E constructed in this naive manner will not even be a Borel set. The trick is to get it to be compact. The construction of K_E is similar to the constructions in Kechris [13, Theorems 33.17 and 37.11].

Theorem 2.1 is a corollary of Theorem 4.1. To see this, just consider a $\boldsymbol{\Sigma}_1^1$ equivalence relation E on $\mathcal{C}$ with $\aleph_1$ and not perfect many equivalence classes, and let K_E be as in Theorem 4.1, for this particular E. It is not hard to show that K_E satisfies 2.1.

Kunen-Starbird [14] proved that there is a $K \in \mathcal{K}(\mathbb{R}^3)$ which has a non-Borel path-component, and asked: Does there exist a $K \in \mathcal{K}(\mathbb{R}^3)$ such that no path-component of K is Borel?

Corollary 4.2 There is a compact set $K \subset \mathbb{R}^3$ such that no path-component of K is Borel.

Proof. It is well known (but apparently unpublished) that there is a $\boldsymbol{\Sigma}_1^1$ equivalence relation E on $\mathcal{C}$ such that no E-equivalence class is Borel. (*Proof.* It will suffice to find such a $\boldsymbol{\Sigma}_1^1$ equivalence relation E' on a standard Borel space. Let S be a $\boldsymbol{\Sigma}_1^1$ non-Borel subset of $\mathbb{R}$, and let $F(\mathbb{R})$ and $F(S)$ be the free groups generated by $\mathbb{R}$ and S, respectively. Let E' be the equivalence relation on $F(\mathbb{R})$ given by the coset decomposition $F(\mathbb{R})/F(S)$.) Let K_E be as in Theorem 4.1, for this particular E. By 4.1(c), if any path-component of K_E was Borel, the corresponding E-equivalence class would be Borel.

Remark. In both 2.1 and 4.2, the K's can be taken to be connected (that is, to

Cambridge University Press
978-1-107-16790-2 — Logic Colloquium '95
Edited by Johann A. Makowsky , Elena V. Ravve
Excerpt
More Information

be *continua*). This is so because the components of the original K are compact and connected, so in 4.2, we can pass from K to any component, and in 2.1, to any component which consists of $\aleph_1$ path-components. Such a component must exist, by 3.4.

There are some very complicated $\boldsymbol{\Sigma}^1_1$ equivalence relations – complicated in both the intuitive sense, and in the precise sense of *definable cardinality*, as explained in Becker-Kechris [4, §8]. One example of a complicated $\boldsymbol{\Sigma}^1_1$ equivalence relation is Turing-equivalence. By 4.1, all this complexity exists in the path-component equivalence relation for compact subsets of $\mathbb{R}^3$.

All of the above results trivially transfer from $\mathbb{R}^3$ to $\mathbb{R}^n$, for $n \geq 3$. What about $n = 2$? Of course, the analog of Theorem 2.1 is false for $\mathbb{R}^2$ (assuming ϵ). The analog of Corollary 4.2 is also false for $\mathbb{R}^2$ (in ZFC); that is, for any compact $K \subset \mathbb{R}^2$, at least one path-component of K is a Borel set. These facts seem to mean that it is not possible to code up arbitrary $\boldsymbol{\Sigma}^1_1$ equivalence relations as the path-component equivalence relation for some $K \in \mathcal{K}(\mathbb{R}^2)$, under any conceivable meaning of "code up". This still leaves open the question of whether $\approx_K$ can ever be "complicated" for $K \in \mathcal{K}(\mathbb{R}^2)$, e.g., can it be as complicated as Turing-equivalence? There are no known examples (from any axioms) of a $K \in \mathcal{K}(\mathbb{R}^2)$ such that $\approx_K$ is not *smooth*, i.e., such that $\approx_K$ is more complicated than the equality relation on $\mathcal{C}$ (see Becker-Kechris [4, §3.4] for definitions and details).

§5. Descriptive set theory and equivalence relations, II: Stern's Theorem

In this section, we consider Borel equivalence relations, which are much better behaved than arbitrary $\boldsymbol{\Pi}^1_1$ equivalence relations. At first glance, Silver's Theorem (3.1) would seem to say that nothing could be better behaved than $\boldsymbol{\Pi}^1_1$ equivalence relations. The problem is that the Silver dichotomy for $\boldsymbol{\Pi}^1_1$ equivalence relations applies only to the *entire Polish space* X. If E is a $\boldsymbol{\Pi}^1_1$ equivalence relation on X, there may well be a simply definable – in fact, $\boldsymbol{\Pi}^1_1$ – E-invariant set $Y \subset X$ such that $E|(Y \times Y)$ does not have either countably many or perfectly many equivalence classes. For example, let E^{***} be the following $\boldsymbol{\Pi}^1_1$ equivalence relation on $\mathcal{C}$:

$$xE^{***}y \Longleftrightarrow [(x \in WO \text{ and } y \in WO \text{ and } |x| = |y|) \text{ or } x = y].$$

Clearly WO is $\boldsymbol{\Pi}^1_1$ and E^{***}-invariant, and $E^{***}|(WO \times WO)$ violates the dichotomy. For Borel equivalence relations, this situation does not occur.

Theorem 5.1 (Stern) Assume ϵ. Let X be a Polish space, let E be a Borel equivalence relation on X and let $Y \subset X$ be an E-invariant $\boldsymbol{\Sigma}^1_2$ set. One of the following two cases holds:

(i) Y has countably many E-equivalence classes;
(ii) Y has perfectly many E-equivalence classes.

Proof. See Stern [19].□

Cambridge University Press
978-1-107-16790-2 — Logic Colloquium '95
Edited by Johann A. Makowsky , Elena V. Ravve
Excerpt
More Information

To put Stern's Theorem in its proper context, the following two remarks may be helpful. First, fix a Borel equivalence relation E on X with perfectly many equivalence classes. Assuming the full axiom of determinacy (which contradicts the axiom of choice), *every* E-invariant set $Y \subset X$ has either countably many or perfectly many E-equivalence classes. This follows from Stern's Theorem together with a result of Harrington-Sami [8, Theorem 2]. Obviously, using the axiom of choice, we can pick out a set of $\aleph_1$ E-equivalence classes; and, in fact, even if CH is true, using choice we can get an E-invariant set $Y \subset X$ with uncountably many but not perfectly many equivalence classes. But such a Y will not be definable. Thus E-invariant sets $Y \subset X$ which violate the dichotomy are like sets of real numbers which are not Lebesgue measurable: Such pathological sets do exist, but one cannot explicitly define an example. That's not provable in ZFC, but all right-thinking people know it is true. Regarding provability, the analogy between sets $Y \subset X$ violating the dichotomy and nonmeasurable sets of reals still holds: Stronger and stronger large cardinal axioms imply larger and larger classes of sets are nonpathological. Stern's Theorem is that the axiom ϵ is sufficient to prove that $\boldsymbol{\Sigma}^1_2$ sets Y are nonpathological.

Second, consider the case where X is the reals and E is equality. For this special case, the conclusion of Theorem 5.1 is that for any $\boldsymbol{\Sigma}^1_2$ set $Y \subset \mathbb{R}$, either Y is countable or Y has a perfect subset. That is, the conclusion of 5.1 is literally the axiom ϵ. So clearly this assumption is necessary. Stern's Theorem says that *if* equality has this property, then every Borel equivalence relation has this property. And as shown by the examples E^* and E^{***}, above, "Borel" is best possible.

§6. Theta-curves

Definition. A *theta-curve* (in $\mathbb{R}^2$) is a 5-tuple $(\mathbf{u}, \mathbf{v}, \gamma_1, \gamma_2, \gamma_3)$ such that $\mathbf{u}, \mathbf{v} \in \mathbb{R}^2$, each γ_i is an arc from $\mathbf{u}$ to $\mathbf{v}$ in $\mathbb{R}^2$, and if $i \neq j$ then $\gamma_i \cap \gamma_j = \{\mathbf{u}, \mathbf{v}\}$.

We sometimes abuse the language and refer to the pointset $Im(\gamma_1) \cup Im(\gamma_2) \cup Im(\gamma_3)$ in $\mathbb{R}^2$ as the "theta-curve". Figure 2 is a picture of a theta-curve in this latter sense.

We need a theorem about the topology of the plane – the theorem says that the picture in Figure 2 is correct. It is actually a very deep theorem, and to motivate it one should first consider the famous Jordan Curve Theorem. A *circle* always means a topological circle. The Jordan Curve Theorem states: If C is any circle embedded in $\mathbb{R}^2$, then $\mathbb{R}^2 \backslash C$ has exactly two components; and furthermore, the boundary of each of the two components is C. There is a similar theorem for theta-curves.

Theorem 6.1 Let $(\mathbf{u}, \mathbf{v}, \gamma_1, \gamma_2, \gamma_3)$ be a theta-curve, and let $\tilde{\gamma}_i = Im(\gamma_i)$. $\mathbb{R}^2 \backslash (\tilde{\gamma}_1 \cup \tilde{\gamma}_2 \cup \tilde{\gamma}_3)$ has exactly three components. The boundary of one component is $\tilde{\gamma}_1 \cup \tilde{\gamma}_2$. The boundary of another component is $\tilde{\gamma}_2 \cup \tilde{\gamma}_3$. And the boundary of the third component is $\tilde{\gamma}_3 \cup \tilde{\gamma}_1$.

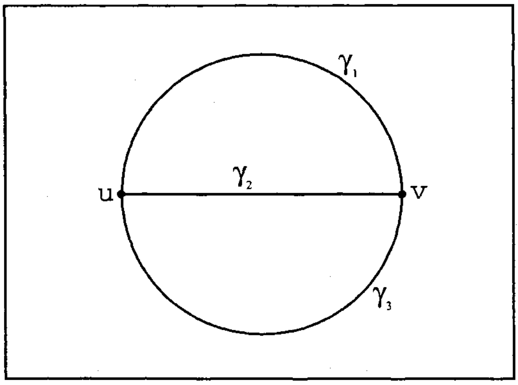

Figure 2

Proof. See Kuratowski [15, §61, II, Theorem 2]. □

Theorem 6.2 Let K be a compact subset of $\mathbb{R}^2$. If there is no theta-curve lying in K, then the equivalence relation $\approx_K$ is Borel.

Corollary 6.3 Assume ϵ. Let K be a compact subset of $\mathbb{R}^2$. If there is no theta-curve lying in K, then for any $\approx_K$-invariant $\boldsymbol{\Sigma}^1_2$ set $Y \subset K$, one of the following two cases holds:

(i) Y has countably many path-components;
(ii) Y has perfectly many path-components.

Proof. This follows from Theorems 5.1 and 6.2. Note that since Y is $\approx_K$-invariant, $\approx_Y$ is $\approx_K |(Y \times Y)$, i.e., every path-component of Y is also a path-component of K. □ In §7, we give a proof of Theorem 2.2 (from a stronger large cardinal axiom than ϵ). That proof uses both Theorem 6.1 and Corollary 6.3. We remark that one could also consider theta-curves in $\mathbb{R}^n$, for any n, and that both 6.2 and 6.3 would still be valid in the n-dimensional case. But the 3-dimensional analog of Theorem 6.1 is obviously false. Theorem 6.1 is the one and only place in the proof of Theorem 2.2 where the hypothesis that $K \subset \mathbb{R}^2$ is used.

The rest of §6 consists of a sketch of the proof of Theorem 6.2. This proof involves *effective* descriptive set theory, that is, recursion theoretic methods. Moschovakis [17] is the reference for this subject.

We work with recursively presented Polish spaces (as defined in Moschovakis [17, Page 128]). The Polish spaces $\mathbb{R}^2, \mathcal{K}(\mathbb{R}^2)$ and $(C[0,1])^2$ are all recursively presented, hence so are all finite products of these spaces. We regard compact subsets of $\mathbb{R}^2$ as points in the space $\mathcal{K}(\mathbb{R}^2)$, and we regard paths in $\mathbb{R}^2$ as points in the space $(C[0,1])^2$. For any recursively presented Polish spaces X and Y, and any points $x \in X$ and $y \in Y$, $x \leq_h y$ means that x is hyperarithmetic-in-y, or equivalently, that x is $\Delta^1_1(y)$. This is defined in Moschovakis [17, Pages 151 and 157].